A SEQUENTIAL STACKELBERG GAME FOR DYNAMIC INSPECTION PROBLEMS

Cristobal Guzman, Javiera Riffo, Claudio Telha, Mathieu Van Vyve

REPRINT | 3194

UCLouvain

CORE | LIDAM
Louvain Institute of Data Analysis and Modeling in economics and statistics

CORE

Voie du Roman Pays 34, L1.03.01

B-1348 Louvain-la-Neuve

Tel (32 10) 47 43 04

Email:lidam-library@uclouvain.be

https://uclouvain.be/en/research-institutes/lidam/core/core-reprints.html

A Sequential Stackelberg Game for Dynamic Inspection Problems [1]

Cristóbal Guzmán[a,1], Javiera Riffo[a,1], Claudio Telha[b,1,*], Mathieu Van Vyve[c,1]

[a] *Institute for Mathematical & Computational Engineering, Facultad de Matemáticas & Escuela de Ingeniería, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Macul, Santiago, Chile*

[b] *Faculty of Engineering and Applied Sciences, Universidad de los Andes, Chile. Monseñor Álvaro del Portillo 12455, Las Condes, Santiago, Chile*

[c] *Center for Operations Research and Econometrics (CORE), Université Catholique de Louvain. Voie du Roman Pays 34/L1.03.01, 1348 Louvain-la-Neuve, Belgium*

Abstract

We introduce an inspection game where one inspector has the role of monitoring a group of inspectees. The inspector has the resources to visit only a few of them. Visits are performed sequentially with no repetitions. The inspectees report and share the sequence of inspections as they occur, but otherwise, they do not cooperate. Our paper focuses on the mathematical structure of the equilibria of this sequential inspection game, where the inspector can perform exactly two visits. We formulate two Stackelberg models, a static game where the inspector commits to play a sequence of visits announced at the start of the game, and a dynamic game where the second visit will depend on who was visited previously.

In the static game, we characterize the (randomized) inspection paths in equilibrium using linear programs. In particular, these inspection paths are solutions to a transportation problem. We use this equivalence to determine an explicit solution to the game and to show that set of inspection path probabilities in equilibrium, projected onto its first and second visit marginals, is convex. In the dynamic game, we determine the inspection paths in equilibrium using backward induction. We discuss how the static and dynamic games relate to each other and how to use these models in practical settings.

*Corresponding author

[1] CG work was partially supported by grant FONDECYT 11160939

Keywords: Game Theory, Inspection Games, Sequential Stackelberg Games

1. Introduction

Inspecting a product or a service to ensure its quality is one of the fundamental tools to monitor the enforcement of an obligation. At the core, an inspection involves two parties: one called the inspector verifies that another, called the inspectee, adheres to specific rules. In a very common scenario, the inspector controls at a given moment that the inspectee is behaving legally in terms of an agreement, while the inspectee has an incentive to not doing so. A concrete scenario, briefly described in this paper, is the case of a restaurant franchising agreement, where the inspectee –a firm with rights to the franchise–could be interested in not adhering to expensive quality requirements of the franchise in order to increase output or to reduce costs. Besides quality control, inspections arise in contexts as diverse as accounting and auditing [10], arms control [44] and the enforcement of enviromental regulations [35].

An essential aspect of inspections is that they can be time-consuming, especially when they are manually performed. The resources available for inspections are limited and so are the opportunities to monitor that the agreement is being followed. A high sanction on an inspectee that is caught during an inspection, and an uncertainty on the time of the inspection are common ways to increase the effectiveness of these inspections [4].

The case of a single inspector and multiples inspectees arises in several contexts, such as police patrol and quality control. If inspections are time-consuming, then the inspector might only be able to check some of the inspectees. This leaves to the inspector not only the task to decide which inspectees should be visited, but also in which order they should be visited, in order to be as efficient as possible.

On the side of the inspectees, there is evidence that they can collaborate sharing non-sensitive information to avoid the consequences of inspections, even if they do not know each other. An example is the case of collaborative maps of police locations made by car drivers via smartphone applications. This information can be used by some drivers (the inspectees) to respect speed driving limits whenever they are near a police checkpoint [23].

The inspector and inspectees have objectives in conflict. If an inspection is made, then the visited inspectee would prefer to abide by the agreement, but in this case, the inspector would prefer not to do the inspection. Therefore, no decision is simultaneously optimal for both parties and it is expected for inspectees to react to the behavior of inspectors and viceversa. The general framework of inspection games [4] has been designed to capture the effects of these interactions and to determine a plausible outcome. A recurrent theme in inspection games is that because the inspector has limited inspection resources, they randomize inspections so that the inspectee sees no expected benefit in behaving illegally [29].

We propose and analyze a theoretical model of one inspector that has to decide how to sequentially inspect two out of n inspectees. Inspectees should perform a task that is costly for them. The task is short enough that it can be executed any time before the inspector shows up at the inspectee's door. The sequence of visits is shared in real-time by the inspectees themselves, who act strategically based on this information. No other information is shared among inspectees, who act otherwise without any coordination. From the perspective of the inspector, we aim to answer the following question: given the visits already performed, who should be the next inspectee to visit? From the perspective of each individual inspectee, we aim to answer the following question: given the visits already performed, should the inspectee make the effort and comply with the agreement? We formulate a game theoretical model to provide analytical answers to these questions.

Several characteristics of our game model mark a difference from classical inspection games. First, inspections are always be carried out but not every inspectee can be controlled. Thus, it is scarcity and not cost what limits inspecting resources. Second, inspectees share information about the route of the inspector, but act otherwise in their own interest. Third, the inspector cannot visit any inspectee twice. All of these features capture even further the asymmetry of the different parties involved in the inspection process.

We also consider the authority or reputation of the inspector and its impact on the game dynamics. More precisely, the inspector acts as the player leader who announces a credible (and randomized) strategy of sequential inspections. In response, each inspectee acts as a follower, best reacting to the strategy announced by the inspector. We propose two Stackelberg game models that differ in how the sequential aspect of the game is tackled

by the players; we call them the backward induction model and the linear programming model. In the backward induction model it is assumed that every potential future situation is played with rational strategies and any equilibrium is subgame perfect. In the linear programming model, players are rational only on paths that are realizable and any equilibrium is of the Nash type. As it is often observed in real life inspections, the inspector will randomize visits in order to induce the desired behavior on the inspectees as efficiently as possible.

Two characteristics of our game and study were shaped by the challenges posed by the problem. The game can be easily defined and extended for more than two stages, but new ideas are needed in order to generalize our results accordingly. The game could also be extended to more general profit functions, but the case we consider here seems to be much more tractable. We work on the *fine collector* variant, where the most profitable scenario for the inspector is to visit inspectees that are not prepared, so that a fine is collected.

2. Literature Review

The origins of the theoretical game models of inspection trace back to the 1960s [18, 7, 8]. Early on, most studies focused on arms control and disarmament [38]. Other applications include economics [45, 9, 42], environmental regulations [25, 43, 55, 24] and crime control [48, 5, 19, 21]. For an extended summary on the variety of applications of inspection games, we refer the reader to [3].

It is natural for inspector to realize multiple inspections over a time horizon. The seminal paper of Dresher [18] is one of the earliest examples of a n-stage, sequential inspection game. The inspector can perform m inspections in n stages, with $m < n$; the game is zero-sum, and it is solved recursively. Mascheler [36] generalized this model to the non-constant sum case. Several extensions followed [48, 5, 20].

The concept of leadership, where one player has the authority to announce a strategy, and then commit to it, was introduced by von Stackelberg [51]. Leadership in inspection games (inspector Stackelberg leadership) was first considered by Mascheler [36], and has been the subject of significant research [2, 52, 15, 53]. The leadership of a patrolling unit has also been studied in the closely related field of network security games [54, 49,

47, 56, 12, 11, 26, 17, 14], were practical implementations of these game-theoretical methods exist (inspection patrolling at Los Angeles International airport [41], inspection patrolling at port infrastructure [46], assignment of Federal Air Marshals to flights [28, 40, 27]).

Our paper focuses on sequential inspection games with leadership. Luh et al. [34] formalized sequential Stackelberg models, pointing out to the issue of temporal consistency in the play of the leader. They propose the notion of robust Stackelberg equilibrium as a solution for the game, emphasizing that other solutions are possible. The Strong Stackelberg Equilibrium SSE [32], is one of the most commonly adopted solution concepts in the literature of non zero-sum inspection games [13, 52, 31]. In the context of inspection games, a SSE is built when the inspector selects an optimal mixed strategy under the assumption that inspectees will choose an optimal response, breaking ties in favor of the inspector. Compared to other solution methods, the SSE has shown desirable properties, regarding existence (a SSE always exists, while robust Stackelberg equilibrium may not [6]) and computability (in two-player normal-form games, computing the Nash equilibrium that maximizes the utility of a player is $\mathcal{NP}$-hard [22, 16], but a SSE can be computed in polynomial time [15]).

Several papers consider fines or punishment in their models. Katiskas et al. [30] analyzed a two-player inspection game. The matrix payoff considers different profits for the inspector, depending of the action of the inspectee, together with punishment fines and inspection costs; in their model inspection always coincides with detection. Nosenzo et al. [37] studied the Nash equilibria of a simplified inspection game with bonuses and fines, concluding that bonuses are more effective that fines to induce the desired behavior. Other models with punishment have been considered in [1, 50]. Our work exhibits several differences with respect to these models, in particular: (i) we consider SSE; (ii) we study a dynamic game; and, (iii) in equilibrium no inspectee is prepared.

Another interesting work motivated by security games is Letchford and Conitzer [33], who investigated static combinatorial security games, where multiple inspectors are simultaneously assigned to multiples spots with the goal of finding an attacker. This problem is modeled as a combinatorial Stackelberg game for which compact linear program formulations are provided; these results are extensions of the Birkhoff-von Neumann Theorem.

Their model and techniques are quite different from ours.

3. The sequential inspection game

3.1. Preliminaries

From now on, we will use the term *operator* to denote an inspectee. We define a sequential, two-stage inspection game with one inspector I and a set V of $n > 2$ operators. The inspector I visits a pair of different targets and has a set of pure strategies $\{(u, v) \in V^2 : u \neq v\}$.

We assume that each operator $v \in V$ may decide at any stage of the game to *prepare*. If and when the operator prepares at some stage, then he or she complies with the inspection criteria and remains prepared for the rest of the game. The set of pure strategies of operator v is $\{0, 1\} \times \{0, 1\}$, where the pair of binary values indicates preparation —at value 1 —at each stage. Despite the notation, an operator who prepares at the first stage remains prepared; in this respect, we may assume the game ends for the operator.

An operator $v \in V$ who is inspected suffers a fine f_v if v is not prepared by the time of inspection. If the operator prepares at any stage, he or she incurs a cost d_v, with $0 < d_v < f_v$. Both events, the inspection of v and the preparation of v end the game for the operator immediately; therefore, the total cost for an operator can only take the values 0 (neither inspection nor preparation occurs by the second stage) d_v (preparation ended the game for v) and f_v (inspection ended the game for v). On the other hand, the inspector is a fine collector. Her total utility is the sum of the fines f_v collected from unprepared operators, during the two inspections.

These payoffs (utilities for the inspector, costs for the operators) induce a dynamics common to inspection games. At any given stage when an operator v is still in the game, the following occurs: if v is inspected, then v strictly prefers to prepare (since $d_v < f_v$); if v is not inspected, then v strictly prefers not to prepare (since $d_v > 0$). Conversely, if an operator v is prepared, then the inspector prefers not to inspect v (this is strictly, if at least one operator is not prepared); if v is not prepared, the inspector can inspect v for positive profit. Operator v will be the optimal choice of the inspector only when v gives the highest payoff f_v among those who have not yet prepared and have not been inspected. However, since the inspector and followers act simultaneously, the inspector must randomize in equilibrium.

3.1.1. Randomized strategies

We introduce some notation regarding mixed strategies. Mixed strategies for the inspector are random variables of the form $X = (X^{(1)}, X^{(2)})$. Mixed strategies for an operator v are random variables of the form $Q_v = (Q_v^{(1)}, Q_v^{(2)})$.

Let $p_{u,v} = \mathbb{P}(X^{(1)} = u, X^{(2)} = v)$ be the probability of the pure strategy (u, v) for the inspector. Let $p_{v|u} = p_{u,v} / \sum_{w \neq u} p_{u,w}$ be the conditional probability of inspecting v in the second stage given a visit to u in the first stage. Finally, let $p_{u\cdot} = \sum_{w \neq u} p_{u,w}$ and $p_{\cdot u} := \sum_{w \neq u} p_{w,u}$, be the marginal probabilities of visiting $u \in V$ at the first and second stage, respectively. For an arbitrary operator v, let $q_{v\cdot} = \mathbb{P}(Q_v^{(1)} = 1)$ be the probability of preparation in the first stage. For $v \neq u$ let $q_{v|u} = \mathbb{P}(Q_v^{(2)} = 1 \mid X^{(1)} = u, Q_v^{(1)} = 0)$ be the conditional probability of preparation in the second stage, given that v did not prepare earlier and that u was inspected in the first stage.

The expected utility function for the inspector and the expected cost function for operator v are denoted by $\mathbb{E}(U_I)$ and $\mathbb{E}(C_v)$, respectively.

3.1.2. Commitment power and information

All players have complete information and perfect recall. The game is Stackelberg. As the leader, the inspector anticipates the best response of the operators and announces her mixed strategy. The operators are followers who decide their strategies based on this announcement.

A subtle, but important, consideration is the issue of temporal consistency in this game. All players decide strategies, whose random actions are partially realized at the end of the first stage, but whether they need to respect their announced strategy after the first stage realization depends on the game model. We will say that the game is static if the inspector follows the conditional second stage strategy, even if it does not lead to a subgame equilibrium. On the other hand, if we require the inspector to play a subgame equilibrium then we will call the game dynamic. In either case, operators are always sequentially rational.

The solution concept used in this game is the Strong Stackelberg Equilibrium (SSE). The inspector will play a utility-maximizing strategy while operators will play the best response that maximizes the utility of the inspector. The latter is equivalent to say that operators will not prepare if they are indifferent on whether preparing or not.

The actions made by players at each stage are simultaneous. The game is of imperfect information, with one exception. In the second stage, all operators know the identity of the operator u visited on the first stage. Actions taken by the players at second stage are independent, conditionally on their previous (private) actions and u.

Remark 1. Our choice of payoffs in this paper was chosen because it promotes that the inspector makes sequence patrols only when operators do not prepare. On the other hand, notice that despite the fine transfer between operator and inspector, this game is not zero-sum, given the preparation costs for operators. This model better represents an "inspection agency," which does not directly profit from the operators preparations.

3.2. Equilibrium conditions in the static and dynamic models

Before we proceed to the analysis, we comment on certain simplifying assumptions and on certain aspects of the model. For convenience, we expand the set of operators to include a "dummy node" v', thus the set of operators in what follows is $V' := V \cup \{v'\}$. It is understood that visiting the dummy v' is equivalent to a no-inspection. For the dummy node, we let $f_{v'} = 0$ and $d_{v'} = \varepsilon > 0$ suitably small. In particular, the dummy node has the (strictly) lowest fine. The role of the dummy is to allow for the inspector to skip visits to induce some desired behavior on the operators. However, this is mostly an artifact: we will show at the end of this section that the dummy node is never used in equilibrium.

We will exclusively concentrate on the case $\sum_{v \in V} d_v / f_v \geq 2$. It turns out that this case corresponds to the situation where the inspection resources do not suffice to inspect all the operators within the two stages of the game effectively. Finally, unless it is explicitly indicated, we assume without loss of generality that operators in V' are sorted in decreasing order of $(f_v)_{v \in V}$.

In defining the equilibrium conditions, the side of the inspector plays a major role. Let us first compute the expected utility of the inspector. Conditioning on $X^{(1)} = u$, the expected fines collected by the inspector are $f_u(1 - q_{u\cdot})$ in the first stage and $\sum_{v \neq u} f_v \; p_{v|u}(1 - q_{v|u})(1 - q_{v\cdot})$ in the second stage. Recalling that for the dummy node, $f_{v'} = 0$, we have

$$\mathbb{E}(U_I) = \sum_{u \in V} p_{u\cdot} \left(f_u(1 - q_{u\cdot}) + \sum_{v \neq u} f_v \; p_{v|u}(1 - q_{v|u})(1 - q_{v\cdot}) \right). \qquad (1)$$

Let us also compute the expected cost of an operator $v \in V$ (the dummy node can be ignored, as it does not affect the utility of the inspector). Charges due to preparation or inspection of v can occur at either stage. In the first stage, the expected charge is $d_v q_{v\cdot} + f_v p_{v\cdot}(1 - q_{v\cdot})$. In second stage a charge can only occur when v does not prepare in the first stage and some other operator u is inspected in the first stage. Conditional on these two events (for a fixed u), the expected cost incurred by v in the second stage is $d_v q_{v|u} + f_v(1 - q_{v|u})p_{v|u}$. We obtain

$$\mathbb{E}(C_v) = d_v q_{v\cdot} + f_v p_v(1 - q_{v\cdot}) + (1 - q_{v\cdot}) \sum_{u \neq v} p_{u\cdot} \left(d_v \; q_{v|u} + f_v(1 - q_{v|u})p_{v|u} \right). \tag{2}$$

Let us now focus on determining the actions of the inspector and operator in the second stage of a SSE. Suppose that operator u is inspected in the first stage. For $v \neq u$ with $p_{u\cdot} > 0$ and $q_{v\cdot} < 1$, the inspector anticipates that if $p_{v|u}$ is announced, then v will optimize the second stage charge $d_v q_{v|u} + f_v(1 - q_{v|u})p_{v|u}$ by setting $q_{v|u} = \mathbb{1}_{\{p_{v|u} > d_v/f_v\}}$.

Playing $p_{v|u} > d_v/f_v$ to induce the response $q_{v|u} = 1$ is always suboptimal for the inspector. Indeed, any strategy S such that $p_{v|u} > d_v/f_v$ can be changed into a strategy identical to S, except that it sets $p_{v|u} = d_v/f_v$ and allocates the remaining probability to the dummy node. From Equation (1), the expected utility $\mathbb{E}(U_I)$ can only increase: changes to $q_{v|u}$ can only increase $\mathbb{E}(U_I)$; and $p_{v|u}$ decreased only in terms that used to contribute 0 (since $1 - q_{v|u}$ used to be 0).

Therefore, in a SSE we have that $p_{v|u} \leq d_v/f_v$ and $q_{v|u} = 0$ for all $v \neq u$ with $p_{u\cdot} > 0$ and $q_{v\cdot} < 1$. A careful use of this property in the second stage charge for operator v (the summation in (2)), simplifies this expression to:

$$\mathbb{E}(C_v) = d_v q_{v\cdot} + f_v(1 - q_{v\cdot}) \left(p_{v\cdot} + \sum_{u \neq v} p_{v|u} p_{u\cdot} \right). \tag{3}$$

With the relevant second stage probabilities determined, we proceed to compute the first stage probabilities. Suppose that v satisfies $p_{v\cdot} > 0$. In a SSE each operator v optimizes Equation (3) while minimizing $q_{v\cdot}$ if indifferent, in order to maximize the utility of the inspector. This is, $q_{v\cdot} = \mathbb{1}_{\{p_{v\cdot} + \sum_{u \neq v} p_{v|u} p_{u\cdot} > d_v/f_v\}}$. This means that v will prepare in the first

stage when the total probability of being inspected is greater than d_v/f_v. Once again, inducing the response $q_{v\cdot} = 1$ is always suboptimal for the inspector, since the total probability of inspection for v can be reduced by reallocation, increasing her expected utility.

Finally in the case of dynamic SSE, notice that the optimization performed by the inspector in the second stage applies to any operator v (and not only to those with $p_{v\cdot} > 0$). This follows from the subgame equilibrium property. In summary, we conclude the following.

Proposition 3.1. *In a SSE, the inspection probabilities satisfy*

$$\begin{aligned} p_{v|u} &\leq \frac{d_v}{f_v} \quad (\forall u \neq v \in V : p_{u\cdot} > 0) \\ p_{v\cdot} + \sum_{u \neq v} p_{v|u} p_{u\cdot} &\leq \frac{d_v}{f_v} \quad (\forall v \in V). \end{aligned}$$

Moreover, in a dynamic SSE the first set of conditions can be strengthened to

$$p_{v|u} \leq \frac{d_v}{f_v} \quad (\forall u \neq v \in V). \tag{4}$$

Remark 2. Let us note that from a static SSE equilibrium S we can construct an alternative equilibrium that also satisfies (4), and gives the same utility for the inspector as S. Just note that the utility of the inspector (1) is unaffected by modifying $p_{v|u}$ when $p_{u\cdot} = 0$.

Remark 3. Note that the constraints obtained in Proposition 3.1 are extendable to the dummy node, having in mind that due to our choice of parameters, $d_{v'}/f_{v'} = +\infty$.

We are now ready to introduce the two game models we consider in the paper.

3.3. The Linear Programming (LP) model

Proposition 3.1 provides necessary conditions for inducing no preparations from the operators, which by Remark 2 can be strengthened to (4). Given that these strategies are the only potential profit maximizers for the inspector, we have the following consequence.

Proposition 3.2. *Consider the linear program*

$$\text{(LP)} \qquad \max \sum_{u \in V} \sum_{v \neq u} \left(p_{u,v} + p_{v,u}\right) \cdot f_u$$

$$\text{s.t}$$

$$\sum_{v \neq u} (p_{u,v} + p_{v,u}) \leq \frac{d_u}{f_u} \quad (\forall u \in V') \tag{5}$$

$$p_{u,v} \leq \frac{d_v}{f_v} \left(\sum_{v \neq u} p_{u,v} \right) \quad (\forall u, v : u \neq v \in V') \tag{6}$$

$$\sum_{u \in V'} \sum_{v \neq u} p_{u,v} = 1 \tag{7}$$

$$p_{u,v} \geq 0 \quad (\forall u, v : u \neq v \in V').$$

Any optimal solution for (LP) is a static SSE.

3.3.1. Backward Induction model

Static models like (LP) may suffer problems of temporal consistency. Namely, after the first stage visit, the inspector may be incentivized to deviate from the conditional strategy announced at the beginning of the game. Depending on the commitment ability of the inspector, this may result in a non-credible strategy.

To resolve the problem of temporal consistency, we propose to compute a *dynamic* SSE, which can be computed by backward induction (BI). A subproblem ($BI^{(2)}(u)$) arises when the inspector visited u in the first stage, and the second stage is about to be played. Unlike the (LP) model, an equilibrium must be played on ($BI^{(2)}(u)$) even if $p_{u\cdot} = 0$, and this will incentivize the inspector to visit operators with a certain order preference. Let $k \in [n]$ be such that $\sum_{m \leq k} d_m/f_m \leq 1$ and $\sum_{m \leq k+1} d_m/f_m > 1$. For $u \in V'$, let $\kappa(u) \in [n] \setminus \{u\}$ be such that $\sum_{m \leq \kappa(u), m \neq u} d_m/f_m \leq 1$ and $\sum_{m \leq \kappa(u)+1, m \neq u} d_m/f_m > 1$.

To solve ($BI^{(2)}(u)$), we use the strengthened characterization of dynamic SSE from Proposition 3.1. In particular, if $q_{v\cdot} \neq 1$ (i.e., if operator v participates in the subgame) for all v, then the expected fines collected by the inspector in the second stage reduces to $\sum_{v \neq u} f_v \; p_{v|u}(1 - q_{v|u})(1 - q_{v\cdot}) =$

$(1 - q_{v\cdot}) \sum_{v \neq u} f_v \, p_{v|u}$, and $p_{v|u} \leq \frac{d_v}{f_v}$ must hold.

$$
\begin{aligned}
(BI^{(2)}(u)) \quad & \max \sum_{v \in V} p_{v|u} f_v \\
& s.t \qquad p_{v|u} \leq \frac{d_v}{f_v} \quad (\forall v \neq u \in V') \\
& \sum_{v \in V'} p_{v|u} = 1 \\
& \qquad p_{v|u} \geq 0 \quad (\forall v \neq u \in V').
\end{aligned}
$$

Model $(BI^{(2)}(u))$ can be optimized by a simple greedy procedure, which leads to the inspector to visit the most profitable operators in the second stage. Let $\left(p^*_{u|v}\right)_{u \neq v}$ be the optimal solution to $(BI^{(2)}(u))$. The proof of Proposition 3.3 and the discussion of the greedy procedure to solve $(BI^{(2)}(u))$ are deferred to Appendix A.1.

Proposition 3.3. *The solution of* $(BI^{(2)}(u))$ *satisfies*

$$
p^*_{v|u} = \begin{cases} \dfrac{d_v}{f_v} & \text{if } v \leq \kappa(u), v \neq u \\ 1 - \displaystyle\sum_{m \leq \kappa(u), m \neq u} \dfrac{d_m}{f_m} & \text{if } v = \kappa(u) + 1, v \neq u \\ 0 & \text{otherwise} \end{cases}
$$

The probabilities of inspection in the first stage of the backward induction model can be obtained by solving

$$
\begin{aligned}
(BI^{(1)}) \quad & \max \sum_{u \in V'} p_{u\cdot} \left(f_u + \sum_{v \neq u} p^*_{v|u} f_v \right) \\
& s.t \qquad p_{u\cdot} \leq \frac{d_u}{f_u} - \sum_{v \neq u} p^*_{u|v} p_{v\cdot} \quad (\forall u \in V') \qquad (8) \\
& \sum_{u \in V'} p_{u\cdot} = 1 \qquad (9) \\
& \qquad p_{u\cdot} \geq 0 \qquad (\forall u \in V').
\end{aligned}
$$

One could naively hope that this problem can be solved by another fractional knapsack (greedy) solution, unfortunately this is not the case: constraints (8) comprise a packing linear program, with nontrivial interactions between the optimization variables. However, the packing linear program $(BI^{(1)})$ is still solvable with a closed form solution. The proof of this result is deferred

to Appendix A.2.

Proposition 3.4. *Suppose* $\sum_{v\in V} d_v/f_v = 2$*, and let* $\gamma := 1 - \sum_{j=1}^{k} d_j/f_j$*. Let* $(p^*_{v|u})_v$ *be an optimal solution of program* $(BI^{(2)}(u))$ *for each* u*, then the following is an optimal solution for program* $(BI^{(1)})$*.*

$$p_{u\cdot} = \begin{cases} 0 & \text{if} \quad u \leq k \\ \dfrac{d_{k+1}/f_{k+1} - \gamma}{1-\gamma} & \text{if} \quad u = k+1 \\ \dfrac{d_u}{f_u} - p^*_{u|k+1}\dfrac{d_{k+1}/f_{k+1} - \gamma}{1-\gamma} & \text{if} \quad u \geq k+2. \end{cases}$$

Notice in the result above that in the case $\gamma = 0$, the optimal solution of the dynamic SSE splits into two fractional knapsack problems: one for the second stage, involving the most profitable nodes $1, \ldots, k$; followed by the subsequent most profitable nodes $k+1, \ldots, n$, which are used in first stage.

Remark 4. If $\sum_{v\in V} \frac{d_v}{f_v} > 2$, then some operators may not be inspected. Let h be the integer such that: $\sum_{v<h} \frac{d_v}{f_v} < 2 \wedge \sum_{v\leq h} \frac{d_v}{f_v} \geq 2$. The probabilities of inspection given in Proposition 3.4 hold provided that we redefine $V = [h]$ and $d_h/f_h \leftarrow \left(2 - \sum_{m<h} d_m/f_m\right)$. Every operator $v > h$ is not inspected.

3.4. Simple numerical example

In order to illustrate the differences between both models, consider an example with a set $V = [4]$ of four operators having $d_v/f_v = 1/2$ for each $v \in V$ and $f_1 > f_2 > f_3 > f_4 > 0$.

For (BI), the optimal equilibrium using backward induction can be obtained as follows. Conditioned on a visit to u in the first stage, the inspector solves $(BI^{(2)}(u))$ by randomizing equally among the two operators $v \neq u$ with the highest values of f_v. In the first stage, the inspector randomizes equally among the operators 3 and 4, as they have the lowest fines f. This is, $p_{u\cdot} = (1/2)\mathbb{1}_{\{3\leq u\leq 4\}}$. The (BI) solution is unique.

On the other hand, one optimal solution to the (LP) model chooses the sequences $(1,3), (1,4), (2,1)$ and $(2,4)$ with equal probability. This strategy satisfies that in the first stage the inspector randomizes equally among operators 1 and 2.

In this example, the solution to (BI) is also a solution to (LP), but not vice-versa. The (LP) solution suffers from temporal inconsistency: if operator $u = 1$ is selected in the first stage, then the inspector would be

Table 1: Conditionals in the (BI) model.

$u \backslash v$	1	2	3	4
1	0.0	0.5	0.5	0.0
2	0.5	0.0	0.5	0.0
3	0.5	0.5	0.0	0.0
4	0.5	0.5	0.0	0.0

Table 2: Conditionals in the (LP) model.

$u \backslash v$	1	2	3	4
1	0.0	0.0	0.5	0.5
2	0.0	0.0	0.5	0.5
3				
4				

Note: The (u, v)-entry of the table corresponds to $p_{v|u}$ in a solution for each model. Empty entries correspond to undefined probabilities.

better off by visiting $v = 2$ in the second stage, however, she is not allowed to.

3.5. Reduced formulation

The (LP) model for the inspection problem involve $O(n^2)$ variables. One can naturally consider a relaxation of the (LP) model, with only $2n$ variables corresponding to the marginal probabilities of inspection at each stage.

$$
\begin{aligned}
(\mathrm{LP}') \qquad \max \quad & \sum_{u \in V} (p_{u\cdot} + p_{\cdot u}) f_u \\
s.t \quad & \sum_{u \in V'} p_{u\cdot} = \sum_{u \in V'} p_{\cdot u} = 1 \\
& p_{u\cdot} + p_{\cdot u} \leq \frac{d_u}{f_u} \quad (\forall u \in V') \\
& p_{u\cdot}, p_{\cdot u} \geq 0 \quad (\forall u \in V').
\end{aligned}
$$

Any solution $(p_{u,v})_{\{u,v \in V : u \neq v\}}$ to (LP) can be converted into a solution to (LP$'$) with the same objective value by computing their marginals. Further, any dynamic SSE solution, such as those obtained by (BI), are feasible solutions for (LP), and thus their marginals are feasible for (LP$'$).

We now show that under certain conditions it is possible to construct optimal strategies for the (BI) and (LP) models, only starting from an optimal solution for (LP$'$). In practice, this allows for a decoupling of the problem in two steps: first, solving the small (LP$'$) model (whose solution can be computed in linear time by a greedy algorithm), and then using a flow algorithm to quickly build an optimal strategy.

The reduced formulation allows us to find explicit solutions which are

needed in the remaining of the section. Note that we can further relax (LP′) using variables $r_u = p_{u\cdot} + p_{\cdot u}$, obtaining

$$\text{(RedLP)} \quad \max\left\{\sum_{u\in V} r_u f_u : \sum_u r_u \leq 2,\, 0 \leq r_u \leq \frac{d_u}{f_u} \quad (\forall u \in V')\right\}. \quad (10)$$

This fractional knapsack problem has an explicit optimal solution. For instance, if $\sum_{u\in V} d_u/f_u = 2$, then the optimum is $\sum_{u\in V} d_u$. In general, the optimal solution to (RedLP) can be turned into a explicit feasible solution $(p_{u\cdot}, p_{\cdot u})$ for (LP′) in different ways. For example, one can set $p_{u\cdot} = p_{\cdot u} = r_u/2$, for all $u \in V'$. This solution turns out to be consistent with (LP), i.e. we will show there exists an (LP) solution whose marginals coincide with the ones above. A second example is to set $p_{u\cdot}, p_{\cdot u}$ as in Proposition 3.4

Remark 5. From Proposition 3.2 and 3.4 we conclude that both the static and dynamic models are such that the set of inspected operators, $V_+ := \{v \in V' : p_{v\cdot} + p_{\cdot v} > 0\}$, is given by $V_+ = \{1, \ldots, \overline{n}\}$, with $\sum_{v<\overline{n}} d_v/f_v < 2$ and $\sum_{v\leq\overline{n}} d_v/f_v \geq 2$. Since we assumed that $\sum_{v\in V} d_v/f_v \geq 2$, we conclude that the dummy node v' is never used in an equilibrium.

Remark 6. The case $\sum_{v\in V} d_v/f_v < 2$ is not analyzed in the paper for the sake of the exposition. To extend the results in this paper to this case we need to allow the inspector to skip one or both inspections. Up to two dummies may be needed to achieve this goal, and the inspector will visit at least one dummy with positive probability.

4. Results

Next we show how to construct mixed strategies from marginals, by two different constructions, which are both inspired by interpreting mixed strategies as flows satisfying supply and demand constraints, given by the marginals.

4.1. Constructing strategies from total probability of inspection via flows

In this section we introduce a transportation problem, whose feasible flows encode assignments of joint probabilities $(p_{u,v})_{u\neq v}$ whose marginals coincide with given probabilities $(p_{u\cdot})_{u\in V}$ and $(p_{\cdot u})_{u\in V}$. This construction will provide a flow algorithm to convert solutions of (RedLP) into solutions of (LP) or (BI).

We start by providing a general description of the transportation model $T = (G(U_1 \cup U_2), c, g)$. Let $G(U_1 \cup U_2)$ be a complete bipartite network on the set of nodes $U_1 \cup U_2$. Each $u \in U_1$ is a source node with g_u units of supply. Each $v \in U_2$ is a demand node with g_v units of demand. The capacity of each arc (u, v) in $U_1 \times U_2$ is $c_{u,v}$. A flow on T is feasible if it satisfies supply and demands and respects the capacities on the arcs.

For $S \subseteq U_1 \cup U_2$, let $N(S)$ denote the set of neighbors of S in $G(U_1 \cup U_2)$. From a theorem of Ore [39] on the existence of $f-$factors on bipartite multigraphs one can derive[3] the following condition for the existence feasible flows

$$T \text{ has a feasible flow} \iff \forall S \subseteq U_1 : \sum_{u \in S} g_u \leq \sum_{v \in N(S)} \min \left\{ g_v, \sum_{w \in S} c_{w,v} \right\}. \tag{11}$$

Let us use this transportation model to find solutions to the inspection game. We consider the set of inspected operators $V_+ = \{1, \ldots, \overline{n}\}$, as in Remark 5 (notice other operators may be ignored, as they are never considered in an equilibrium). In this case, we have $\sum_{v \in V_+} d_v/f_v \geq 2$: it is convenient for now to think of this sum to be exactly 2, but for the proof this assumption will not be necessary.

With the assumption $\sum_{v \in V_+} d_v/f_v \geq 2$, the optimal solution for (RedLP) satisfies $r_v = d_v/f_v$, for all $v \in V_+$. Given a set of marginals $p_{\cdot u}, p_{u\cdot}$ with $p_{\cdot u} + p_{u\cdot} = r_u$, we define a transportation model T on a bipartition $U_1 \cup U_2$, where U_1 and U_2 are two distinguishable copies of V_+. The marginals of each stage will be used as the supply/demand on the nodes, this is, $g_u = p_{u\cdot}$ for each $u \in U_1$, and $g_v = p_{\cdot v}$ for each $v \in U_2$. We write $T = T(\{p_{u\cdot}\}_{u \in V_+}, \{p_{\cdot u}\}_{u \in V_+})$, to emphasize the dependency on the chosen marginals. The nodes in U_1 and U_2 represent potential inspections in the first and second stage, respectively. The capacity $c_{u,v}$ on arcs $u \neq v$ are set to $(d_v/f_v)p_{u\cdot}$, otherwise they are set to zero. Figure 1 illustrates the construction.

A feasible flow $(p_{u,v})_{u \in V_+, v \in V_+}$ in $T(\{p_{u\cdot}\}_{u \in V_+}, \{p_{\cdot u}\}_{u \in V_+})$ is a set of values that can be interpreted as probabilities of pairs of inspections, whose marginals coincide with given probabilities $(p_{u\cdot})_{u \in V_+}$ and $(p_{\cdot u})_{u \in V_+}$. More-

[3]Technically, the Equation (11) follows directly from [39] when c and g are integer (or rational) valued. However, the proof in [39] is based on duality and does not depend on integrality.

Figure 1: Capacitated graph.

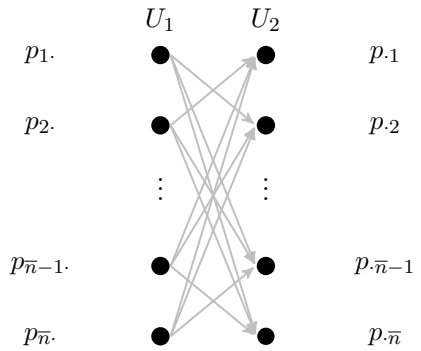

over, the capacity constraints ensure that constraint (6) is satisfied and that repeated visits do not occur. Thus, this defines a feasible and optimal solution for (LP).

The following is a structural property regarding the existence of solutions to (LP), written in terms of the equivalent transportation model.

Theorem 4.1. *Fix an instance of the sequential inspection game. Let* $(r_v)_{v\in V_+}$ *be the optimal solution of (RedLP), given in* (10)*. Then, the set* $M = \left\{\{p_{u\cdot}, p_{\cdot u}\}_{u\in V_+} : \text{ the problem } T\left(\{p_{u\cdot}\}_{u\in V_+}, \{p_{\cdot u}\}_{u\in V_+}\right) \text{ has a solution}\right\}$ *is convex.*

Proof. For $i = 0, 1$, let $\left\{p^i_{u\cdot}, p^i_{\cdot u}\right\}_{u\in V_+}$ be two marginals in M. For $\lambda \in [0, 1]$, let $\{p^\lambda_{u\cdot}, p^\lambda_{\cdot u}\}_{u\in V_+}$ be the convex combination defined by $p^\lambda_{u\cdot} = (1-\lambda)p^0_{u\cdot} + \lambda p^1_{u\cdot}$ and $p^\lambda_{\cdot u} = \lambda(1-\lambda)p^0_{\cdot u} + p^1_{\cdot u}$, for $u \in V_+$.

Equation (11) implies

$$\forall i \in \{0, 1\}, \forall S \subseteq U_1 : \sum_{u\in S} p^i_{u\cdot} \leq \sum_{v\in N(S)} \min\left\{p^i_{\cdot v}, r_v \sum_{w\in S, w\neq v} p^i_{w\cdot}\right\}.$$

For fixed $S \subseteq U_1$, we now combine these two inequalities (for $i = 0, 1$) to derive the same property for any $i = \lambda \in [0, 1]$, and this will conclude the proof (by (11)). For all $S \subseteq U_1$ we bound $\sum_{u\in S} \lambda p^1_{u\cdot} + (1-\lambda)p^0_{u\cdot}$ by

$$\sum_{v\in N(S)} \left[\min\left\{\lambda p^1_{\cdot v},\ r_v \sum_{w\in S, w\neq v} \lambda p^1_{w\cdot}\right\} + \min\left\{(1-\lambda)p^0_{\cdot v},\ r_v \sum_{w\in S, w\neq v} (1-\lambda)p^0_{w\cdot}\right\}\right]$$

which in turn it is upper bounded by

$$\sum_{v\in N(S)} \min\left\{\lambda p^1_{\cdot v} + (1-\lambda)p^0_{\cdot v},\ r_v \sum_{w\in S, w\neq v} \left(\lambda p^1_{w\cdot} + (1-\lambda)p^0_{w\cdot}\right)\right\},$$

as desired.

□

Theorem 4.1 implies that the set of marginals for which we can build a solution to (LP) is convex. A more constructive proof of Theorem 4.1 exists: the convex combination of the solutions to $T\left(\{p^i_{u\cdot}\}_{u\in V_+}, \{p^i_{\cdot u}\}_{u\in V_+}\right)$, $i \in \{0,1\}$ can be shown to be a solution of $T\left(\{p^\lambda_{u\cdot}\}_{u\in V_+}, \{p^\lambda_{\cdot u}\}_{u\in V_+}\right)$ on the corresponding convex combination of the marginals.

In the following result, we show at least two different marginals for which a solution to (LP) can be constructed. One of them can be seen as a non-constructive proof of Theorem 4.4 that will be described in Section 4.3.

Theorem 4.2. *Fix an instance of the sequential inspection game. Let $(r_v)_{v\in V_+}$ be the optimal solution of (RedLP), given in* (10)*, and consider the following choices of marginal probabilities:*

(i) $p_{u\cdot} = p_{\cdot u} = r_u/2$, for all $u \in V_+$.
(ii) $\{p_{u\cdot}, p_{\cdot u}\}_{u\in V_+}$ as in Proposition 3.4.

Then there exists a joint probability distribution $(p_{u,v})_{u\neq v}$ whose marginals coincide with $\{p_{u\cdot}, p_{\cdot u}\}_{u\in V_+}$. Moreover, one such solution can be computed from the capacitated transportation model $T = T(\{p_{u\cdot}\}_{u\in V_+}, \{p_{\cdot u}\}_{u\in V_+})$.

Proof. Case (ii) follows from the fact that problem (BI) has a solution, and it is given explicitly by Proposition 3.4. We should note, however, that T has other solutions which may not be compatible with backward induction.

For Case (i), we will show that the condition for the existence of solutions stated in (11) is satisfied in T. Before doing this, we slightly modify the capacities in the transportation graph. Notice that by assumption $r_v = d_v/f_v$ for all $v \in V_+ \setminus \{\overline{n}\}$, and that $r_{\overline{n}} \leq d_{\overline{n}}/f_{\overline{n}}$. If $r_{\overline{n}} < d_{\overline{n}}/f_{\overline{n}}$, then we will reduce the capacities of arcs $c_{u,\overline{n}} = (d_{\overline{n}}/f_{\overline{n}})p_{u\cdot} \to r_{\overline{n}}p_{u\cdot}$. If we show that this modified instance has a feasible flow, then clearly the starting instance has one as well. Finally, recall that by the optimality conditions of (RedLP), $\sum_{v\in V_+} r_v = 2$.

We now proceed to verify that the condition on the right side of (11) holds. First, we consider $S = \{u\} \subseteq U_1$. Then, $N(S) = U_2 \setminus \{u\}$ and therefore

$$\sum_{v\in N(S)} \min\{g(v), \sum_{w\in S} c_{w,v}\} = \sum_{v\neq u} \min\{p_{\cdot v},\ p_{u\cdot} r_v\} = \frac{1}{2}\sum_{v\neq u} r_v \min\{1, r_u\}.$$

Since $r_u \leq 1$, the last term is equal to $\frac{r_u}{2}\sum_{v\neq u} r_v = \frac{r_u}{2}\big(2-r_u\big) \geq \frac{r_u}{2} = g(u)$.

Next we consider the case $|S| > 1$. Then, $N(S) = U_2$, so re-arranging terms, we have that (11) is equivalent to

$$\sum_{v\in V_+} r_v \min\left\{1, \sum_{w\in S\setminus\{v\}} r_w\right\} - \sum_{u\in S} r_u \geq 0. \tag{12}$$

Let $B = \{v \in V_+ : \sum_{w\in S\setminus\{v\}} r_w < 1\}$. Let also $b = \sum_{v\in B} r_v$ and $s = \sum_{v\in S} r_v$. We can decompose the expression on the left-hand side of (12), depending on whether $v \in B$, obtaining:

$$\sum_{v\in V_+} r_v \min\left\{1, \sum_{w\in S\setminus\{v\}} r_w\right\} - \sum_{u\in S} r_u = \sum_{v\in V_+\setminus B} r_v + \sum_{v\in B} r_v\left(s - r_v \mathbb{1}_{\{v\in S\}}\right) - s \tag{13}$$

$$= 2 - b - s + sb - \sum_{v\in S\cap B} r_v^2. \tag{14}$$

To conclude we will prove that (14) is non-negative:

- If $s \leq 1$ or $b \leq 1$: Using that $r_v^2 \leq r_v$ for all $v \in V_+$, we obtain $\sum_{v\in B\cap S} r_v^2 \leq \min\{s, b\}$. Therefore,

  $$2 - b - s + sb - \sum_{v\in S\cap B} r_v^2 \geq \max\{(2-b)(1-s), (1-b)(2-s)\} \geq 0,$$

  where we used that both b and s are no larger than 2.

- If $s > 1$ and $b > 1$ then $B \subseteq S$, by definition of B (if there is $v \in B \setminus S$ then $s < 1$) and therefore $b \leq s$. Moreover, by minimizing (14) with respect to s we obtain:

  $$2 - b - s + sb - \sum_{v\in B} r_v^2 \geq 2 - 2b + b^2 - \sum_{v\in B} r_v^2 = 1 + (b-1)^2 - \sum_{v\in B} r_v^2.$$

  To conclude, note that if $\sum_{v\in B} r_v = b$ is constant in the right hand side, then $\sum_{v\in B} r_v^2$ is maximum when $r_1 = 1, r_2 = b-1$ and the other terms are zero. In this case, $1 + (b-1)^2 - \sum_{v\in B} r_v^2 = 0$, as needed.

□

Combining Theorem 4.2 with Theorem 4.1 we obtain an infinite set of marginals for which we can build a solution to (LP). On the other hand, it

is easy to build marginals for which (LP) has no solution (see example in Appendix B).

4.2. Comparing (BI) and (LP)

To solve the SSE we have proposed two different models, (BI) and (LP), based on different principles. On the one hand, (LP) addresses a static model for inspections, given by profit maximization over pairs of visits, however it ignores issues of temporal consistency that may be faced by the inspector at the second stage; this results in a linear program in dimension $O(n^2)$. On the other hand, (BI) addresses temporal consistency (the inspector is not incentivized to deviate from her strategy after the first stage), at the cost of solving n linear programs of dimension $O(n)$.

We will now compare both models in terms of their resulting profit. Let z_{LP} (resp. z_{BI}) be the optimal value of the (LP) (resp. (BI)) model.

Theorem 4.3. *The* (BI) *and* (LP) *models have the same optimal value.*

Proof. Suppose $\sum_{v\in V} d_v/f_v = 2$. From Proposition 3.4, the values of the marginals in the optimal (BI) solution satisfy $p_{u\cdot} + p_{\cdot u} = \frac{d_u}{f_u}$ for all $u \leq j$. Otherwise, they take value 0. Therefore, $z_{BI} = \sum_{u\in V'}(p_{u\cdot} + p_{\cdot u})f_u = \sum_{u:u\in V} d_u =: z$.

We prove that z_{LP} is both lower and upper bounded by z. Since feasible solutions for (BI) are feasible for (LP), $z \leq z_{LP}$. On the other hand, we showed in Section 3.5 that the reduced and relaxed model (RedLP) in Equation (10) has optimal value z and it is an upper bound on z_{LP}. Therefore $z_{LP} \leq z$.

If $\sum_{v\in V} \frac{d_v}{f_v} > 2$, then some operators may not be inspected. Let h be the integer such that: $\sum_{v<h} \frac{d_v}{f_v} < 2 \wedge \sum_{v\leq h} \frac{d_v}{f_v} \geq 2$. Let us redefine $V = [h]$ and

$$(d_h, f_h) \leftarrow \left(\left(2 - \sum_{m<h} d_m/f_m\right) f_h, f_h\right).$$

The greedy solution to (RedLP) does not change, except that $r_h = 2 - \sum_{m<h} d_m/f_m$ and $r_v = 0$ for $v > h$. Therefore $r_h f_h = d_h$, the contribution of h to the objective value in (RedLP), is correctly captured by z.

□

4.3. An explicit equilibrium

In this section we go one step further and provide one explicit solution $(p_{u,v})_{\{u,v\in V:u\neq v\}}$ for model (LP). With these explicit probabilities, the inspector only needs to compute the solution along the sequence which is effectively played. Some of the terms $p_{u,v}$ will require the computation of a series.

However, the rate of convergence is at least exponential, in practice reaching computer precision quite quickly. Operators will not prepare in equilibrium. Throughout this section we focus on the especial case $\sum_u \alpha_u = 1$, where $\alpha_u = \frac{d_u}{2f_u} < 1/2$. The construction of the explicit solution can be reduced to this especial case and this is discussed in Appendix C.

We first provide some intuition behind the construction of this explicit solution. The choice of independent inspection visits with probabilities $p_{u,v} = \alpha_u \alpha_v$ for all $u, v \in V$ "almost" works as a solution to (LP), fulfilling all the restrictions if repeated visits (u, u) were allowed. A nice feature of this "solution" is that it has significant slack in the second-stage probability constraints (6), but we failed to find a method to distribute the probability of repeated visits $p_{u,u} = \alpha_u^2$ to valid inspection paths without violating other constraints. Instead, our construction emulates the properties of this independent solution, avoiding the repetition of visits altogether using an elaborated recursive procedure.

We anticipate several key aspects behind the construction of the solution. The probabilities will be constructed to be symmetric, satisfying $p_{u,v} = p_{v,u}$ for all $u \neq v$. It is convenient to think of $p_{u,v}$ as flow variables on a bipartite graph, imitating what we did in Section 3.5, Figure 1. We set from the start $p_{u\cdot} = \alpha_u$ for all u; this will be the supply amount of node u on one side of the bipartition. It follows that $p_{\cdot u} = \alpha_u$ after symmetry is checked; this will be the demand amount of node u on the other side of the bipartition.

The procedure will build the flow $p_{u,v}$ incrementally, starting from $p_{u,v} \equiv 0$, and eventually satisfying supply and demands at every node. Whenever the flow $p_{u,v}$ is incremented, the flow $p_{v,u}$ will be incremented by the same amount. Therefore the procedure creates partial flows that drain the supply and demand nodes symmetrically. The relative order of the values α_u is important, it is assumed that $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$ are ordered decreasingly.

The procedure has two major steps. In the first step, called the *correcting* step, some flow will be routed in very specific quantities. For the moment, what is relevant is that the supply and demand not yet served by this flow will decrease from the original values $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$ to new values $\alpha'_1 = \alpha'_2 \geq \alpha'_3 \geq \ldots \geq \alpha'_n$, where not only the relative order is preserved, but also the two largest supply values $\alpha'_1 = \alpha'_2$ will be equal. In the second step, called the *recursive step*, we imitate the characteristics of the independent solution, routing flow proportional to $\alpha'_u \alpha'_v$ on every arc (u, v) with $u \neq v$. The supply

and demand that remains to be sent at the nodes will reduce to their squares $\alpha_1'^2 = \alpha_2'^2 \geq \alpha_3'^2 \geq \ldots \geq \alpha_n'^2$, and this is then sent recursively. A careful analysis is needed to show that the aggregated flows obtained through these two steps satisfy the second stage conditional probability constraints (6).

Correcting step. We assume that $\alpha_1 > \alpha_2$ as otherwise we proceed directly to the recursive step. The correction consists in sending for some $\beta \geq 0$ to be defined, $f_{1v}^{\text{cor}} = f_{v1}^{\text{cor}} = \beta\alpha_1\alpha_v$ units of flow in the arcs $(1, v)$ and $(v, 1)$, this for every other node $v \neq 1$. The supply yet to be served by node $v \neq 1$ reduces to $\alpha_v(1 - \beta\alpha_1)$, while the the supply yet to be served by node 1 reduces to $\alpha_1(1 - \beta(1 - \alpha_1))$. As β increases, note that all nodes $v \neq 1$ decrease at the same rate, while node 1 decreases at a higher rate, since $1 - \alpha_1 > 1/2 > \alpha_1$. In particular, the relative order of the remaining supplies will be preserved until β is large enough for nodes 1 and 2 to have the same remaining supply.

Equating $\alpha_1(1 - \beta(1 - \alpha_1)) = \alpha_2(1 - \beta\alpha_1)$ we determine the value

$$\beta = \frac{\alpha_1 - \alpha_2}{\alpha_1(1 - \alpha_1 - \alpha_2)}$$

at which the remaining supply at nodes 1 and 2 coincide. It is easy to check that $\beta \in (0, 1/(1 - \alpha_1)]$ when α_2 ranges in $[0, \alpha_1)$. At this value of β, the remaining supply is $\alpha_1' = \alpha_2' \geq \alpha_3' \geq \ldots \geq \alpha_n'$. Proceed to the next step.

Recursive step. Misusing the notation, we now assume that $\alpha_1 = \alpha_2 \geq \alpha_3 \geq \ldots \geq \alpha_n$ is the supply to satisfy, with $\sum_w \alpha_w$ not necessarily equal to 1. In this step we send $f_{uv} = f_{vu} = \frac{\alpha_u\alpha_v}{\sum_w \alpha_w}$ units of flow on the arcs (u, v) and (v, u), this for every pair $u \neq v$ of nodes. Node u sends a total of $\left(\sum_{v \neq u} \alpha_u\alpha_v\right) / \sum_w \alpha_w$ units of flow. We will then send the remaining supply at node u recursively.

The remaining supply at node u by the end of the first step of the recursion is just $\frac{\alpha_i^2}{\sum_k \alpha_k}$. At the end of the t-th step of the recursion it is

$$\frac{\alpha_u^{2^t}}{\sum_w \alpha_w \sum_w \alpha_w^2 \ldots \sum_w \alpha_w^{2^{t-1}}}.$$

When $\alpha_1 = \alpha_2$, we can twice bound this quantity by $\frac{\alpha_i^{2^t}}{2\alpha_1 \cdot 2\alpha_1^2 \cdot \ldots \cdot 2\alpha_1^{2^{t-1}}} \leq \frac{\alpha_i}{2^{t-1}}$, bounds that converge to 0 for large t.

The total flow f_{uv}^{rec} sent in the arc (u, v) throughout the recursion is.

$$f_{uv}^{\text{rec}} = \frac{\alpha_u \alpha_v}{\sum_w \alpha_w} + \frac{\alpha_u^2 \alpha_v^2}{\sum_w \alpha_w \sum_w \alpha_w^2} + \frac{\alpha_u^4 \alpha_v^4}{\sum_w \alpha_w \sum_w \alpha_w^2 \sum_w \alpha_w^4} + \frac{\alpha_u^8 \alpha_v^8}{\sum_w \alpha_w \sum_w \alpha_w^2 \sum_w \alpha_w^4 \sum_w \alpha_w^8} + \dots \tag{15}$$

These flows exhaust the supply at every node u.

Let us bound f_{uv}^{rec} compared to $\alpha_u \alpha_v$. Let r_t be the t-th term in the series (15). We have that

$$r_{t+1} = \frac{\alpha_u^{2^t} \alpha_v^{2^t}}{\sum_w \alpha_w \sum_w \alpha_w^2 \dots \sum_w \alpha_w^{2^t}} = r_t \frac{\alpha_u^{2^{t-1}} \alpha_v^{2^{t-1}}}{\sum_w \alpha_w^{2^t}} \leq r_t \frac{\alpha_u^{2^{t-1}} \alpha_v^{2^{t-1}}}{(\alpha_u^{2^{t-1}})^2 + (\alpha_v^{2^{t-1}})^2} \leq \frac{r_t}{2}$$

From this, it follows that $f_{uv}^{\text{rec}} \leq r_1(1 + 1/2 + 1/4 + \dots) = 2\alpha_u \alpha_v / \sum_w \alpha_w$.

Analysis of the combined procedure. Let us call $f_{uv} = f_{uv}^{\text{cor}} + f_{uv}^{\text{rec}}$ the total flow sent on the arc (u, v) using the correction and recursive step. It is easy to show that all the constraints in (LP), except the second-stage conditional probability constraints (6), are satisfied by $p_{u,v} = f_{uv}$. To complete the validity of the construction we want to prove the bound $f_{uv} \leq 2\alpha_u \alpha_v$ in order to satisfy (6).

Note that the correction step induces a flow only when $u = 1$ or $v = 1$, and we focus on this case first. Suppose $u = 1$ (the case $v = 1$ follows by symmetry). We have that $f_{1v}^{\text{cor}} = \beta \alpha_1 \alpha_v$ and also $f_{1v}^{\text{rec}} \leq 2\alpha_1' \alpha_v' / \sum_w \alpha_w'$, where we denote by α' the remaining supply post-correction. Substituting the values α' we obtain

$$f_{1v}^{\text{cor}} + f_{1v}^{\text{rec}} \leq \alpha_1 \alpha_v \left(\beta + \frac{2(1 - \beta(1 - \alpha_1))(1 - \beta \alpha_1)}{1 - 2\beta \alpha_1 (1 - \alpha_1)} \right) = \frac{2 - \beta}{1 - 2\beta \alpha_1 (1 - \alpha_1)} \alpha_1 \alpha_v.$$

Recall that β lies in the interval $(0, 1/(1 - \alpha_1)]$. In this range, it is easy to show that the fraction multiplying $\alpha_1 \alpha_v$ is between $1/(1 - \alpha_1)$ and 2. Thus $f_{1v} \leq 2\alpha_1 \alpha_v$ as desired.

Now we focus on the case where $u, v \neq 1$. Although $f_{uv}^{\text{cor}} = 0$ in this case, the supply of the nodes is still modified. We have that $f_{uv}^{\text{rec}} \leq 2\alpha_1' \alpha_v' / \sum_w \alpha_w'$, where we denote by α' the supply remaining post-correction. Substituting the values α' we obtain

$$f_{uv}^{\text{rec}} \leq 2\alpha_u \alpha_v \frac{(1 - \beta \alpha_1)^2}{1 - 2\beta \alpha_1 (1 - \alpha_1)}.$$

In the range $\beta \in (0, 1/(1-\alpha_1)]$, the fraction multiplying $2\alpha_u\alpha_v$ on the right term lies in the range $\left[1 - \frac{\alpha_1^2}{(1-\alpha_1)^2}, 1\right)$. Thus $f_{uv} \leq 2\alpha_u\alpha_v$ as desired.

Algorithm 1 Explicit solution algorithm from Section 4.3

procedure ExplicitSolution($\alpha_1 \geq \alpha_2 \geq \ldots, \geq \alpha_n$)
  $\beta \leftarrow \frac{\alpha_1 - \alpha_2}{\alpha_1(1-\alpha_1-\alpha_2)}$
  $f^{\text{cor}}_{uv} \leftarrow 0, \quad u, v \in \{1, \ldots, n\}$
  $f^{\text{cor}}_{1v} \leftarrow f^{\text{cor}}_{v1} \leftarrow \beta\alpha_1\alpha_v, \quad v \in \{1, \ldots, n\}$
  $\alpha_v \leftarrow \alpha_v - f^{\text{cor}}_{v\cdot}, \quad v \in \{1, \ldots, n\}$
  $f^{\text{rec}} = \text{Recursive}(0, \alpha_1, \alpha_2, \ldots, \alpha_n)$
  return $f^{\text{cor}} + f^{\text{rec}}$
end procedure
procedure Recursive(f, $\alpha_1 = \alpha_2 \geq \ldots, \geq \alpha_n$)
  $f^{rec}_{uv} \leftarrow \frac{\alpha_u\alpha_v}{\sum_w \alpha_w}, \quad u, v \in \{1, \ldots, n\}$
  if $|f^{rec}| \leq \epsilon$ **then**
    return $f + f^{rec}$
  else
    return Recursive($f + f^{rec}$, $\frac{1}{\sum_w \alpha_w}\left(\alpha_1^2, \alpha_2^2 \ldots, \alpha_n^2\right)$)
  end if
end procedure

Theorem 4.4. *Algorithm 1 converges (if $\epsilon \to 0$) to an optimal solution to the* (LP) *model.*

Remark 7. Both steps are needed in the proof. The recursive step does not work when $\alpha_1 > \alpha_2$ because the remaining supply at node 1 does not converge to 0. Intuitively, the powers of α_1 become the dominant terms in the denominator way too quickly. Thus the correcting step is needed. On the other hand, the correcting step alone does not work in isolation. If $n = 3$ and $\alpha_1 = \alpha_2 = 0.45$, $\alpha_3 = 0.1$, it is not possible to send flow from nodes 1 and 2 in order to equate the remaining flow $\alpha'_1 = \alpha'_2 = \alpha'_3$, because α_3 is not big enough to receive the incoming flow from the nodes 1 and 2.

5. A more practical example

In this section, we describe one example to expose our model from a more practical point of view. We focus on how a decision-maker can calibrate the model in practice. Our example is based on real data from a restaurant chain, where preparation costs are fines are not easy to quantify. We focus mostly on the methodology, as the source data is confidential. In particular, we look at the outcomes of the (LP) and (BI) model.

The chain had eight stores $V = \{1, \ldots, 8\}$, and we will assume that the inspector can only visit two stores on a given day. One difficulty for the decision-maker is to estimate the preparation costs of each store. We propose the use of Key Performance Indicators (KPIs) to estimate them.

The decision-maker had access to KPIH_v and KPIS_v, two percentages related to the hygiene evaluation at the store and the variability in sales. In both indicators, the higher the percentage, the better the performance. These indicators are linked to organization and order and can be loosely correlated with preparation costs.

One possible estimate for d_v uses a linear combination of these KPIs:

$$d_v = \beta_{\mathrm{KPIH}} \cdot (1 - \mathrm{KPIH}_v) + \beta_{\mathrm{KPIS}} \cdot (1 - \mathrm{KPIS}_v),$$

where the weights $\beta_{\mathrm{KPIH}}, \beta_{\mathrm{KPIS}} \geq 0$ represent the relative importance of the different indicators.

The decision-maker can set fines based on the indicators related to the purpose of the inspection (the chain did not apply monetary fines). One proposal is to use KPIC_v and KPIL_v, two percentages related to customer evaluation and variability on budget, respectively. In both indicators, the higher the percentage, the better the performance. If the decision-maker needs to prioritize some stores, an additional term XP_v can reflect the strategic importance of store v for the chain, according to some objective measure. To guarantee that $d_v < f_v$, the fines can be built starting from the preparation costs:

$$f_v = d_v + \beta_{\mathrm{XP}} \cdot \mathrm{XP}_v + \beta_{\mathrm{KPIC}} \cdot (1 - \mathrm{KPIC}_v) + \beta_{\mathrm{KPIL}} \cdot \mathrm{KPIL}_v$$

Details on the construction of each KPI are specific to this chain. For instance, the opinions of clients captured via surveys were the base to compute KPIC. The KPIs value evolved over time, and so did the inspection routes in equilibrium. The appendix describes one inspection game that had to be solved, and the solution with the (BI) and (LP) models.

This methodology has several practical advantages. First, the inspection routes do not generate obvious patterns and can be repeated safely. Second, the use of these randomized routes and KPIs significantly reduces the inspector bias. Finally, a problem with a few inspections and a decent number of inspectees can be solved in any computer.

References

[1] Andreozzi, L. (2004). Rewarding policemen increases crime. another surprising result from the inspection game. *Public Choice*, 121(1):69–82.

[2] Avenhaus, R. and Okada, A. (1992). Statistical criteria for sequential inspector-leadership games. *Journal of The Operations Research Society of Japan*, 35:134–151.

[3] Avenhaus, R., von Stengel, B., and Zamir, S. (2001). Inspection games. In Aumann, R. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 51, pages 1947–1987. Elsevier, 1 edition.

[4] Avenhaus, R., Von Stengel, B., and Zamir, S. (2002). *Chapter 51 Inspection games*, volume 3, page 1947–1987. Elsevier.

[5] Baston, V. J. and Bostock, F. A. (1991). A generalized inspection game. *Naval Research Logistics (NRL)*, 38(2):171–182.

[6] Başar, T. and Olsder, G. J. (1998). *Dynamic Noncooperative Game Theory, 2nd Edition*. Society for Industrial and Applied Mathematics.

[7] Bierlein, D. (1968). Direkte inspektionssysteme. In *Operations Research Verfahren*.

[8] Bierlein, D. (1969). Auf bilanzen und inventuren basierende safeguards-systeme. In *Operations Research Verfahren*, volume 8, page 36–43.

[9] Borch, K. (1982). Insuring and auditing the auditor. In Deistler, M., Fürst, E., and Schwödiauer, G., editors, *Games, Economic Dynamics, and Time Series Analysis*, pages 117–126, Heidelberg. Physica-Verlag HD.

[10] Borch, K. H., Aase, K. K., and Sandmo, A. (1990). *Economics of insurance*. Advanced textbooks in economics. North-Holland; Distributors for the U.S. and Canada, Elsevier Science Pub. Co.

[11] Borndörfer, R., Buwaya, J., Sagnol, G., and Swarat, E. (2013). Optimizing toll enforcement in transportation networks: a game-theoretic approach. *Electronic Notes in Discrete Mathematics*, 41:253–260.

[12] Borndörfer, R., Omont, B., Sagnol, G., and Swarat, E. (2012). A stackelberg game to optimize the distribution of controls in transportation networks. In *International Conference on Game Theory for Networks*, pages 224–235. Springer.

[13] Breton, M., Alj, A., and Haurie, A. (1988). Sequential stackelberg equilibria in two-person games. *Journal of Optimization Theory and Applications*, 59:71–97.

[14] Bucarey, V., Casorrán, C., Labbé, M., Ordoñez, F., and Figueroa, O. (2019). Coordinating resources in stackelberg security games. *European Journal of Operational Research*.

[15] Conitzer, V. and Sandholm, T. (2006). Computing the optimal strategy to commit to. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, EC '06, page 82–90, New York, NY, USA. Association for Computing Machinery.

[16] Conitzer, V. and Sandholm, T. (2008). New complexity results about nash equilibria. *Games and Economic Behavior*, 63(2):621–641.

[17] Correa, J., Harks, T., Kreuzen, V., and Matuschke, J. (2014). Fare evasion in transit networks. *Operations Research*, 65.

[18] Dresher, M. (1962). A sampling inspection problem in arms control agreements: A game-theoretic analysis. Technical report, RAND Corporation Santa Monica, Calif.

[19] Feichtinger, G. (1983). A differential games solution to a model of competition between a thief and the police. *Management Science*, 29(6):686–699.

[20] Ferguson, T. S. and Melolidakis, C. (1998). On the inspection game. *Naval Research Logistics (NRL)*, 45(3):327–334.

[21] Filar, J. (1985). Player aggregation in the traveling inspector model. *IEEE Transactions on Automatic Control*, 30(8):723–729.

[22] Gilboa, I. and Zemel, E. (1989). Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93.

[23] Gold, M. (2019). Google and waze must stop sharing drunken-driving checkpoints, new york police demand. *New York Times*, page 23.

[24] Güth, W. and Pethig, R. (1992). Illegal pollution and monitoring of unknown quality—a signaling game approach—. In *Conflicts and cooperation in managing environmental resources*, pages 275–332. Springer.

[25] Höpfinger, E. (1979). Dynamic standard setting for carbon dioxide. In *Applied Game Theory*, pages 373–389. Springer.

[26] Jain, M., Conitzer, V., and Tambe, M. (2013). Security scheduling for real-world networks. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 215–222. International Foundation for Autonomous Agents and Multiagent Systems.

[27] Jain, M., Kardes, E., Kiekintveld, C., Ordonez, F., and Tambe, M. (2010a). Optimal defender allocation for massive security games: A branch and price approach. In *AAMAS 2010 Workshop on Optimisation in Multi-Agent Systems (OptMas)*.

[28] Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rathi, S., Tambe, M., and Ordóñez, F. (2010b). Software assistants for randomized patrol planning for the lax airport police and the federal air marshal service. *Interfaces*, 40:267–290.

[29] John Canty, M., Rothenstein, D., and Avenhaus, R. (2001). Timely inspection and deterrence. *European Journal of Operational Research*, 131(1):208–223.

[30] Katsikas, S., Kolokoltsov, V., and Yang, W. (2016). Evolutionary inspection and corruption games. *Games*, 7:31.

[31] Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Ordóñez, F., and Tambe, M. (2009). Computing optimal randomized resource allocations for massive security games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '09, page 689–696, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

[32] Leitmann, G. (1978). On generalized Stackelberg strategies. *Journal of Optimization Theory and Applications*, 26:637–643.

[33] Letchford, J. and Conitzer, V. (2013). Solving security games on graphs via marginal probabilities. In *AAAI*.

[34] Luh, P. B., Chang, S.-C., and Chang, T.-S. (1984). Solutions and properties of multi-stage stackelberg games. *Automatica*, 20(2):251 – 256.

[35] Magat, W. A. and Viscusi, W. K. (1990). Effectiveness of the epa's regulatory enforcement: The case of industrial effluent standards. *The Journal of Law and Economics*, 33(2):331–360.

[36] Maschler, M. (1966). A price leadership method for solving the inspector's non-constant-sum game. *Naval Research Logistics Quarterly*, 13(1):11–33.

[37] Nosenzo, D., Offerman, T., Sefton, M., and van der Veen, A. (2013). Encouraging Compliance: Bonuses Versus Fines in Inspection Games. *The Journal of Law, Economics, and Organization*, 30(3):623–648.

[38] O'Neill, B. (1994). Game theory models of peace and war. In Aumann, R. and Hart, S., editors, *Handbook of Game Theory with Economic Applications*, volume 2, chapter 29, pages 995–1053. Elsevier, 1 edition.

[39] Ore, O. (1956). Studies on directed graphs, i. *The Annals of Mathematics*, 63(3):383.

[40] Pita, J., Bellamane, H., Jain, M., Kiekintveld, C., Tsai, J., Ordóñez, F., and Tambe, M. (2009a). Security applications: Lessons of real-world deployment. *SIGecom Exch.*, 8(2).

[41] Pita, J., Jain, M., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., and Kraus, S. (2009b). Using game theory for los angeles airport security. *AI Magazine*, 30:43–57.

[42] Reinganum, J. F. and Wilde, L. L. (1986). Equilibrium verification and reporting policies in a model of tax compliance. *International Economic Review*, pages 739–760.

[43] Russell, C. (1990). Game models for structuring monitoring and enforcement systems. *Natural Resource Modeling*, 4(2):143–173.

[44] Schelling, T. C. and Halperin, M. H. (1985). *Strategy and arms control.* Pergamon-Brassey's.

[45] Schleicher, H. (1971). A recursive game for detecting tax law violations. *Economies et Sociétés*, 5(8):1421.

[46] Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., Maule, B., and Meyer, G. (2012). Protect: A deployed game theoretic system to protect the ports of the united states. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 13–20.

[47] Tambe, M. (2011). *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge university press.

[48] Thomas, M. and Nisgav, Y. (1976). An infiltration game with time dependent payoff. *Naval Research Logistics Quarterly*, 23(2):297–302.

[49] Tsai, J., Rathi, S., Kiekintveld, C., Ordonez, F., and Tambe, M. (2009). Iris-a tool for strategic security allocation in transportation networks. *AAMAS (Industry Track)*, pages 37–44.

[50] Tsebelis, G. (1990). Penalty has no impact on crime:: A game-theoretic analysis. *Rationality and Society*, 2(3):255–286.

[51] Von Stackelberg, H. (1934). *Marktform und Gleichgewicht*. J. Springer.

[52] Von Stengel, B. and Zamir, S. (2004). Leadership with commitment to mixed strategies. Technical report, Technical Report LSE-CDAM-2004-01, CDAM Research Report.

[53] Von Stengel, B. and Zamir, S. (2010). Leadership games with convex strategy sets. *Games and Economic Behavior*, 69:0–457.

[54] Washburn, A. and Wood, K. (1995). Two-person zero-sum games for network interdiction. *Operations research*, 43(2):243–251.

[55] Weissing, F. and Ostrom, E. (1991). Irrigation institutions and the games irrigators play: Rule enforcement without guards. In *Game equilibrium models II*, pages 188–262. Springer.

[56] Yin, Z., Jiang, A. X., Johnson, M. P., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., Tambe, M., and Sullivan, J. P. (2012). Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *IAAI*.

Appendix A. Omitted details and proofs for Section 3

We first recall that in Propositions 3.3 and 3.4 we are assuming that $\sum_{u\in V}\frac{d_u}{f_u}=2$ and we consider the following indices:

(i) Let $k\in[n]$ be such that: $\sum_{m\leq k} d_m/f_m \leq 1 \wedge \sum_{m\leq k+1} d_m/f_m > 1$;

(ii) Let $\kappa(u)\in[n]\setminus\{u\}$ be such that:

$$\sum_{m\leq\kappa(u),m\neq u} d_m/f_m \leq 1 \wedge \sum_{m\leq\kappa(u)+1,m\neq u} d_m/f_m > 1$$

.

Appendix A.1. Proof of Proposition 3.3

Proof. Given $u\in V'$, we claim that $(BI^{(2)}(u))$ model can be solved by the greedy algorithm. Let define $x_{v|u}$ as $x_{v|u} := \frac{f_v}{d_v}\, p_{v|u}$ $(\forall u,v\in V)$ then $(BI^{(2)}(u))$ can be written as follows.

$$\max\left\{\sum_{v\neq u} x_{v|u}\, d_v : x_{v|u}\leq 1,\, \sum_{v\neq u} x_{v|u}\,\frac{d_v}{f_v}\leq 1,\, x_{v|u}\geq 0 \quad (\forall v\neq u\in V')\right\}.$$

The previous model has a fractional knapsack structure so its solution is given by

$$p^*_{v|u} = \begin{cases} \dfrac{d_v}{f_v} & \textit{if } v\leq\kappa(u), v\neq u \qquad \text{(A.1)}\\[2ex] 1-\displaystyle\sum_{m\leq\kappa(u),m\neq u}\frac{d_m}{f_m} & \textit{if } v=\kappa(u)+1, v\neq u \qquad \text{(A.2)}\\[2ex] 0 & \text{otherwise.} \qquad \text{(A.3)}\end{cases}$$

□

Appendix A.2. Proof of Proposition 3.4

Proof. Let us proceed by cases:

- $u\leq k$: We start by noticing that by Proposition 3.3, for any node $u\leq k$ and $v\neq u$, $p^*_{u|v}=d_u/f_u$. In particular, by constraint (8),

$$p_{u\cdot}\leq\frac{d_u}{f_u}-\sum_{v\neq u}p^*_{u|v}p_{v\cdot}=\frac{d_u}{f_u}\Big(1-\sum_{v\neq u}p_{v\cdot}\Big)=0.$$

  In conclusion, $p_{u\cdot}=0$.

- $u = k+1$: Using the property proved in the previous case, we have by imposing equality in constraint (8),

$$\begin{aligned} & p_{(k+1)\cdot} + \textstyle\sum_{v \geq k+2} p^*_{(k+1)|v} p_{v\cdot} = \frac{d_{k+1}}{f_{k+1}} \\ \Longrightarrow \quad & p_{(k+1)\cdot} + \gamma(1 - p_{(k+1)\cdot}) = \frac{d_{k+1}}{f_{k+1}}, \end{aligned}$$

  where at the last implication we used the characterization $p^*_{(k+1)|v} = \gamma$, from Proposition 3.3. Hence, $p_{(k+1)\cdot} = (d_{k+1}/f_{k+1} - \gamma)/(1-\gamma)$.

- $u \geq k+2$: In this case, we also impose constraint (8) with equality, obtaining

$$p_{u\cdot} + \sum_{v \geq (k+1), v \neq u} p^*_{u|v} p_{v\cdot} = \frac{d_u}{f_u}.$$

  Notice further, by Proposition 3.3 that since $u \geq k+2$, if $v > k+1$ then $p^*_{u|v} = 0$; that, together with the formula above for $p_{(k+1)\cdot}$, leads to

$$p_{u\cdot} + p^*_{u|(k+1)} \frac{d_{k+1}/f_{k+1} - \gamma}{1-\gamma} = \frac{d_u}{f_u},$$

  which implies the claimed formula for $p_{u\cdot}$.

We now verify that other constraints of the problem are satisfied:

- Constraint (9). Substituting $p_{u\cdot}$, we obtain

$$\sum_u p_{u\cdot} = \frac{d_{k+1}/f_{k+1} - \gamma}{1-\gamma} + \sum_{u > k+1} \left(\frac{d_u}{f_u} - p^*_{u|k+1} \left(\frac{d_{k+1}/f_{k+1} - \gamma}{1-\gamma} \right) \right)$$

  By Proposition 3.3, $\sum_{u>k+1} p^*_{u|k+1} = 1 - \sum_{u \leq k} p^*_{u|k+1} = \gamma$. The expression on the right-hand side simplifies to

$$\sum_u p_{u\cdot} = \sum_{u \geq k+1} \frac{d_u}{f_u} - \gamma = \sum_{u \geq k+1} \frac{d_u}{f_u} + \sum_{u \leq k} \frac{d_u}{f_u} - 1 = 1.$$

- Non-negativity of variables. First of all, notice that $1 \geq d_{k+1}/f_{k+1} \geq \gamma$ by definition, therefore by inspection $p_{u\cdot} \leq d_u/f_u$ for all u. Non-negativity for $u = k+1$ is obtained by the same considerations, whereas for $u > k+1$ it is proved as follows: since $p^*_{u|k+1} \leq d_u/f_u$

$$p_{u\cdot} \quad \geq \quad \frac{d_u}{f_u}\Big(1 - \frac{d_{k+1}/f_{k+1} - \gamma}{1-\gamma}\Big) = \frac{d_u}{f_u} \frac{1 - d_{k+1}/f_{k+1}}{1-\gamma} \geq 0.$$

Finally, since $(BI^{(1)})$ is a packing linear program, and we are satisfying all packing constraints with equality, the optimality follows. □

Appendix B. About the generalization of Theorem 4.2 to other marginals

Theorem 4.2 cannot be generalized to every possible set of marginals satisfying $p_{v\cdot} + p_{\cdot v} = r_v$, where $(r_v)_{v \in V}$ is a solution to (RedLP). Consider a set $V = [3]$ of three operators having $d_1/f_1 = 0.8$, $d_2/f_2 = 0.5$, $d_3/f_3 = 0.7$, with $f_1 > f_2 > f_3 > 0$. In this example, $r_v = d_v/f_v$ for all $v \in V$. Consider the following marginals $p_{1\cdot} = 0, p_{\cdot 1} = 0.8, p_{2\cdot} = 0.5, p_{\cdot 2} = 0, p_{3\cdot} = 0.5, p_{\cdot 3} = 0.2$. We will show that the transportation model T has no feasible solutions.

Indeed, there is only one solution in T that satisfies supply, demand, non-repetition ($p_{uu} = 0, \forall u$) and non-negativity. This solution is shown in Figure B.2. However, this solution does not satisfy the capacity constraint $p_{3,1} \leq (d_1/f_1)\ p_{3\cdot}$ in the model T. Hence, T is infeasible.

Figure B.2: Unfeasible flow.

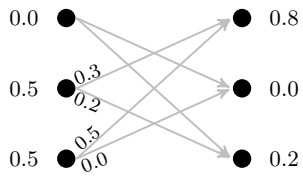

Appendix C. Addressing the case $\sum_u \alpha_u > 1$ in Section 4.3.

Algorithm 1 is an explicit solution for model (LP), under the assumption that $\sum_u \alpha_u = 1$.

The construction of the explicit solution when $\sum_u \alpha_u > 1$ can be reduced to this especial case. Assuming that the values f_u are sorted decreasingly, we can limit the support of the inspections in this solution to the first j nodes, where j is the smallest index for which $\sum_{u \leq j} \alpha_u > 1$.

If $\sum_{u \leq j} \alpha_u > 1$, we can define $\bar{\alpha}$ so that $\bar{\alpha}_u = \alpha_u$ for $u < j$, and $\bar{\alpha}_j$ is set so that $\sum_{u \leq j} \bar{\alpha}_u = 1$. We can use Algorithm 1 on $\bar{\alpha}$ to build an explicit solution, understanding that $p_{uv} = 0$ if $u > j$ or $v > j$. Using the fact that $\bar{\alpha}_u \leq \bar{\alpha}_u$ for all u, it is easy to check that all constraints in the original model will also be satisfied by this solution. And it is optimal since it gives the same objective value as the optimum of (RedLP).

Appendix D. Supplementary tables for Section 5

Table D.3 gives the parameters $(f_u)_{u \in V}$ and $(d_u/f_u)_{u \in V}$.

Table D.3: Parameters f_u and d_u/f_u for each store u.

u	f_u	d_u/f_u
1	4.23	0.83
2	3.60	0.95
3	4.60	0.76
4	5.43	0.81
5	3.00	0.82
6	5.17	0.86
7	7.77	0.89
8	2.20	0.82

Tables D.4 - D.5 shows the linear probabilities. The inspector gives positive probabilities for store 4, 6 and 7 with the highest value $\boldsymbol{f}$.

Table D.4: Probabilities of inspection paths p_{uv} for the (BI) model.

Stores	4	6	7
4	0.000	0.087	0.700
6	0.023	0.000	0.190
7	0.000	0.000	0.000

Table D.5: Probabilities of inspection paths p_{uv} for the (LP) model.

Stores	4	6	7
4	0.000	0.065	0.340
6	0.045	0.000	0.105
7	0.360	0.085	0.000