

Copyright
by
Michael Victor Nehme
2009

The Dissertation Committee for Michael Victor Nehme
certifies that this is the approved version of the following dissertation:

**Two-Person Games for Stochastic Network Interdiction:
Models, Methods, and Complexities**

Committee:

David P. Morton, Supervisor

Jonathan F. Bard

Constantine Caramanis

John Hasenbein

Elmira Popova

R. Kevin Wood

**Two-Person Games for Stochastic Network Interdiction:
Models, Methods, and Complexities**

by

Michael Victor Nehme, B.S.E.; M.S.

DISSERTATION

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

December 2009

To my grandparents.

Acknowledgments

I would like to begin by thanking my advisor and financial supporter Professor David P. Morton. His classes and his uncanny ability to come up with a great answer on the spot to almost any question made my studies significantly easier than they would have been otherwise. He also imparted enough of his speaking and writing style onto me that people would sometimes guess that he was my advisor before being told. I've been told this is a mostly good thing.

I would like to thank my dissertation committee members, Professors Jonathan F. Bard, Constantine Caramanis, John Hasenbein, Elmira Popova, and Kevin Wood. In particular, I thank Dr. Popova for providing financial support near the beginning of my program and for having the utmost faith in my abilities. Also, I thank Dr. Wood for providing the basis for much of my research and for providing many suggestions for the improvement of my dissertation.

A special thanks goes to my high school math teachers Mrs. Michelle Haubert and Mr. Jeff Limber for teaching ridiculously good classes, and in particular Mr. Limber for ensuring that I would always remember how to spell the word "parallel." Finally, I would like to thank Professor Gary Ybarra of Duke University for making my first engineering class a great one, for greatly

improving my work ethic, and for forcing us to write everything in L^AT_EX.

Two-Person Games for Stochastic Network Interdiction: Models, Methods, and Complexities

Publication No. _____

Michael Victor Nehme, Ph.D.
The University of Texas at Austin, 2009

Supervisor: David P. Morton

We describe a stochastic network interdiction problem in which an interdictor, subject to limited resources, installs radiation detectors at border checkpoints in a transportation network in order to minimize the probability that a smuggler of nuclear material can traverse the residual network undetected. The problems are stochastic because the smuggler's origin-destination pair, the mass and type of material being smuggled, and the level of shielding are known only through a probability distribution when the detectors are installed. We consider three variants of the problem. The first is a Stackelberg game which assumes that the smuggler chooses a maximum-reliability path through the network with full knowledge of detector locations. The second is a Cournot game in which the interdictor and the smuggler act simultaneously. The third is a "hybrid" game in which only a subset of detector locations is revealed to the smuggler.

In the Stackelberg setting, the problem is NP-complete even if the interdicator can only install detectors at border checkpoints of a single country. However, we can compute wait-and-see bounds in polynomial time if the interdicator can only install detectors at border checkpoints of the origin and destination countries. We describe mixed-integer programming formulations and customized branch-and-bound algorithms which exploit this fact, and provide computational results which show that these specialized approaches are substantially faster than more straightforward integer-programming implementations. We also present some special properties of the single-country case and a complexity landscape for this family of problems.

The Cournot variant of the problem is potentially challenging as the interdicator must place a probability distribution over an exponentially-sized set of feasible detector deployments. We use the equivalence of optimization and separation to show that the problem is polynomially solvable in the single-country case if the detectors have unit installation costs. We present a row-generation algorithm and a version of the weighted majority algorithm to solve such instances. We use an exact-penalty result to formulate a model in which some detectors are visible to the smuggler and others are not. This may be appropriate to model “decoy” detectors and detector upgrades.

Table of Contents

Acknowledgments	v
Abstract	vii
List of Tables	xii
List of Figures	xiii
Chapter 1. Introduction	1
1.1 Game Theory	1
1.2 Two-Person Zero-Sum Cournot Games	6
1.3 The Weighted Majority Algorithm	9
1.4 Benders' Decomposition	14
1.5 Network Interdiction	15
1.6 Attacker-Defender Models and Exact Penalty Results	19
1.7 Overview of the Contents	23
Chapter 2. Bipartite Stochastic Network Interdiction Problem	25
2.1 Introduction	25
2.2 Problem Description	28
2.3 Scenario Aggregation	33
2.4 Step Inequalities	36
2.5 Reformulation	42
2.6 Efficient Nested Solutions	50
2.7 Preprocessing	62
2.8 Branching	64
2.9 Computational Results	69

Chapter 3. Two-Person Zero-Sum Games for Network Interdiction	77
3.1 Introduction	77
3.2 TPZSCGs with Exponentially Many Strategies	78
3.2.1 Motivation	78
3.2.2 Single-Evader Model	79
3.2.3 Multiple-Evader Model	83
3.2.4 Cardinality-Constrained Case	86
3.2.5 Weighted Majority	89
3.3 Optimal Design of a Two-Person Zero-Sum Cournot Game	92
3.3.1 Investing in Premium Strategies	93
3.3.2 Tightening the Formulation	95
3.3.3 Removing an Opponent's Premium Strategies	98
3.3.4 Complexity	101
3.4 Transparent and Non-transparent Assets	105
3.4.1 Knapsack-Constrained Non-transparent Assets	105
3.4.2 Cardinality-Constrained Non-transparent Assets	111
3.4.3 Reformulation for $p = q$	113
3.4.4 Reformulation for $p \neq q$	116
Chapter 4. Two-Country Network Interdiction	119
4.1 Introduction	119
4.2 Two-Country Deterministic Network Interdiction Problem	120
4.2.1 Formulation	120
4.2.2 Simplifying the Model	123
4.2.3 Solution Techniques	125
4.3 Two-Country Stochastic Network Interdiction Problem	130
4.3.1 Identifying Additional Non-binding Constraints	132
4.3.2 Reformulation	134
4.3.3 Branching	137
4.3.4 Computational Results	142
4.4 Complexity	148

4.4.1	Two-Country Stochastic Network Interdiction Problem with Dualized Budget Constraint	148
4.4.2	Three-Country Deterministic Network Interdiction Problem	155
Chapter 5.	Conclusion	163
	Appendices	167
Appendix A.	Customized Branch-and-Bound for BiSNIP	168
Appendix B.	Customized Branch-and-Bound for the Two-Country Stochastic Network Interdiction Problem	171
	Bibliography	174
	Vita	180

List of Tables

2.1	Optimality gaps as percentages of optimal unconditional evasion probabilities for US model instances with perfectly reliable detectors, $q = 0$	74
2.2	Solution times in seconds for US model instances with perfectly reliable detectors, $q = 0$. \times indicates that the solution time exceeded 2 hours.	74
4.1	Solution times in seconds for Russia-US model instances with perfectly reliable detectors. \times indicates that the solution time exceeded 2 hours.	144
4.2	Complexity landscape of the maximum-reliability network interdiction problem.	162

List of Figures

2.1	(a) Network topology after preprocessing for the single country problem. Only border crossing arcs can receive detectors, so the network between the source and the tails of the checkpoint arcs and the heads of the checkpoint arcs and the destination can be reduced. (b) The equivalent bipartite network.	29
2.2	An example efficient frontier of solutions. Solutions corresponding to points A , B , and C can be found in polynomial time and are nested.	52
2.3	The figure shows 136 motor-crossing checkpoints from Canada and Mexico into the continental United States, and groups the checkpoints into four clusters.	71
2.4	The figure shows the improvement factor as a function of the number of detectors installed for the US model. The four plots correspond to different levels of effectiveness of the detectors, specifically, with $q = 0.75, 0.50, 0.25$ and 0 in $q_k^w = qp_k$	72
2.5	Part (a) of the figure shows the optimal solution to the US model instance with perfectly reliable detectors, $q = 0$, and with a budget to install detectors at $b = 10$ locations. Part (b) of the figure is identical but for $b = 11$. Note that the full number of checkpoints are not visible in the map due to their close proximity.	75
2.6	Part (a) of the figure shows the optimal solution to the US model instance with perfectly reliable detectors, $q = 0$, and with a budget to install detectors at $b = 96$ locations. Part (b) of the figure is identical but for $f = 97$. Part (c) of the figure is for $b = 97$ and $q = 0.75$	76
4.1	Topology of the preprocessed network for the two-country stochastic network interdiction problem. Only border crossing arcs can receive detectors.	121
4.2	Motor-crossing checkpoints for a two-country instance	143
4.3	The improvement factor as a function of the number of detectors installed for the better of the two one-country solutions and the two-country solution for the Russia-US model.	145

4.4	Optimal solutions to the Russia-US model instance with perfectly reliable detectors.	146
4.5	Optimal solution to the Russia-US model for $b = 88$	147
4.6	Transformed instance of (4.24) for $V = \{1, 2, 3\}$ and $E = \{(1, 2), (2, 3)\}$. Each arrow indicates a path that the smuggler has access to under the given scenario. All available paths have nominal evasion probability 1.	151
4.7	Network corresponding to (4.34) for $V = \{1, 2, 3\}$ and $E = \{e_1, e_2\}$ where $e_1 = (1, 2)$ and $e_2 = (2, 3)$. Solid arrows indicate arcs with reliability 1, while dotted arrows indicate arcs with reliability ϵ	159

Chapter 1

Introduction

1.1 Game Theory

In this dissertation, we describe several models for deploying radiation detectors in order to minimize the probability that a smuggler of nuclear material can avoid detection. In a standard mathematical programming model, a decision maker chooses values for a set of variables known as *decision variables* in order to minimize or maximize some function known as an *objective function* or simply the *objective*. Many such models are well studied and are appropriate if the objective depends only on decision variables controlled by the decision maker. For our models, it is prudent to assume that the smuggler is a strategic thinker and makes decisions in order to maximize the probability that he can avoid detection. We therefore need to borrow heavily from game theory, which was developed to model situations in which multiple parties with differing and possibly conflicting objectives can affect each other's objectives. In a game, each party or player has a set of "strategies" from which to choose and an objective function which typically depends on the strategies chosen by the other players. A strategy can be defined as "a rule for action so complete and detailed that a player need not actually be present once his strategy is known" [38]. A game can be either zero-sum or general-sum; in the former

case the sum of all the players' gains is a constant and in the latter the sum depends on which strategies the players select. Players may choose, and reveal, their strategies either simultaneously or sequentially; we refer to the former case as a "Cournot" game and the latter as a "Stackelberg" game.

Most credit game theory's inception to von Neumann and Morgenstern [37] who show that in any two-person zero-sum game, there exists an equilibrium in which neither player has an incentive to change his own strategy unilaterally. That is, neither player wishes to change his strategy even if he is aware of the other player's strategy. In many cases this equilibrium is only guaranteed to exist if we allow the players to randomize their strategies. That is, instead of choosing a single "pure" strategy, a player may choose a "mixed" strategy which assigns a probability to each pure strategy. In a game such as Rock, Paper, Scissors, for example, it is clear that no pure-strategy equilibrium exists but if we allow randomization then there exists a mixed-strategy equilibrium in which each player uniformly and randomly selects one of the three available strategies. Nash [28] shows that such an equilibrium exists even in general-sum games and thus permanently attached his name to the concept.

Zero-sum games became a natural model for many military applications as it is generally wise to assume your enemy is a strategic thinker. One of the earliest and most well known of such models is the Colonel Blotto game [6, 35] in which two players allocate soldiers to N independent battlefields. On each battlefield the player who allocates the most soldiers wins and the player who wins the most battlefields wins the overall game. The classic Blotto game

has become the baseline model for much of the recent research regarding the allocation of resources to defend against terrorist threats, although many of the latest models are not zero-sum.

Bier et al. [5] consider a model in which a defender allocates resources to two sites and an attacker can choose to attack exactly one of the sites. While the attacker is assumed to know the defender's valuations of the sites, the defender only has a distribution function for the attacker's valuations. The defender seeks to minimize the expected loss due to an attack plus cost of defenses and the attacker seeks to maximize the expected payoff of the attack. They prove the existence of a pure strategy Nash equilibrium for both players. Zhuang and Bier [44] describe a model for allocating resources to defend against the dual threats of terrorism and natural disasters. In the model, a defender can continuously allocate resources to defend against both threats at multiple sites while an attacker continuously allocates resources to attack those sites. Both simultaneous and sequential games are studied and, perhaps surprisingly, the authors show that the defender may actually do better by revealing his strategy to the attacker if the attacker's optimal response is unique. This so-called "first-mover advantage" is only possible, however, in a general-sum setting. They also remark that pure strategy equilibria are common in their game because continuous decision variables are used and the probability of a successful attack is assumed to be convex in the defender's allocations and concave in the attacker's.

Powell [30] describes a similar model where the defender again allocates

continuous resources to several potential targets. These defensive resources decrease the probability of a successful attack. The defender and the attacker place possibly different values on each target, and the attacker chooses the target with the highest expected payoff, defined as the product of the probability of success and the attacker's value of the target. In the zero-sum setting in which the attacker and defender have the same valuations of the targets, he shows that the simultaneous and sequential game both have the same payoff. He also shows that the defender's optimal strategy is to allocate resources to the target with the highest value until the expected payoff equals that of the target with the second highest value. Then resources are allocated to both targets until their expected payoff equals that of the third highest value and so on, until the defender exhausts his resources. He also claims that this strategy is optimal even if the attacker and defender have different valuations of the targets.

Though intuitively it seems that a defender benefits by keeping his allocations a secret, in the papers discussed previously, secrecy is either irrelevant or may even hurt the defender. This may be the case in a general-sum game but is never the case in a zero-sum game. Consider the following representation of a two-person zero-sum game:

$$\min_{x \in X} \max_{y \in Y} f(x, y).$$

Here the player who chooses x must choose first and reveal this decision to the other player, who then chooses y . The game is zero-sum because one player

seeks to maximize $f(x, y)$ and the other seeks to maximize $-f(x, y)$. We can then make the following observation.

Proposition 1. *Let*

$$v_1^* = \min_{x \in X} \max_{y \in Y} f(x, y)$$

and

$$v_2^* = \max_{y \in Y} \min_{x \in X} f(x, y).$$

Then $v_1^* \geq v_2^*$.

Proof. Let x_1^* and y_1^* solve the “min-max” problem and x_2^* and y_2^* solve the “max-min” problem. Then

$$v_1^* = f(x_1^*, y_1^*) \geq f(x_1^*, y_2^*) \geq f(x_2^*, y_2^*) = v_2^*.$$

□

So in a zero-sum game, the second-mover may have an advantage. So-called *minimax* theorems identify conditions under which $v_1^* = v_2^*$. One of the most famous such theorems is the following, due to Fan [11].

Proposition 2. *If X and Y are compact, convex sets and if $f(\cdot, y)$ is convex on X for all $y \in Y$ and $f(x, \cdot)$ is concave on Y for all $x \in X$ then $v_1^* = v_2^*$.*

If the conditions of Proposition 2 hold then neither player can benefit from secrecy and thus the payoff of the game is the same whether the game is sequential or simultaneous. For the games considered in [44], these conditions

are satisfied due to continuous decision variables and the assumption of diminishing returns. In our work, some decisions are naturally discrete and we can only achieve convexity of the feasible regions by allowing mixed strategies.

1.2 Two-Person Zero-Sum Cournot Games

In a two-person zero-sum Cournot game (TPZSCG), each player has a set of “pure” strategies from which to choose, and strategies are assumed to be selected simultaneously by the two players. The assumption that players act simultaneously is valid provided neither player can ascertain the other player’s selection. In the zero-sum case, which we focus on here, there exists an objective function, sometimes called a *payoff* function, which depends on both players’ selections and which one player seeks to maximize and the other seeks to minimize [38]. In the so-called *normal form* representation of a TPZSCG we use a payoff matrix to store all possible payoffs of the game [37]. This matrix contains a row for each pure strategy available to the minimizing player, who we call the “row player”, and a column for each pure strategy available to the maximizing player, who we call the “column player.” When playing pure strategies the row player picks a row, the column player simultaneously picks a column, and the entry of the matrix that is at the intersection of the row and column picks, becomes the payoff for the game. The normal form is valid whenever both players have a finite number of pure strategies.

To address the simultaneous nature of the players’ decisions, instead of selecting a single pure strategy, each player is motivated to select a *mixed*

strategy, a probability distribution over the available pure strategies. Then, when the game is played each player randomly selects a pure strategy according to the chosen probability distribution.

To formulate the normal form TPZSCG as a mathematical program, let $i \in I$ index the row players pure strategies and $j \in J$ index those of the column player. Then the row player's decision variables, $x_i, i \in I$, represent the probability that pure strategy i is selected and the column player's decision variables, $y_j, j \in J$, represent the probability that pure strategy j is selected. If the payoff matrix is given by $A_{ij}, i \in I, j \in J$, then the TPZSCG can be formulated as follows:

$$v^* = \max_{y \in Y} \min_{x \in X} \sum_{i \in I} \sum_{j \in J} A_{ij} x_i y_j, \quad (1.1)$$

where $Y = \{y \in \mathbb{R}_+^{|J|} : \sum_{j \in J} y_j = 1\}$ and $X = \{x \in \mathbb{R}_+^{|I|} : \sum_{i \in I} x_i = 1\}$. The objective function is simply the expected payoff conditioned on each player's mixed strategy. The ordering of the "max" and "min" seems to imply that the column player selects his mixed strategy first and reveals his selection to the row player. However, the fact that X and Y are convex and compact, coupled with the fact that the objective function in (1.1) is concave in y for fixed x and convex in x for fixed y , implies that, by Proposition 2, we obtain the same optimal value, v^* , if we exchange the order of the "max" and "min", i.e., if the row player selects his mixed strategy first and reveals his selection to the column player. Due to this symmetry, we can view the formulation as requiring that the two players choose their mixed strategies simultaneously. This also

implies that the optimal solution of this problem gives a Nash equilibrium, that is, even if one player knows the other's mixed strategy, that player still has no incentive to deviate from his original mixed strategy. If in model (1.1), we (a) fix $y \in Y$ to create a linear program (LP) with variables x , (b) define λ as the dual variable for the single structural constraint in that LP, (c) take the dual of the LP, and (d) then release y , we obtain the following equivalent problem:

$$v^* = \max_{y, \lambda} \quad \lambda \tag{1.2a}$$

$$\text{s.t.} \quad \lambda \leq \sum_{j \in J} A_{ij} y_j \quad : x_i, \quad i \in I \tag{1.2b}$$

$$\sum_{j \in J} y_j = 1 \tag{1.2c}$$

$$y_j \geq 0, \quad j \in J. \tag{1.2d}$$

Here, the column player takes a convex combination of the columns of A , and the new decision variable λ is the smallest entry in the resulting column vector. Thus λ represents the value of the row player's best response to the column player's mixed strategy. From the column player's perspective this is a (seemingly) pessimistic view since it assumes the row player knows his mixed strategy and acts accordingly. Equivalently, however, we can return to model (1.1) and do the following: (a) interchange the "max" and "min", (b) fix $x \in X$ to create an LP with variables y , (c) define θ as the dual variable for the single structural constraint in that LP, (d) take the dual of the LP, and

(e) release y . This results in the following problem:

$$v^* = \min_{x, \theta} \theta \tag{1.3a}$$

$$\text{s.t.} \quad \theta \geq \sum_{i \in I} A_{ij} x_i \quad : y_j, \quad j \in J \tag{1.3b}$$

$$\sum_{i \in I} x_i = 1 \tag{1.3c}$$

$$x_i \geq 0, \quad i \in I. \tag{1.3d}$$

Here, the row player takes a convex combination of the rows of A and θ falls onto the largest value of the resulting row vector. This would seem pessimistic from the point of view of the row player. But problems (1.2) and (1.3) are each other's duals and therefore have the same optimal value according to strong duality. This proves Proposition 2 for the special case in which the payoff function is bilinear and the players' feasible regions are unit simplices. Also note that either LP gives an optimal mixed strategy for both players; the primal solution gives probabilities for one player and the dual variables the probabilities for the other.

1.3 The Weighted Majority Algorithm

For a two-person zero-sum Cournot game in which each player has a modest number of strategies, the linear programs described in the previous section provide an easy way to compute the Nash equilibrium. For the network interdiction models considered in this dissertation, however, the interdicator may have an exponentially-sized set of strategies. We could use standard row- or column-generation techniques to overcome this difficulty, but instead turn

to a generalization of the “weighted majority” algorithm originally developed by Littlestone and Warmuth [22] to solve on-line allocation models. Freund and Schapire show that a slight modification of this algorithm can be used to approximately solve games with a large or even unknown payoff matrix [12].

The on-line allocation model can be defined as follows. For every time step $t = 1, \dots, T$, an allocation agent must assign a probability p_i^t to each strategy $i \in I$. After each time period, the agent observes a loss $l_i^t \in [0, 1]$ for each strategy, and after T time periods incurs an expected cumulative loss of

$$L = \sum_{t=1}^T \sum_{i \in I} l_i^t p_i^t.$$

The agent’s goal is to minimize the *net loss*, defined as the difference between the expected cumulative loss and the minimum cumulative loss over all strategies. That is, the agent seeks to minimize

$$L - \min_{i \in I} L_i.$$

where $L_i = \sum_{t=1}^T l_i^t$. The weighted majority algorithm prescribes a choice for p_i^t , $i \in I$, $t \in \{1, \dots, T\}$, that we describe shortly. The net loss of the weighted majority algorithm is bounded by $O(\sqrt{T \log |I|})$ even if we allow the loss vectors l^t to be chosen in an adversarial fashion and to depend on the agent’s choice of the distribution p^t [13].

The algorithm is easy to implement. At each time period t , the agent keeps a non-negative weight w_i^t on each strategy and chooses a probability

distribution over the strategies by normalizing the weights as follows:

$$p_i^t = \frac{w_i^t}{\sum_{i' \in I} w_{i'}^t}, \quad i \in I, \quad t = 1, \dots, T.$$

The agent, after observing the losses for each strategy, then updates the weights by the following rule:

$$w_i^{t+1} = w_i^t \beta^{l_i^t},$$

where $\beta \in (0, 1)$ is an appropriately chosen constant. The following result from [13] relates the loss incurred by the weighted majority algorithm to the loss of the best strategy.

Theorem 3. *Let L_β be the cumulative loss incurred by an agent applying the weighted majority algorithm for a given value of β . For any sequence of loss vectors l^1, \dots, l^T and for any $i \in I$, we have*

$$L_\beta \leq \frac{-\ln(w_i^1) - L_i \ln \beta}{1 - \beta}.$$

In particular if we let $i \in \operatorname{argmin}_{i \in I} L_i$ and choose $w_i^1 = 1/|I|$, $i \in I$, we get the following bound:

$$L_\beta \leq \frac{\min_{i \in I} L_i \ln(1/\beta) + \ln |I|}{1 - \beta}.$$

To choose β , suppose that we have an upper bound \tilde{L} on the loss incurred by the best strategy, i.e., $\min_{i \in I} L_i \leq \tilde{L}$. Then if we choose

$$\beta = \frac{1}{1 + \sqrt{\frac{2 \ln |I|}{\tilde{L}}}}$$

we achieve the following upper bound on the average loss incurred by the weighted majority algorithm:

$$\frac{L_\beta}{T} \leq \min_{i \in I} \frac{L_i}{T} + \frac{\sqrt{2\tilde{L} \ln |I|}}{T} + \frac{\ln |I|}{T}.$$

Now consider a two-person zero-sum Cournot game with a payoff matrix A and pure strategies for the row and column player indexed by $i \in I$ and $j \in J$, respectively. Every time period t , which we interpret as iterations in an algorithmic solution procedure, the probabilities p_i^t give a valid mixed-strategy x_i^t over the row player's pure strategies $i \in I$. Suppose we choose the loss for each strategy $i \in I$ as $l_i^t = \sum_{j \in J} A_{ij} y_j^t$, where $(y_j^t)_{j \in J}$ is the column player's best response to the mixed strategy $(x_i^t)_{i \in I}$ given by an optimal solution to

$$\max_{y \in Y} \sum_{i \in I} \sum_{j \in J} A_{ij} x_i^t y_j.$$

If we let $\bar{x}_i = \frac{1}{T} \sum_{t=1}^T x_i^t$ and $\bar{y}_j = \frac{1}{T} \sum_{t=1}^T y_j^t$, then \bar{x} and \bar{y} are valid mixed-strategies for the row and column player, respectively. Following the analysis of [12], we can show that \bar{x} and \bar{y} approximate the Nash equilibrium of the game within $\Delta_T = \frac{\sqrt{2\tilde{L} \ln |I|}}{T} + \frac{\ln |I|}{T}$. First, we show that $\frac{L_\beta}{T}$ gives an upper

bound on the row player's loss if he uses the mixed-strategy \bar{x} .

$$\begin{aligned}
\max_{y \in Y} \sum_{i \in I} \sum_{j \in J} A_{ij} \bar{x}_i y_j &= \max_{y \in Y} \sum_{i \in I} \sum_{j \in J} A_{ij} \left(\frac{1}{T} \sum_{t=1}^T x_i^t \right) y_j \\
&\leq \frac{1}{T} \sum_{t=1}^T \max_{y \in Y} \sum_{i \in I} \sum_{j \in J} A_{ij} x_i^t y_j \\
&= \frac{1}{T} \sum_{t=1}^T \sum_{i \in I} \sum_{j \in J} A_{ij} x_i^t y_j^t \\
&= \frac{L_\beta}{T}.
\end{aligned}$$

Also, we show that $\min_{i \in I} \frac{L_i}{T}$ is exactly the column player's payoff if he uses the mixed-strategy \bar{y} and the row player responds optimally.

$$\begin{aligned}
\min_{i \in I} \frac{L_i}{T} &= \min_{i \in I} \frac{1}{T} \sum_{t=1}^T l_i^t \\
&= \min_{i \in I} \frac{1}{T} \sum_{t=1}^T \sum_{j \in J} A_{ij} y_j^t \\
&= \min_{x \in X} \frac{1}{T} \sum_{t=1}^T \sum_{i \in I} \sum_{j \in J} A_{ij} x_i y_j^t \\
&= \min_{x \in X} \sum_{i \in I} \sum_{j \in J} A_{ij} x_i \bar{y}_j.
\end{aligned}$$

Using the fact that $\frac{L_\beta}{T} \leq \min_{i \in I} \frac{L_i}{T} + \Delta_T$ we can now say the following:

$$\max_{y \in Y} \sum_{i \in I} \sum_{j \in J} A_{ij} \bar{x}_i y_j \leq \min_{x \in X} \sum_{i \in I} \sum_{j \in J} A_{ij} x_i \bar{y}_j + \Delta_T. \quad (1.6)$$

Inequality (1.6) implies that if the row player uses the mixed-strategy \bar{x} , the loss he incurs never exceeds the Nash equilibrium loss by more than Δ_T . Likewise, subtracting Δ_T from both sides of (1.6), we see that if the column player

uses the mixed-strategy \bar{y} , the payoff he receives falls short of the Nash equilibrium payoff by at most Δ_T . It is in this sense that \bar{x} and \bar{y} approximate the Nash equilibrium of the game.

1.4 Benders' Decomposition

Consider an LP of the following form:

$$z^* = \min_{x \geq 0, y \geq 0} \quad cx + fy \quad (1.7a)$$

$$\text{s.t.} \quad Ax = b \quad (1.7b)$$

$$-Bx + Dy = d. \quad (1.7c)$$

Assume that this LP has a finite optimal solution and that for every x that satisfies $Ax = b, x \geq 0$, there exists a y that satisfies $Dy = Bx + d, y \geq 0$. Then (1.7) can be rewritten as:

$$z^* = \min_{x \geq 0} \quad cx + h(x) \quad (1.8a)$$

$$\text{s.t.} \quad Ax = b, \quad (1.8b)$$

where

$$h(x) = \min_{y \geq 0} \quad fy \quad (1.9a)$$

$$\text{s.t.} \quad Dy = Bx + d \quad : \pi. \quad (1.9b)$$

The dual feasible region of (1.9) is $\Pi = \{\pi : \pi D \leq f\}$ and the dual objective is $\pi(Bx + d)$. Therefore with $\pi^{(1)}, \dots, \pi^{(L)}$ denoting the extreme points of Π ,

$h(x) = \max_{1 \leq i \leq L} \pi^{(i)}(Bx + d)$. This implies that $h(x) \geq \pi^{(i)}(Bx + d), i = 1, \dots, L$, and so (1.8) can be written as:

$$z^* = \min_{x \geq 0, \theta} \quad cx + \theta \quad (1.10a)$$

$$\text{s.t.} \quad Ax = b \quad (1.10b)$$

$$\theta \geq \pi^{(i)}Bx + d, \quad i = 1, \dots, L. \quad (1.10c)$$

Constraints (1.10c) are called “optimality cuts.” Of course, enumerating all of the extreme points of Π is grossly inefficient so we aim to solve relaxations of the master problem (1.10) with optimality cuts generated for a small subset of the extreme points of Π . We generate additional optimality cuts by solving instances of the subproblem (1.9) with x fixed to an optimal solution of a relaxation of (1.10). This identifies the most violated constraint in (1.10c) which can then be added to the relaxation of the master problem. Since Π has a finite number of extreme points this process is guaranteed to terminate in a finite number of iterations. The above algorithm forms an outer-linearization of $h(x)$. In stochastic programming this algorithm is known as the L-shaped method [36]. In integer programming a similar method is due to Benders [4] and in nonlinear programming a related method is due to Kelley [19].

1.5 Network Interdiction

Network interdiction deals with problems in which an interdictor, subject to one or more resource constraints, can structurally or parametrically alter a given network. The interdictor does so knowing that an adversary

solves a network optimization problem in the resulting network, and the interdicator's goal is to optimally degrade the adversary's performance. Many models for network interdiction have been proposed, varying in the objectives of the interdicator and adversary, the manner in which the interdicator may change the network, and the underlying network optimization problem.

Wollmer considers problems in which the interdicator can remove a fixed number of arcs from the network with the goal of either minimizing the maximum flow [41] or maximizing the minimum cost flow [42]. In Corley and Chang [8], the interdicator removes a fixed number of nodes, and all arcs incident to those nodes, in order to minimize the maximum flow. Fulkerson and Harding [14] show that if the interdicator may continuously increase the lengths of the arcs in a shortest path problem subject to a linear cost, then the problem is equivalent to a minimum cost flow problem. Bayrak and Bailey [3] consider a shortest path interdiction problem with asymmetric information in that the adversary does not know the true lengths of the arcs. The authors formulate the problem as a mixed-integer nonlinear program, which they then linearize.

Wood [43] proves that the problem of minimizing the maximum flow subject to a cardinality constraint on the number of arcs that can be removed is NP-complete in the strong sense. He also shows that the problem can be formulated as a mixed-integer program by using the equivalence between the max-flow and min-cut problems and suggests several valid inequalities that can be used to strengthen the formulation. Lim and Smith [21] consider two multicommodity versions of this problem; in the first arcs are completely de-

stroyed when interdicted, and in the second arc capacities can be continuously decreased. Smith et al. [33] considers the problem of designing a multicommodity flow network that is robust with respect to an intelligent attacker. The authors propose a three-stage model; in the first stage the network designer constructs a network, in the second stage an attacker reduces the capacity of some arcs in the network, and in the third stage the designer solves a multicommodity flow problem on the residual network. The authors consider three models of attacker behavior; in the first two the attacker destroys arcs in a greedy fashion and in the third the attacker seeks to minimize the maximum post-interdiction profit. Each case is formulated as a mixed-integer program, and a cutting-plane algorithm is suggested for the third.

Israeli and Wood [18] develop a decomposition scheme for the problem of maximizing the adversary's shortest path when the interdictor can discretely increase the length of some subset of arcs. Specifically, a master problem first finds an interdiction plan that maximizes the minimum length of some subset of the adversary's feasible paths, then a subproblem generates the best response path to this interdiction plan and adds it to the subset of paths considered by the master. The optimal value of the master problem and the length of the best response path provide upper and lower bounds, respectively, to allow termination with an ϵ -optimal solution. The authors also introduce the notion of a *super-valid inequality*, an inequality which is guaranteed not to remove all optimal solutions to a problem unless the optimal solution has already been found. The idea behind these super-valid inequalities is as follows.

The subproblem generates the adversary's best response to a feasible interdiction plan and consequently produces a lower bound to the master problem. To achieve a maximum shortest path greater than this lower bound, the interdictor must remove at least one arc on the shortest path generated by the subproblem. Several variants of the super-valid inequality are discussed and are shown to considerably speed solution time.

Washburn and Wood [39] consider a model in which the interdictor chooses an arc k to interdict and the adversary simultaneously chooses a path to traverse. If the adversary's path includes the interdicted arc, then the adversary is detected with probability p_k , otherwise the adversary is not detected. They show that optimal mixed strategies for both players can be found by solving a minimum-cut problem. Interestingly, the optimal mixed strategy for the interdictor is to interdict arcs along this minimum-cut with a probability that is inversely proportional to the detection probability on that arc. Several extensions to the basic model are considered. In particular, under certain conditions the problem in which the interdictor can choose multiple arcs to interdict can also be solved via a minimum-cut problem.

Morton et al. [26] look at stochastic network interdiction problems with a focus on the prevention of nuclear smuggling. In the basic model, every arc in a transportation network has an associated evasion probability, and the interdictor may discretely decrease the evasion probability on some subset of those arcs by installing detectors with the goal of minimizing the evasion probability of the maximum-reliability path. The model is stochastic

in that the interdicator does not know the adversary’s origin and destination but instead has a probability distribution over a number of potential origin-destination pairs. A simplified model is discussed in which the interdicator may only install detectors at border checkpoints of a single country, and valid inequalities known as *step inequalities* are developed to tighten the resulting mixed-integer program. Morton and Pan [25] develop an enhanced L-shaped decomposition method for solving the general model. Step inequalities are also developed to tighten the linear programming relaxation of the associated master problem.

1.6 Attacker-Defender Models and Exact Penalty Results

Here we describe a canonical problem which we use as the basis for many of our models. Brown et al. [7] propose an “attacker-defender” model in which the defender operates a system according to the following LP:

$$\min_{y \geq 0} \quad cy \tag{1.11a}$$

$$\text{s.t.} \quad Ay = b \tag{1.11b}$$

$$Fy \leq u. \tag{1.11c}$$

Constraints (1.11b) include operation requirements of the defender’s system and (1.11c) are capacity constraints for each of the defender’s assets that is vulnerable to attack. An attacker, subject to his own resource constraints, attacks some subset of these assets with the goal of maximizing the defender’s

cost of operating the resulting system. Assuming an attack on an asset is guaranteed to destroy the entire capacity of that asset, the attacker’s problem can be formulated as:

$$\max_{x \in X} h(x), \tag{1.12}$$

where

$$h(x) = \min_{y \geq 0} cy \tag{1.13a}$$

$$\text{s.t. } Ay = b \tag{1.13b}$$

$$Fy \leq U(e - x) \quad : -\pi(x). \tag{1.13c}$$

Here $U = \text{diag}(u)$, e is the vector of all 1s, and $x \in X$ contains binary restrictions on x as well as the attacker’s resource constraints. If a component of x takes value 1 then the corresponding capacity drops to 0, and otherwise the capacity remains at its nominal value. Since X is not a convex set, the order of the “max” and “min” cannot be interchanged. The defender then operates the system with any assets which were not destroyed. A natural way to reformulate (1.12) is to convert the inner linear program to a maximization problem by taking its dual. Unfortunately this yields a term containing the product of x and π and thus a nonlinear mixed-integer program.

The underlying problem here is that since x appears on the right-hand side of the constraints, $h(\cdot)$ is a convex function over the convex hull of X . But since we are maximizing over x , we would prefer $h(\cdot)$ to be concave. This can be achieved if we can relax constraints (1.13c) and add an appropriate

penalty term to the objective. The following refinement of an exact-penalty result from [27] does exactly this.

Proposition 4. *Assume*

$$v^* = \min_{y \geq 0} \quad cy \tag{1.14}$$

$$\text{s.t.} \quad Ay = b \quad : \gamma \tag{1.15}$$

$$Fy \leq u \quad : -\pi, \tag{1.16}$$

has a finite optimal solution. If $\bar{\pi} \geq \pi^$ for some optimal dual subvector π^* and*

$$v^{**} = \min_{y \geq 0} \quad cy + \bar{\pi}(Fy - u)^+ \tag{1.17}$$

$$\text{s.t.} \quad Ay = b, \tag{1.18}$$

where $x^+ = \max(x, 0)$, then $v^ = v^{**}$.*

Proof. Problem (1.17) can be reformulated as:

$$v^{**} = \min_{y \geq 0, \theta \geq 0} \quad cy + \bar{\pi}\theta$$

$$\text{s.t.} \quad Ay = b \quad : \gamma$$

$$Fy - \theta \leq u \quad : -\pi.$$

The dual of this reformulation is:

$$v^{**} = \max_{\gamma, \pi} \quad \gamma b - \pi u \tag{1.19a}$$

$$\text{s.t.} \quad \gamma A - \pi F \leq c \quad : y \tag{1.19b}$$

$$\pi \leq \bar{\pi} \quad : \theta \tag{1.19c}$$

$$\pi \geq 0. \tag{1.19d}$$

This is exactly the dual of (1.15) with the addition of constraint (1.19c). However, by assumption $\bar{\pi} \geq \pi^*$ for some optimal π^* , so the addition of (1.19c) does not remove all optimal solutions to the dual of (1.15). Therefore by strong duality $v^* = v^{**}$. \square

Applying Proposition 4 to (1.13) results in a penalty term of the form $\bar{\pi}(Fy - U(e - x))^+$. A standard linearization of this term moves x back into the constraints. But the fact that the components of x are binary coupled with Fy being nominally bounded above by u allows us to keep x in the objective and out of the constraints. To see why, we write the k th component of $(Fy - U(e - x))^+$ as $(F_{k \cdot}y - u_k(1 - x_k))^+$, which equals 0 if $x_k = 0$ and equals $F_{k \cdot}y$ if $x_k = 1$. Thus, we have $\bar{\pi}(Fy - U(e - x))^+ = x^T \bar{\Pi} Fy$ for all $x \in X$ where $\bar{\Pi} = \text{diag}(\bar{\pi})$.

We now define a new function $\bar{h}(x)$ as follows:

$$\bar{h}(x) = \min_{y \geq 0} \quad cy + x^T \bar{\Pi} Fy \quad (1.20a)$$

$$\text{s.t.} \quad Ay = b \quad (1.20b)$$

$$Fy \leq u, \quad (1.20c)$$

where $\bar{\pi} \geq \pi^*(x)$ for all $x \in X$ and $\pi^*(x)$ is an optimal dual subvector to (1.13). By Proposition 4, $h(x) = \bar{h}(x)$ for all $x \in X$, but $\bar{h}(\cdot)$ is concave over the convex hull of X , making it more amenable to maximization. In particular we can either take the dual of (1.20) to obtain a mixed-integer linear program, or apply a decomposition scheme such as Benders' decomposition.

1.7 Overview of the Contents

Pan [29] considers a stochastic network interdiction problem in which the interdictor seeks to minimize the maximum-reliability path in a network. The interdictor, subject to a budget constraint, installs radiation detectors on some subset of arcs in the network. A detector has the effect of decreasing the reliability of the arc at which it is installed. The adversary is a smuggler of nuclear material whose characteristics, such as his origin and destination, the mass and type of material being smuggled, and the thickness of the lead shielding, are known only via a probability distribution at the time the interdictor installs the detectors. This dissertation augments and extends that work as follows.

In Chapter 2, we consider a special case in which the interdictor can only install detectors at border checkpoints of a single country. We assume that the smuggler chooses a path with knowledge of the detector locations and arc reliabilities. While [29] shows that the problem is NP-complete subject to a cardinality or knapsack constraint on the number of detectors that are installed, we describe a formulation of this problem which has a totally unimodular constraint matrix when the budget constraint is relaxed. As a result, the budget-constrained version of the problem can be solved in polynomial time under certain conditions. We show that the solutions that can be found in polynomial time are *nested*, that is, for any pair of such solutions, the checkpoints which receive detectors in the solution with the smaller budget will be a subset of those which receive detectors in the other. We also describe

a customized branch-and-bound algorithm which performs well in practice.

Chapter 3 considers a variant of the single-country problem in which the interdictor and the smuggler act simultaneously. The resulting model is a two-person zero-sum Cournot game in which the interdictor can have exponentially-many strategies. We discuss the complexity of the model and describe a solution technique based on the weighted majority algorithm. In the cardinality-constrained case, we show that the value of the game can be found by solving a polynomially-sized linear program. We also suggest two Cournot-Stackelberg “hybrid” models in which one player may either purchase additional pure strategies for his own use or remove pure strategies from his opponent’s strategy set. The latter model is used as the basis for an interdiction model in which some detectors are transparent to the smuggler and others are not.

Chapter 4 considers a network interdiction problem in which the interdictor may install detectors along border checkpoints of two countries, typically the origin and destination countries of a smuggler. We show that this problem can be solved in polynomial time if the smuggler characteristics are known before detectors are installed. We also show that in the stochastic setting, many of the solution techniques described in Chapter 2 for the single-country problem have natural extensions into the two-country version and are very effective at reducing computational effort. We conclude with some complexity results which show that some polynomially-solvable interdiction problems become NP-complete when we add an additional country to the problem.

Chapter 2

Bipartite Stochastic Network Interdiction Problem

2.1 Introduction

We describe a stochastic network interdiction model designed to locate radiation detectors, which detect gamma and neutron emissions from nuclear material, at critical border crossings. The goal is to locate the detectors on an underlying transportation network to minimize the probability of a successful smuggling attempt. We focus on the development of a strengthened mixed-integer programming formulation and a customized branch-and-bound algorithm which reduce the required computational effort.

We model two adversaries, an interdictor and a smuggler, and a transportation network $G(N, A)$. The smuggler starts at origin node $o \in N$ and wishes to reach destination node $d \in N$. The probability that the smuggler will evade detection while traversing arc $(i, j) \in A$ is q_{ij} if the interdictor installs a detector on (i, j) and $p_{ij} > q_{ij}$ otherwise. At most one detector may be installed per arc. A smuggler can be caught by indigenous law enforcement without detection equipment, and so $p_{ij} < 1$. Detection events on distinct arcs are assumed to be mutually independent. The smuggler chooses an o - d path to

maximize his evasion probability. With limited resources, the interdicator must select arcs on which to install detectors in order to minimize this probability.

The *threat scenario*, indexed by $\omega \in \Omega$, specifies the origin-destination pair, (o^ω, d^ω) , as well as other details about the nuclear material being smuggled and the manner in which it is shielded. So, the probability a smuggler evades detection if a detector is installed on arc (i, j) is scenario dependent, i.e., q_{ij} depends on ω , denoted q_{ij}^ω . In general, the indigenous evasion probabilities, p_{ij} , could also depend on the threat scenario. The bulk of what we present is valid when $p_{ij} = p_{ij}^\omega$, but in Section 2.3 we discuss a computationally valuable variable-aggregation scheme that arises naturally when p_{ij} does not depend on ω . The threat scenario is unknown when detectors are installed, but is governed by a probability mass function, $p^\omega, \omega \in \Omega$, which is assumed to be known. In what follows, “threat scenario ω ” will often be shortened to simply “smuggler ω .”

The timing of the interdicator’s and smuggler’s decisions and the realization of the threat scenario is as follows: First, the interdicator installs detectors on a subset of the network’s arcs subject to a budget constraint. Then, a threat scenario is revealed and the smuggler selects a path that solves

$$\max_{P \in \mathcal{P}_{o^\omega, d^\omega}} \prod_{(i,j) \in P} [p_{ij}^\omega(1 - x_{ij}) + q_{ij}^\omega x_{ij}], \quad (2.1)$$

where $\mathcal{P}_{o^\omega, d^\omega}$ is the set of all o^ω - d^ω paths and where $x_{ij} = 1$ if a detector is installed on arc (i, j) and $x_{ij} = 0$ otherwise. We conservatively assume the smuggler selects a path with full knowledge of the detector locations and eva-

sion probabilities. For a given $\omega \in \Omega$ and installation plan $x_{ij}, (i, j) \in A$, the value of (2.1) gives the probability that the smuggler traverses the network undetected, conditional on the realization of threat scenario ω . The interdicator seeks to minimize the sum of these conditional evasion probabilities, each weighted by p^ω , over all threat scenarios. Problem (2.1) is a maximum-reliability path problem since both the evasion probabilities and the installation plan are fixed and known to the smuggler by the time he selects a path. Then we can view the interdicator's problem as minimizing the expected value of the maximum-reliability path.

Morton et al. [26] formulate the problem on a general network as a mixed-integer program and Pan and Morton [25] develop an enhanced L-shaped decomposition method and use valid inequalities when solving the associated master problem. See [18] for the deterministic version of this problem, [3] for a variant with asymmetric information, and [9] for a variant in which interdiction successes are uncertain. We give a more extensive review of network interdiction research in Section 1.5. This chapter develops mixed-integer programming formulations and enhanced branch-and-bound algorithms for the special case of this problem in which a smuggler encounters at most one detector. These enhancements significantly reduce computational effort. Also, while this special case is known to be NP-complete, we describe a condition under which we can solve an instance in polynomial time.

2.2 Problem Description

We restrict attention to a special case that arises when we can only place detectors at border-crossing checkpoints of a single country. The key to simplifying the formulation in this case is that each o^ω - d^ω path has exactly one arc on which the smuggler could encounter a detector. Let K be the set of *checkpoint* arcs, i.e., arcs that a smuggler could traverse depending on the selected path, that could contain a detector. For each ω , we compute the value of the maximum-reliability path from o^ω to the tail of each checkpoint arc and the value of the maximum-reliability path from the head of each checkpoint arc to d^ω . Call the product of these two probabilities γ_k^ω , $k = (i, j) \in K$. Then the evasion probability for smuggler ω if he traverses checkpoint k is either $\gamma_k^\omega p_k^\omega$ or $\gamma_k^\omega q_k^\omega$, depending on whether a detector is installed. Figure 2.1 shows the topology of the preprocessed network and the transformed bipartite network. We can then formulate the bipartite stochastic network interdiction problem (BiSNIP) as follows.

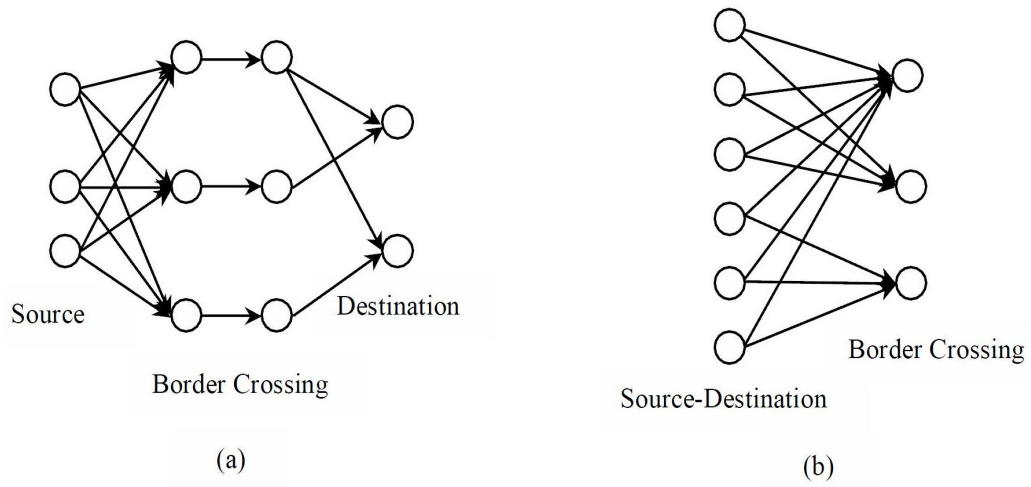


Figure 2.1: (a) Network topology after preprocessing for the single country problem. Only border crossing arcs can receive detectors, so the network between the source and the tails of the checkpoint arcs and the heads of the checkpoint arcs and the destination can be reduced. (b) The equivalent bipartite network.

Sets:

K set of border checkpoints

Data:

b total budget for installing detectors

c_k cost of installing a detector at border checkpoint $k \in K$

Random elements:

$\omega \in \Omega$ sample point and sample space for threat scenarios

p^ω probability mass function

p_k^ω probability smuggler ω can traverse checkpoint k undetected with no detector installed

$q_k^\omega < p_k^\omega$ probability smuggler ω can traverse checkpoint k undetected with a detector installed

γ_k^ω product of the values of the maximum-reliability paths from o^ω to the tail of arc k and from the head of arc k to d^ω for smuggler ω

Interdictor's decision variables:

x_k 1 if a detector is installed at checkpoint k and 0 otherwise

Smuggler's decision variables:

θ^ω evasion probability for smuggler ω

Formulation:

$$\min_{x, \theta} \sum_{\omega \in \Omega} p^\omega \theta^\omega \quad (2.2a)$$

$$\text{s.t. } x \in X \quad (2.2b)$$

$$\theta^\omega \geq \gamma_k^\omega p_k^\omega (1 - x_k), \quad k \in K, \omega \in \Omega \quad (2.2c)$$

$$\theta^\omega \geq \gamma_k^\omega q_k^\omega x_k, \quad k \in K, \omega \in \Omega, \quad (2.2d)$$

where $X = \{x \in \mathbb{B}^{|K|} : \sum_{k \in K} c_k x_k \leq b\}$.

BiSNIP (2.2) may be viewed on a bipartite network with arcs (ω, k) linking each threat scenario with its checkpoints. Variable θ^ω is the conditional

probability the smuggler avoids detection, given ω , and model (2.2) minimizes the (unconditional) probability the smuggler avoids detection. Constraints (2.2c) and (2.2d) force the evasion probability for each smuggler to equal that of the maximum reliability path, i.e., $\theta^\omega = \max_{k \in K} \{\gamma_k^\omega p_k^\omega (1 - x_k), \gamma_k^\omega q_k^\omega x_k\}$.

Our initial attempts to solve the BiSNIP model (2.2) using a branch-and-bound solution method indicated that BiSNIP's linear-programming (LP) relaxation can produce very weak lower bounds. The following proposition tightens constraints (2.2c) and effectively eliminates constraints (2.2d) in the BiSNIP model.

Proposition 5. *Consider the BiSNIP model (2.2), let $q_{\max}^\omega \equiv \max_{k \in K} \gamma_k^\omega q_k^\omega$, and assume $0 \leq q_k^\omega \leq p_k^\omega \leq 1$ and $0 \leq \gamma_k^\omega \leq 1$ for all $k \in K$, $\omega \in \Omega$. Then the inequalities*

$$\theta^\omega \geq \gamma_k^\omega p_k^\omega - (\gamma_k^\omega p_k^\omega - q_{\max}^\omega) x_k, \quad k \in K, \omega \in \Omega \quad (2.3a)$$

$$\theta^\omega \geq q_{\max}^\omega, \quad \omega \in \Omega \quad (2.3b)$$

are valid for BiSNIP.

Proof. Let $k^* \in \operatorname{argmax}_{k \in K} \gamma_k^\omega q_k^\omega$ for some $\omega \in \Omega$. If $x_{k^*} = 1$, then constraint (2.2d) dominates (2.2c) and yields $\theta^\omega \geq \gamma_{k^*}^\omega q_{k^*}^\omega = q_{\max}^\omega$. And, if $x_{k^*} = 0$ then constraint (2.2c) dominates (2.2d) and yields $\theta^\omega \geq \gamma_{k^*}^\omega p_{k^*}^\omega \geq \gamma_{k^*}^\omega q_{k^*}^\omega = q_{\max}^\omega$. This proves the validity of (2.3b). Now for any $k \in K$, if $x_k = 1$ then (2.3a) becomes (2.3b), and if $x_k = 0$ then (2.3a) is equivalent to (2.2c). Thus (2.3a) is valid as well. \square

We can view the right-hand side of (2.3b) as providing an optimistic bound, from the interdicator's perspective, on the evasion probability of smuggler ω . Then (2.3a) is simply a strengthened version of (2.2c) in which the right-hand side drops down to the lower bound q_{\max}^ω instead of zero when $x_k = 1$.

We can replace constraints (2.2c) and (2.2d) in BiSNIP with (2.3a) and (2.3b) since every constraint in the former set is dominated by some constraint in the latter. In doing so we obtain a model with half as many structural constraints and at least as strong an LP relaxation. Furthermore, defining $\bar{\theta}^\omega = \theta^\omega - q_{\max}^\omega$ and $r_k^\omega = (\gamma_k^\omega p_k^\omega - q_{\max}^\omega)^+$, where $(\cdot)^+ = \max(\cdot, 0)$, we can transform BiSNIP into a model in which $\bar{\theta}^\omega$ has simple lower bounds of zero:

$$\min_{x, \bar{\theta}} \quad \sum_{\omega \in \Omega} p^\omega \bar{\theta}^\omega \quad (2.4a)$$

$$\text{s.t.} \quad x \in X \quad (2.4b)$$

$$\bar{\theta}^\omega \geq r_k^\omega (1 - x_k), \quad k \in K, \quad \omega \in \Omega. \quad (2.4c)$$

Here $\gamma_k^\omega p_k^\omega \leq q_{\max}^\omega$ implies that $r_k^\omega = 0$ and in this case, the corresponding constraint (2.4c) reduces to a non-negativity constraint. This occurs when smuggler ω prefers a checkpoint with evasion probability q_{\max}^ω to that of checkpoint k . Model (2.4) implicitly ignores such checkpoint-smuggler pairs.

Model (2.4) is equivalent to BiSNIP in that both models have the same set of optimal solutions for locating the detectors, but their objective functions differ by the constant $\sum_{\omega \in \Omega} p^\omega q_{\max}^\omega$. We can view this as a transformation to a model in which the radiation detectors are perfectly reliable, i.e., model (2.4)

has the form of model (2.2) with $q_k^\omega = 0$. We conclude by noting that in some instances (2.4) can be further tightened by replacing (2.4c) with

$$\bar{\theta}^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)x_k, \quad k \in K, \quad \omega \in \Omega, \quad (2.5)$$

where

$$\underline{\theta}^\omega = \min_{x \in X} \max_{k \in K} r_k^\omega (1 - x_k), \quad \omega \in \Omega. \quad (2.6)$$

We can efficiently compute $\underline{\theta}^\omega$ for each $\omega \in \Omega$ by sorting the checkpoints in decreasing order by r_k^ω , then greedily allocating detectors until the interdiction budget b is depleted. This is equivalent to solving the wait-and-see problem, that is, each $\underline{\theta}^\omega$ is computed under the assumption that we know in advance that threat scenario ω is realized and allocate the detectors accordingly.

2.3 Scenario Aggregation

We now focus our discussion on the transformed model (2.4) but suppress the “bar” notation on θ^ω for simplicity. Suppose that for some pair of smugglers $\omega, \omega' \in \Omega$ we can index the checkpoints in K , $k_1, k_2, \dots, k_{|K|}$, such that $r_{k_1}^\omega \geq r_{k_2}^\omega \geq \dots \geq r_{k_{|K|}}^\omega$ and $r_{k_1}^{\omega'} \geq r_{k_2}^{\omega'} \geq \dots \geq r_{k_{|K|}}^{\omega'}$. That is to say, both of these particular smugglers may have different evasion probabilities at some or all checkpoints, but they can rank-order the checkpoints in an identical manner.

The motivation for considering the above situation in the context of the BiSNIP model arises as follows. Suppose the indigenous evasion probabilities

do not depend on the threat scenario. Consider two smugglers, ω and ω' , that are identical in every way, including their origin-destination pair, the mass and type of material they smuggle, etc., except that smuggler ω shields his material better than does smuggler ω' . Then, for each checkpoint the indigenous evasion probabilities associated with traveling from origin to destination via that checkpoint will be identical for both smugglers, $p_k^\omega \gamma_k^\omega = p_k^{\omega'} \gamma_k^{\omega'}$ for all $k \in K$. And, the evasion probability at each checkpoint will be larger for the smuggler with better shielding, $q_k^\omega > q_k^{\omega'}$ for all $k \in K$. This then results in smugglers ω and ω' ordering their checkpoints in an identical manner. As suggested above, there may be fewer positive values of r_k^ω , $k \in K$, than of $r_k^{\omega'}$, $k \in K$, but they will still satisfy the requisite (inclusive) ordering condition. The same result can arise, for example, when the two smugglers are carrying different masses of nuclear material, and it can arise for distinct origin-destination pairs, typically in close geographic proximity. It can also arise when the indigenous evasion probabilities depend on the threat scenario, as long as their ordering is identical. Specifically, since $r_k^\omega = (\gamma_k^\omega p_k^\omega - q_{\max}^\omega)^+$, two smugglers with different but identically ordered $\gamma_k^\omega p_k^\omega$ values satisfy the ordering condition.

Fix an interdiction plan for BiSNIP, $x = (x_k)_{k \in K} \in X$. Then for ω and ω' satisfying the identical-ordering assumption we have $\theta^\omega = r_{k^*}^\omega$ and $\theta^{\omega'} = r_{k^*}^{\omega'}$, where $k^* \in \operatorname{argmax}_{k \in K} \gamma_k^\omega (1 - x_k)$ and $k^* \in \operatorname{argmax}_{k \in K} \gamma_k^{\omega'} (1 - x_k)$ can be taken to be the same checkpoint. The contribution of θ^ω and $\theta^{\omega'}$ to the objective function (2.4a) is given by $p^\omega \theta^\omega + p^{\omega'} \theta^{\omega'}$. Of course, we do not know x ahead of time but we can replace ω and ω' with a single scenario, say $\bar{\omega}$. The objective

function coefficient of $\theta^{\bar{\omega}}$ is equal to $p^\omega + p^{\omega'}$, and the evasion probability at each checkpoint $k \in K$ is

$$\frac{p^\omega r_k^\omega + p^{\omega'} r_k^{\omega'}}{p^\omega + p^{\omega'}}$$

for scenario $\bar{\omega}$. Extending these ideas to an arbitrary number of scenarios yields the following proposition.

Proposition 6. *Consider model (2.4), let $x \in X$, and let $\theta^\omega = \max_{k \in K} r_k^\omega (1 - x_k)$. Suppose there exists a partition, $\Omega^n, n \in \mathcal{N}$, of Ω such that every smuggler in a particular subset Ω^n orders his evasion probabilities in an identical fashion. That is, for each $n \in \mathcal{N}$ there exists $k_1^n, k_2^n, \dots, k_{|K|}^n$ such that $r_{k_1^n}^\omega \geq r_{k_2^n}^\omega \geq \dots \geq r_{k_{|K|}^n}^\omega$ for all $\omega \in \Omega^n$. Let $\theta^{\omega^n} = \max_{k \in K} r_k^{\omega^n} (1 - x_k)$ where $r_k^{\omega^n} = \sum_{\omega \in \Omega^n} p^\omega r_k^\omega / p^{\omega^n}$ and where $p^{\omega^n} = \sum_{\omega \in \Omega^n} p^\omega$. Then $p^{\omega^n} \theta^{\omega^n} = \sum_{\omega \in \Omega^n} p^\omega \theta^\omega$.*

Proof. Under the ordering assumption for r_k^ω , $\omega \in \Omega$, for each $x \in X$ and $n \in \mathcal{N}$, there exists a k^* such that $r_{k^*}^\omega = \max_{k \in K} r_k^\omega (1 - x_k)$, $\forall \omega \in \Omega^n$. Since $p^\omega \geq 0, \forall \omega$, k^* also maximizes $\sum_{\omega \in \Omega^n} p^\omega r_k^\omega (1 - x_k)$. Thus,

$$\begin{aligned} p^{\omega^n} \theta^{\omega^n} &= \max_{k \in K} p^{\omega^n} r_k^{\omega^n} (1 - x_k) \\ &= \max_{k \in K} \sum_{\omega \in \Omega^n} p^\omega r_k^\omega (1 - x_k) \\ &= \sum_{\omega \in \Omega^n} p^\omega \max_{k \in K} r_k^\omega (1 - x_k) \\ &= \sum_{\omega \in \Omega^n} p^\omega \theta^\omega. \end{aligned}$$

□

Corollary 7. *Under the hypotheses of Proposition 6, the following model is equivalent to model (2.4):*

$$\begin{aligned}
\min_{x, \theta} \quad & \sum_{n \in \mathcal{N}} p^{\omega_n} \theta^{\omega_n} \\
\text{s.t.} \quad & x \in X \\
& \theta^{\omega_n} \geq r_k^{\omega_n} (1 - x_k), \quad k \in K, \quad n \in \mathcal{N}.
\end{aligned} \tag{2.7}$$

In the equivalent aggregated model (2.7), $r_k^{\omega_n}$ and θ^{ω_n} are still conditional evasion probabilities but are now conditioned on the event $\omega \in \Omega^n$ whereas their counterparts in (2.4) were conditioned on the realization of a single threat scenario. Similarly $p^{\omega_n} = P(\Omega^n)$ is the probability that a threat scenario in Ω^n is realized.

2.4 Step Inequalities

Previous work (e.g., [26, 29]) tightening the LP relaxation of (2.4) involves the development of a class of valid inequalities known as step inequalities. To motivate these step inequalities, we first define some notation which we will use for the remainder of this chapter. Let $k(i, \omega) \in K$ be an index mapping of the checkpoints such that $r_{k(i, \omega)}^{\omega} \geq r_{k(i+1, \omega)}^{\omega}$ for all $i = 1, \dots, |K| - 1$ and $\omega \in \Omega$ and that $\cup_{i=1}^{|K|} \{k(i, \omega)\} = K$ for all $\omega \in \Omega$. Then we define a set K_k^{ω} for every $k \in K$ and $\omega \in \Omega$ which satisfies

$$K_{k(i, \omega)}^{\omega} = \{k(i', \omega) : 1 \leq i' < i\}, \quad 1 \leq i \leq |K|, \quad \omega \in \Omega. \tag{2.8}$$

We can view K_k^ω as the set of all checkpoints which smuggler ω ranks higher than checkpoint k , with ties between checkpoints being resolved arbitrarily. So every $k' \in K_k^\omega$ satisfies $r_{k'}^\omega \geq r_k^\omega$ and every $k' \notin K_k^\omega$ satisfies $r_{k'}^\omega \leq r_k^\omega$. Now let $T(\omega) = \{k_1, \dots, k_l\} \subseteq K$ satisfy $k_i \in K_{k_{i+1}}^\omega$ for $i = 1, \dots, l-1$. Then we must have $r_{k_1}^\omega \geq r_{k_2}^\omega \geq \dots \geq r_{k_l}^\omega$ and can define a step inequality on $T(\omega)$ as follows:

$$\theta^\omega \geq r_{k_1}^\omega - (r_{k_1}^\omega - r_{k_2}^\omega)x_{k_1} - \dots - (r_{k_l}^\omega - r_{k_{l+1}}^\omega)x_{k_l}, \quad (2.9)$$

where $r_{k_{l+1}}^\omega \equiv 0$. The number of step inequalities for every scenario can be exponential in $|K|$ and so adding all possible step inequalities to (2.4) is out of the question. Instead, we iteratively solve the linear-programming relaxation of (2.4) and add step inequalities on an as-needed basis. The *separation problem* for step inequalities requires that given $(\hat{x}, \hat{\theta})$, a feasible solution to the LP relaxation of (2.4), we either identify a most violated step inequality for each ω or determine that none are violated. To find a most violated step inequality for some ω , if it exists, we must find a $T(\omega)$ which maximizes the right-hand side of (2.9). That is, we must solve:

$$z^\omega = \max_{T(\omega) \subseteq K} r_{k_1}^\omega - (r_{k_1}^\omega - r_{k_2}^\omega)\hat{x}_{k_1} - \dots - (r_{k_l}^\omega - r_{k_{l+1}}^\omega)\hat{x}_{k_l}, \quad (2.10)$$

for each $\omega \in \Omega$. The following results show that we can determine whether a particular k should be included in an optimal solution to (2.10) by sorting. For this result and for the remainder of the chapter we reserve the “hat” notation, as in $(\hat{x}, \hat{\theta})$, for solutions of the LP relaxation.

Proposition 8. *Let $\omega \in \Omega$ and $\hat{x} \in [0, 1]^{|K|}$ be given. There exists an optimal solution $T^*(\omega)$ to (2.10) in which $k \in T^*(\omega)$ if and only if either $K_k^\omega = \emptyset$ or*

$$\hat{x}_k < \min_{k' \in K_k^\omega} \hat{x}_{k'}. \quad (2.11)$$

Proof. We first show that we can construct an optimal solution to (2.10) which contains every k which satisfies either $K_k^\omega = \emptyset$ or (2.11). We then show that any checkpoint k which satisfies neither $K_k^\omega = \emptyset$ nor (2.11) can be removed from this optimal solution without decreasing the objective.

Suppose that $T^*(\omega)$ is an optimal solution to (2.10) and z^ω is the corresponding optimal value. We first consider the case in which $K_k^\omega = \emptyset$ for some $k \in K \setminus T^*(\omega)$,

$$\begin{aligned} z^\omega &= r_{k_1}^\omega - \dots - (r_{k_l}^\omega - r_{k_{l+1}}^\omega) \hat{x}_{k_l} \\ &= r_k^\omega - (r_k^\omega - r_{k_1}^\omega) - \dots - (r_{k_l}^\omega - r_{k_{l+1}}^\omega) \hat{x}_{k_l} \\ &\leq r_k^\omega - (r_k^\omega - r_{k_1}^\omega) \hat{x}_k - \dots - (r_{k_l}^\omega - r_{k_{l+1}}^\omega) \hat{x}_{k_l}. \end{aligned} \quad (2.12)$$

The inequality in (2.12) holds since $\hat{x}_k \leq 1$ and since $K_k^\omega = \emptyset$ implies $r_k^\omega \geq r_{k_1}^\omega$. But since (2.12) is simply the objective function of (2.10) evaluated at $T(\omega) = T^*(\omega) \cup \{k\}$, adding k to an optimal solution of (2.10) will maintain optimality. We turn to the case in which there exists $k \in K \setminus T^*(\omega)$ such that $K_k^\omega \neq \emptyset$ and $\hat{x}_k < \min_{k' \in K_k^\omega} \hat{x}_{k'}$. If $T^*(\omega) \cap K_k^\omega = \emptyset$, then adding k to $T^*(\omega)$ cannot decrease the objective function as shown above. So we assume $T^*(\omega) \cap K_k^\omega \neq \emptyset$ and let $i' = \max_{\{i: k_i \in T^*(\omega) \cap K_k^\omega\}} i$. Then $k_{i'} \in K_k^\omega$ and for $i' \neq l$,

$k_{i'+1} \notin K_k^\omega$. We can then add k to the right-hand side of the step inequality between $k_{i'}$ and $k_{i'+1}$ as follows:

$$\begin{aligned}
z^\omega &= r_{k_1}^\omega - \dots - (r_{k_{i'}}^\omega - r_{k_{i'+1}}^\omega) \hat{x}_{k_{i'}} - \dots \\
&= r_{k_1}^\omega - \dots - (r_{k_{i'}}^\omega - r_k^\omega) \hat{x}_{k_{i'}} - (r_k^\omega - r_{k_{i'+1}}^\omega) \hat{x}_{k_{i'}} - \dots \\
&\leq r_{k_1}^\omega - \dots - (r_{k_{i'}}^\omega - r_k^\omega) \hat{x}_{k_{i'}} - (r_k^\omega - r_{k_{i'+1}}^\omega) \hat{x}_k - \dots. \tag{2.13}
\end{aligned}$$

The inequality in (2.13) holds since $k_{i'} \in K_k^\omega$ implies $\hat{x}_k < \hat{x}_{k_{i'}}$ and $k_{i'+1} \notin K_k^\omega$ implies $r_k^\omega \geq r_{k_{i'+1}}^\omega$. Note that the above still holds if $i' = l$ since we defined $r_{k_{l+1}}^\omega$ to be 0 and $r_k \geq 0$. But (2.13) is simply the objective function of (2.10) evaluated at $T(\omega) = T^*(\omega) \cup \{k\}$. So, starting with an arbitrary optimal solution to (2.10), we can add any checkpoint satisfying either $K_k^\omega = \emptyset$ or (2.11) and maintain optimality. This proves the reverse direction.

Now suppose $T^*(\omega)$ contains all k satisfying either $K_k^\omega = \emptyset$ or (2.11) but also contains some k with $K_k^\omega \neq \emptyset$ which does not satisfy (2.11). It must hold that at least one element of $\operatorname{argmin}_{k' \in K_k^\omega} \hat{x}_{k'}$ is in $T^*(\omega)$. Then there exists $k_i, k_{i+1} \in T^*(\omega)$ such that $\hat{x}_{k_i} \leq \hat{x}_{k_{i+1}}$, and so,

$$\begin{aligned}
z^\omega &= r_{k_1}^\omega - \dots - (r_{k_i}^\omega - r_{k_{i+1}}^\omega) \hat{x}_{k_i} - (r_{k_{i+1}}^\omega - r_{k_{i+2}}^\omega) \hat{x}_{k_{i+1}} - \dots \\
&\leq r_{k_1}^\omega - \dots - (r_{k_i}^\omega - r_{k_{i+1}}^\omega) \hat{x}_{k_i} - (r_{k_{i+1}}^\omega - r_{k_{i+2}}^\omega) \hat{x}_{k_i} - \dots \\
&= r_{k_1}^\omega - \dots - (r_{k_{i+2}}^\omega - r_{k_i}^\omega) \hat{x}_{k_i} - \dots. \tag{2.14}
\end{aligned}$$

But (2.14) is simply the objective function of (2.10) evaluated at $T(\omega) = T^*(\omega) \setminus k_{i+1}$. So we can remove k_{i+1} without decreasing the objective function. But k_{i+1} satisfies neither $K_{k_{i+1}}^\omega = \emptyset$ nor (2.11) since $k_i \in K_{k_{i+1}}^\omega$. If we iteratively

remove all such checkpoints, after a finite number of iterations we obtain an optimal solution which only includes checkpoints satisfying either $K_k^\omega = \emptyset$ or (2.11). \square

Proposition 8 provides an algorithm for generating an optimal solution to (2.10) in polynomial time. In fact, we can sort r_k^ω , $k \in K$, for each $\omega \in \Omega$, prior to beginning the iterative separation process, then we can solve the separation problem for some $\omega \in \Omega$ by making at most $|K|$ comparisons. Defining $\min_{k' \in \emptyset} \hat{x}_{k'}$ to be 1, we can also obtain an analytical form for z^ω in terms of \hat{x} as follows.

Proposition 9. *Let $\omega \in \Omega$ and $\hat{x} \in [0, 1]^{|K|}$ be given. Then the optimal value of (2.10) is given by:*

$$z^\omega = \sum_{k \in K} r_k^\omega \left(\min_{k' \in K_k^\omega} \hat{x}_{k'} - \hat{x}_k \right)^+. \quad (2.15)$$

Proof. Rearranging the terms in the objective function of (2.10) we obtain:

$$z^\omega = \max_{T(\omega) \subseteq K} \sum_{k_i \in T(\omega)} r_{k_i}^\omega (\hat{x}_{k_{i-1}} - \hat{x}_{k_i}), \quad (2.16)$$

where $\hat{x}_{k_0} \equiv 1$. By Proposition 8 we have that we can form an optimal solution $T^*(\omega)$ to (2.10) via the rule $k \in T^*(\omega)$ if and only if either $K_k^\omega = \emptyset$ or $\hat{x}_k < \min_{k' \in K_k^\omega} \hat{x}_{k'}$. For such an optimal set, $T^*(\omega)$, we claim that

$$\hat{x}_{k_{i-1}} = \min_{k' \in K_{k_i}^\omega} \hat{x}_{k'}. \quad (2.17)$$

Equation (2.17) holds for $i = 1$ since $\hat{x}_{k_0} = 1$ and $K_{k_1}^\omega = \emptyset$. Suppose that (2.17) does not hold for some $i > 1$. Then there exists a non-empty set

$K' = \{k' \in K_{k_i}^\omega : \hat{x}_{k'} < \hat{x}_{k_{i-1}}\}$. If any element of K' also lies in $K_{k_{i-1}}^\omega$, then k_{i-1} would not satisfy condition (2.11) and thus would not be included in $T^*(\omega)$. If none of the elements of K' lie in $K_{k_{i-1}}^\omega$, then at least one element of K' would have been included in $T^*(\omega)$ between k_{i-1} and k_i . In either case we have a contradiction, and so (2.17) must hold. Applying (2.17) we obtain:

$$\begin{aligned}
z^\omega &= \sum_{k_i \in T^*(\omega)} r_{k_i}^\omega (\hat{x}_{k_{i-1}} - \hat{x}_{k_i}) \\
&= \sum_{k_i \in T^*(\omega)} r_{k_i}^\omega \left(\min_{k' \in K_{k_i}^\omega} \hat{x}_{k'} - \hat{x}_{k_i} \right) \\
&= \sum_{k \in K} r_k^\omega \left(\min_{k' \in K_k^\omega} \hat{x}_{k'} - \hat{x}_k \right)^+, \tag{2.18}
\end{aligned}$$

for all $\omega \in \Omega$. The equality in (2.18) is due to the fact that $\hat{x}_k \geq \min_{k' \in K_k^\omega} \hat{x}_{k'}$ for any $k \in K \setminus T^*(\omega)$. \square

Corollary 10. *Consider the LP relaxation of model (2.4) in which all of the step inequalities have been added. If $(\hat{x}, \hat{\theta})$ denotes an optimal solution to this linear program, then its optimal value satisfies*

$$\sum_{\omega \in \Omega} p^\omega \hat{\theta}^\omega = \sum_{\omega \in \Omega} \sum_{k \in K} r_k^\omega \left(\min_{k' \in K_k^\omega} \hat{x}_{k'} - \hat{x}_k \right)^+. \tag{2.19}$$

Note that we can say that $\hat{\theta}^\omega = z^\omega$ even if none of the step inequalities are violated for a particular ω since the original inequalities (2.4c) are special cases of the step inequality. That is, if no step inequalities are violated for some $\omega \in \Omega$, z^ω will simply equal the right-hand side of some binding constraint in (2.4c).

2.5 Reformulation

We now describe a polynomially-sized reformulation of BiSNIP, whose LP relaxation is as strong as that of (2.4) with every step inequality added. Our motivation for developing this reformulation is three-fold. First, a compact reformulation may be preferable in terms of ease of implementation. Second, the reformulation reveals some interesting theoretical results which we present in Section 2.6. Third, the reformulation is amenable to a customized branch-and-bound scheme which we describe in Section 2.8.

Let decision variable v_k^ω equal 1 if smuggler ω traverses checkpoint k and 0 otherwise, and let $v_{k_0}^\omega$ equal 1 if all checkpoints are interdicted and 0 otherwise, where $k_0 \notin K$ is an additional dummy index. These variables allow us to explicitly encode the smuggler's preferences as follows:

$$\min_{x, \theta, v} \sum_{\omega \in \Omega} p^\omega \theta^\omega \quad (2.20a)$$

$$\text{s.t.} \quad x \in X \quad (2.20b)$$

$$\theta^\omega = \sum_{k \in K} r_k^\omega v_k^\omega, \quad \omega \in \Omega \quad (2.20c)$$

$$x_k \geq v_{k_0}^\omega + \sum_{k' \in K \setminus \bar{K}_k^\omega} v_{k'}^\omega, \quad k \in K, \quad \omega \in \Omega \quad (2.20d)$$

$$v_{k_0}^\omega + \sum_{k \in K} v_k^\omega = 1, \quad \omega \in \Omega \quad (2.20e)$$

$$0 \leq v_k^\omega \leq 1, \quad k \in K \cup \{k_0\}, \quad \omega \in \Omega, \quad (2.20f)$$

where $\bar{K}_k^\omega = K_k^\omega \cup \{k\}$ and we take $\sum_{k' \in \emptyset} v_{k'}^\omega$ as 0. Constraint (2.20d) says that a smuggler will only traverse a checkpoint ranked lower than checkpoint k

if a detector is installed at checkpoint k . Constraint (2.20e) requires that each smuggler traverses exactly one checkpoint and (2.20c) selects the appropriate evasion probability for each smuggler. We do not enforce integrality constraints on v since, as we show later, there is no incentive to fractionalize v when x is binary.

To see that model (2.20) properly computes the conditional evasion probabilities $\theta^\omega, \omega \in \Omega$, let $x \in X$ be a feasible installation plan and $k^{\omega*} \in \operatorname{argmax}_{\{k: x_k=0\}} r_k^\omega$ be an optimal response for smuggler ω given x . Since $x_{k^{\omega*}} = 0$, by (2.20d) v_k^ω must be 0 for all k ranked lower than $k^{\omega*}$. The variable $v_{k^{\omega*}}^\omega$, along with any v_k^ω corresponding to a k ranked higher than $k^{\omega*}$, is not bound by constraint (2.20d). For each $\omega \in \Omega$, in order to minimize θ^ω and satisfy constraint (2.20e), we choose from amongst the free v_k^ω one of the variables with the smallest r_k^ω coefficient to be 1. So we can choose $v_{k^{\omega*}}^\omega = 1$ which yields the appropriate evasion probability $\theta^\omega = r_{k^{\omega*}}^\omega$. For the trivial case in which $x_k = 1, \forall k \in K$, we can set $v_{k_0}^\omega = 1$ and achieve $\theta^\omega = 0$. In practice, the budget is typically small enough so that it is impossible to install a detector at every checkpoint and in that case we can fix $v_{k_0}^\omega = 0$.

We can informally comment on the strength of (2.20) as follows. Consider the LP relaxation of (2.20), and suppose that $x_k = 0$ for some $k \in K$. Then $v_{k'}^\omega = 0$ for all $k' \in K$ such that $r_{k'}^\omega < r_k^\omega$ by (2.20d) and $\theta^\omega \geq r_k^\omega$, with equality holding only if $v_k^\omega = 1$. But $v_k^\omega = 1$ forces $x_{k'} = 1$ for all k' such that $r_{k'}^\omega > r_k^\omega$, again by (2.20d). Contrast this with the LP relaxation of (2.4), in which we can achieve $\theta^\omega = r_k^\omega$ by setting $x_{k'} = \frac{r_k^\omega - r_{k'}^\omega}{r_k^\omega}$ for every k' such that

$r_{k'}^\omega > r_k^\omega$. We next comment more formally on the strength of (2.20) by first showing that the optimal value of its LP relaxation is equal to that of (2.4) with all step inequalities added.

Lemma 11. *There exists an optimal solution to the LP relaxation of (2.20) $(\hat{x}, \hat{\theta}, \hat{v})$, in which*

$$\hat{v}_k^\omega = \left(1 - \hat{x}_k - \sum_{k' \in K_k^\omega} \hat{v}_{k'}^\omega \right)^+,$$

for all $k \in K, \omega \in \Omega$.

Proof. By constraint (2.20e), we can rewrite constraint (2.20d) as:

$$x_k \geq 1 - v_k^\omega - \sum_{k' \in K_k^\omega} v_{k'}^\omega, \quad k \in K, \omega \in \Omega. \quad (2.21)$$

Rearranging the terms of (2.21) and incorporating the non-negativity constraints on v we obtain:

$$v_k^\omega \geq \left(1 - x_k - \sum_{k' \in K_k^\omega} v_{k'}^\omega \right)^+ \equiv \underline{v}_k^\omega, \quad k \in K, \omega \in \Omega. \quad (2.22)$$

Let $\omega \in \Omega$ and suppose (2.22) is strict for some $k \in K$ at an optimal solution $(\hat{x}, \hat{v}, \hat{\theta})$ to the LP relaxation of (2.20). Let $k_1 \in \operatorname{argmin}_{\{k': \hat{v}_{k'}^\omega > \underline{v}_{k'}^\omega\}} |K_{k'}^\omega|$. And, let $k_2 \in \operatorname{argmin}_{\{k': k_1 \in K_{k'}^\omega\}} |K_{k'}^\omega|$ if $K_{k_1}^\omega \cup \{k_1\} \neq K$ and $k_2 = k_0$ otherwise. Now we construct a perturbed solution identical to $(\hat{x}, \hat{v}, \hat{\theta})$ but with $v_{k_1}^\omega = \hat{v}_{k_1}^\omega - \epsilon$, $v_{k_2}^\omega = \hat{v}_{k_2}^\omega + \epsilon$, and $\theta^\omega = \hat{\theta}^\omega - \epsilon(r_{k_1}^\omega - r_{k_2}^\omega)$, where $\epsilon = \hat{v}_{k_1}^\omega - \underline{v}_{k_1}^\omega$ and $r_{k_0}^\omega \equiv 0$. To show that this perturbed solution satisfies (2.20d), we equivalently show that it satisfies (2.22). For $k \in K \setminus \{k_1, k_2\}$, both sides of (2.22) remain unchanged

since for all such k , either $k_1, k_2 \in K_k^\omega$ or $k_1, k_2 \notin K_k^\omega$. (If $k_2 = k_0$ the latter is always true since k_1 must be the smuggler's lowest ranked checkpoint and $k_0 \notin K$.) Inequality (2.22) remains satisfied for $k = k_1$ since the right-hand side remains unchanged and $\epsilon = \hat{v}_{k_1}^\omega - \underline{v}_k^\omega$. Finally, for $k = k_2 \neq k_0$, both sides of (2.22) increase by ϵ . So inequality (2.22) holds which implies (2.20d) holds as well. The rest of the constraints in (2.20) are also satisfied by construction. This perturbed solution is optimal since $r_{k_1}^\omega \geq r_{k_2}^\omega$ and has the property that (2.22) is tight for all $k' \in K_{k_1}^\omega \cup \{k_1\}$. Repeating this perturbation at most $|K||\Omega|$ times we obtain an optimal solution for which (2.22) is tight for all $k \in K$. \square

Theorem 12. *The optimal value of the LP relaxation of (2.4) with all step inequalities added is equal to that of the LP relaxation of (2.20).*

Proof. By Proposition 9, if we fix $\hat{x} \in [0, 1]^{|K|}$ in (2.4), add the most violated step inequalities for each scenario, and optimize over θ we obtain:

$$\theta^\omega = z^\omega = \sum_{k \in K} r_k^\omega \left(\min_{k' \in K_k^\omega} \hat{x}_{k'} - \hat{x}_k \right)^+,$$

where z^ω is the largest right-hand side of all step inequalities for scenario ω . Since the objective function of (2.4) is $\sum_{\omega \in \Omega} p^\omega \theta^\omega$ and is identical to that of (2.20), it suffices to show that $\hat{v}_k^\omega = (\min_{k' \in K_k^\omega} \hat{x}_{k'} - \hat{x}_k)^+$ at an optimal solution to the LP relaxation of (2.20) for the same fixed $x = \hat{x}$.

Let $\omega \in \Omega$ and assume without loss of generality that $K_k^\omega = \{k' \in K : k' < k\}$. Then under the indexing for this particular ω , $\min_{k' \in K_k^\omega} \hat{x}_{k'} =$

$\min_{k' < k} \hat{x}_{k'}$. By induction we now show that:

$$\min_{k' < k+1} x_{k'} = x_1 - (x_1 - x_2)^+ - \dots - \left(\min_{k' < k} x_{k'} - x_k \right)^+ \quad (2.23a)$$

$$\hat{v}_k^\omega = \left(\min_{k' < k} \hat{x}_{k'} - \hat{x}_k \right)^+ \quad (2.23b)$$

holds for all $k \geq 2$. For the base case $k = 2$ we have:

$$\min(x_1, x_2) = x_1 - (x_1 - x_2)^+ \quad (2.24a)$$

$$\hat{v}_2^\omega = (\hat{x}_1 - \hat{x}_2)^+. \quad (2.24b)$$

Equation (2.24a) holds since $(x_1 - x_2)^+ = (x_1 - x_2)$ if $x_1 > x_2$ and $(x_1 - x_2)^+ = 0$ otherwise. Equation (2.24b) holds since by Lemma 11, $\hat{v}_1^\omega = 1 - \hat{x}_1$ and $\hat{v}_2^\omega = (1 - \hat{x}_2 - \hat{v}_1^\omega)^+ = (\hat{x}_1 - \hat{x}_2)^+$. Now assume that (2.23a) holds for an arbitrary $k \geq 2$. Then:

$$\begin{aligned} \min_{k' < k+1} x_{k'} &= \min_{k' < k+1} x_{k'} - \left(\min_{k' < k+1} x_{k'} - x_{k+1} \right)^+ \\ &= x_1 - (x_1 - x_2)^+ - \dots - \left(\min_{k' < k+1} x_{k'} - x_{k+1} \right)^+ \end{aligned}$$

by the same argument we made for the base case (2.24a). Thus (2.23a) holds for all $k \geq 2$. Finally, assume (2.23b) holds for an arbitrary $k \geq 2$. Lemma 11 then yields:

$$\begin{aligned} \hat{v}_{k+1}^\omega &= (1 - \hat{x}_{k+1} - \hat{v}_1^\omega - \hat{v}_2^\omega - \dots - \hat{v}_k^\omega)^+ \\ &= \left(1 - \hat{x}_{k+1} - (1 - \hat{x}_1) - (\hat{x}_1 - \hat{x}_2)^+ - \dots - \left(\min_{k' < k} \hat{x}_{k'} - \hat{x}_k \right)^+ \right)^+ \\ &= \left(\min_{k' < k+1} \hat{x}_{k'} - \hat{x}_k \right)^+. \end{aligned}$$

□

Finally, we show that the non-dominated auxiliary variables of (2.20) correspond exactly to the extreme points of the convex hulls of the polyhedra induced by constraints (2.4c) for each $\omega \in \Omega$.

Proposition 13. *Let $\Theta^\omega = \{(x, \theta^\omega) : \theta^\omega \geq r_k^\omega(1 - x_k), k \in K, \theta^\omega \in \mathbb{R}_+, x \in \mathbb{Z}_+^{|K|}\}$, where $0 \leq r_k^\omega \leq 1$, and let $(\hat{x}, \hat{\theta}^\omega)$ be an extreme point of the convex hull of Θ^ω . Then either $\hat{\theta}^\omega = r_k^\omega$ for some $k \in K$ or $\hat{\theta}^\omega = 0$. Moreover, $\hat{x}_k = 1$ if $r_k^\omega > \hat{\theta}^\omega$ and $\hat{x}_k = 0$ otherwise.*

Proof. Suppose that $(\hat{x}, \hat{\theta}^\omega)$ is an extreme point of $\text{conv}(\Theta^\omega)$. Then

$$\hat{\theta}^\omega \geq r_{\max}^\omega \equiv \max(\max_{k \in K} r_k^\omega(1 - \hat{x}_k), 0). \quad (2.25)$$

Now suppose that $\hat{\theta}^\omega > r_{\max}^\omega$. Then the points $(\hat{x}, \hat{\theta}^\omega + \epsilon)$ and $(\hat{x}, \hat{\theta}^\omega - \epsilon)$ where $\epsilon = \hat{\theta}^\omega - r_{\max}^\omega$ are both in Θ^ω . But $(\hat{x}, \hat{\theta}^\omega)$ is a convex combination of these points and thus cannot be an extreme point of $\text{conv}(\Theta^\omega)$. So $\hat{\theta}^\omega = r_{\max}^\omega$ and consequently $\hat{\theta}^\omega$ must be either r_k^ω for some $k \in K$ or 0. This proves the first claim.

To prove the second claim, note that if $(\hat{x}, \hat{\theta}^\omega) \in \Theta^\omega$, then $\hat{x}_k \geq 1$ for every k with $r_k^\omega > \hat{\theta}^\omega$. Now suppose that for some $k' \in K$, either $x_{k'} \geq 2$ and $r_{k'}^\omega > \hat{\theta}^\omega$ or $x_{k'} \geq 1$ and $r_{k'}^\omega \leq \hat{\theta}^\omega$. Then $(\hat{x} + e_{k'}, \hat{\theta}^\omega)$ and $(\hat{x} - e_{k'}, \hat{\theta}^\omega)$, where e_k is the unit vector with the k th component equal to 1, are both in Θ^ω . Since $(\hat{x}, \hat{\theta}^\omega)$ is a convex combination of these two points, we must have that $\hat{x}_k = 1$ for all k with $r_k^\omega > \hat{\theta}^\omega$ and $\hat{x}_k = 0$ otherwise. \square

The significance of Proposition 13 is that the convex hull of Θ^ω has at most $|K| + 1$ extreme points. We can also easily enumerate the extreme directions of $\text{conv}(\Theta^\omega)$; they are simply $(e_k, 0), k \in K$, and $(0, 1)$. Writing $\text{conv}(\Theta^\omega)$ as a convex combination of its extreme points and a non-negatively weighted linear combination of its extreme directions results in constraints (2.20c)-(2.20f) above. If some of the r_k^ω are equal to each other or 0, we have fewer than $|K| + 1$ extreme points since each extreme point corresponds to a unique value of $\hat{\theta}^\omega$. This may seem to be a discrepancy since we always have $|K| + 1$ auxiliary variables per scenario, but can be explained as follows. If multiple v_k^ω variables have the same coefficient in (2.20c), the variable corresponding to the smallest $|K_k^\omega|$ will dominate, and only the dominating variable corresponds to an extreme point. So the strength of the LP relaxation of (2.20) lies in that it does not allow points (x, θ) feasible to the LP relaxation of (2.4) but for which (x, θ^ω) does not lie in $\text{conv}(\Theta^\omega)$. See [24] for a survey of tight formulations for mixed-integer sets similar to that of BiSNIP.

We conclude this section with a transformed version of (2.20) which has the same LP relaxation value but a sparser constraint matrix. We introduce a new decision variable u_k^ω which equals 1 if smuggler ω traverses a checkpoint with a lower evasion probability than that of checkpoint k and 0 otherwise. Recall that $k(i, \omega)$ is smuggler ω 's i th best checkpoint. Then we can relate v to u as follows

$$v_{k(i, \omega)}^\omega = u_{k(i-1, \omega)}^\omega - u_{k(i, \omega)}^\omega, \quad (2.26)$$

where $u_{k(0, \omega)}^\omega \equiv 1$ and $v_{k_0}^\omega = u_{k(|K|, \omega)}^\omega$, $\omega \in \Omega$. We can now replace v with u in

model (2.20) starting with constraint (2.20d) as follows

$$\begin{aligned}
x_{k(i,\omega)} &\geq v_{k_0}^\omega + \sum_{i'=i+1}^{|K|} v_{k(i',\omega)}^\omega \\
&= u_{k(|K|,\omega)}^\omega + \sum_{i'=i+1}^{|K|} (u_{k(i'-1,\omega)}^\omega - u_{k(i',\omega)}^\omega) \\
&= u_{k(i,\omega)}^\omega.
\end{aligned}$$

Constraint (2.20e) is satisfied automatically since $v_{k_0}^\omega + \sum_{k \in K} v_k^\omega = u_{k(0,\omega)}^\omega = 1$.

Non-negativity constraints on v simply translate into

$$u_{k(i-1,\omega)}^\omega \geq u_{k(i,\omega)}^\omega, i = 1, \dots, |K|, \omega \in \Omega.$$

Finally, we can substitute out v from constraint (2.20c) as follows

$$\begin{aligned}
\theta^\omega &= \sum_{i=1}^{|K|} r_{k(i,\omega)}^\omega (u_{k(i-1,\omega)}^\omega - u_{k(i,\omega)}^\omega) \\
&= r_{k(1,\omega)}^\omega + \sum_{i=1}^{|K|} (r_{k(i+1,\omega)}^\omega - r_{k(i,\omega)}^\omega) u_{k(i,\omega)}^\omega,
\end{aligned}$$

where $r_{k(|K|+1,\omega)}^\omega \equiv 0$. With

$$s_{k(i,\omega)}^\omega \equiv r_{k(i,\omega)}^\omega - r_{k(i+1,\omega)}^\omega,$$

minimizing $\sum_{\omega \in \Omega} p^\omega \theta^\omega$ is equivalent to maximizing $\sum_{\omega \in \Omega} \sum_{k \in K} p^\omega s_k^\omega u_k^\omega$. Thus,

we arrive at the following model:

$$\max_{x,u} \quad \sum_{\omega \in \Omega} \sum_{k \in K} p^\omega s_k^\omega u_k^\omega \quad (2.29a)$$

$$\text{s.t.} \quad x \in X \quad (2.29b)$$

$$x_k \geq u_k^\omega, \quad k \in K, \omega \in \Omega \quad (2.29c)$$

$$u_{k(i-1,\omega)}^\omega \geq u_{k(i,\omega)}^\omega, \quad i = 2, \dots, |K|, \omega \in \Omega \quad (2.29d)$$

$$0 \leq u_k^\omega \leq 1, \quad k \in K, \omega \in \Omega. \quad (2.29e)$$

We do not include constraint (2.29d) for $i = 1$ since that would simply reduce to a simple upper bound $1 \geq u_{k(1,\omega)}^\omega$. Note that this model is equivalent to (2.20) since given v , u can be uniquely determined and vice-versa. While model (2.29) has roughly twice as many structural constraints as (2.20), it has a sparser constraint matrix for moderate- and large-scale instances since every constraint other than the budget constraint has only two non-zero terms. We focus our attention on model (2.29) for the remainder of this chapter. An advantage to taking the perspective of model (2.29) is that it is identical to, minus the budget constraint, the shared fixed cost problem introduced by Rhys [32].

2.6 Efficient Nested Solutions

BiSNIP is strongly NP-complete, even with unit interdiction costs and equally likely threat scenarios [29]. We now describe a family of instances of BiSNIP that can be solved in polynomial time by solving a modest number of linear programs. Given a probability mass function over the scenarios and an

evasion probability for every scenario-checkpoint pair, we can plot the maximum decrease in evasion probability, relative to the uninterdicted system, as a function of the budget by solving a BiSNIP instance for every possible budget level. We refer to this curve as the *efficient frontier* of solutions, provided it increases strictly. In this section we show that the solutions corresponding to extreme points of the concave envelope of the efficient frontier can be found in polynomial time (see Figure 2.2). And, given a pair of solutions which both correspond to extreme points, the checkpoints interdicted in one solution are a subset of those interdicted in the other. We refer to this as the *nestedness* property. Hochbaum [17] shows that this property holds for a related problem by exploiting the solution properties of an equivalent parametric max-flow problem. In this section, we show that the nestedness property hinges on a supermodularity property of the objective function and a submodularity property of the cost function. In general, solutions to BiSNIP are not necessarily nested; see for example [23].

To elaborate further, consider a biobjective integer program in which one objective is to maximize the decrease in the evasion probability, and the other is to minimize detector installation costs. A solution for which it is both impossible to decrease the expected evasion probability without increasing installation costs and impossible to decrease installation costs without increasing the expected evasion probability is said to be *Pareto efficient*. Kuhn and Tucker [20] show that we can find some Pareto efficient solutions by solving a single-objective program whose objective is a weighted sum of the objec-

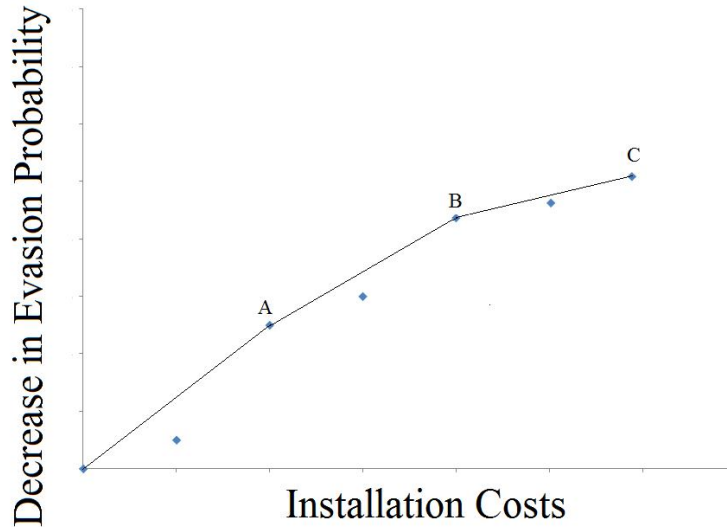


Figure 2.2: An example efficient frontier of solutions. Solutions corresponding to points A , B , and C can be found in polynomial time and are nested.

tives of the biobjective program. More specifically, every solution to this single-objective program corresponds to a point on the concave envelope of the efficient frontier [31].

For our problem, the resulting single-objective program is simply the Lagrangian relaxation of (2.29) with the budget constraint dualized, which we now show has an LP relaxation with integral extreme points. We first present a result linking \hat{u} to \hat{x} in an extreme point solution.

Lemma 14. *Consider the LP relaxation of model (2.29) with the budget constraint $\sum_{k \in K} c_k x_k \leq b$ relaxed and let (\hat{x}, \hat{u}) be an extreme point solution.*

Then if $\hat{u}_k^\omega > 0$,

$$\hat{u}_k^\omega = \min_{k' \in \bar{K}_k^\omega} \hat{x}_{k'}.$$

Proof. By constraints (2.29c) and (2.29d) we have that

$$\hat{u}_{k(i,\omega)}^\omega \leq \min(\hat{x}_{k(i,\omega)}, \hat{u}_{k(i-1,\omega)}^\omega) = \min_{i' \leq i} \hat{x}_{k(i',\omega)}, \quad i = 1, \dots, |K|, \quad \omega \in \Omega. \quad (2.30)$$

Equivalently, we have $\hat{u}_k^\omega \leq \min_{k' \in \bar{K}_k^\omega} \hat{x}_{k'}$ for all $k \in K$ and $\omega \in \Omega$. If this inequality is strict for some $k \in K$ and $\omega \in \Omega$ and $\hat{u}_k^\omega > 0$, then we can form two points feasible to the LP relaxation of (2.29) by perturbing \hat{u}_k^ω by $\epsilon = \min\left(\min_{k' \in \bar{K}_k^\omega} \hat{x}_{k'} - \hat{u}_k^\omega, \hat{u}_k^\omega\right)$. Since (\hat{x}, \hat{u}) can be written as a strict convex combination of these two points, it cannot be an extreme point. \square

Note that the above argument also holds if the budget constraint is not relaxed, a fact we use later.

Proposition 15. *Let XU be the feasible region of the LP relaxation of (2.29) with the budget constraint $\sum_{k \in K} c_k x_k \leq b$ relaxed. Then every extreme point of XU is integer valued.*

Proof. Suppose not. Then there exists $(\hat{x}, \hat{u}) \in XU$ with some component of \hat{x} or \hat{u} being fractional. By Lemma 14, if any component of \hat{u} is fractional then some component of \hat{x} is also fractional. Therefore, there exists a non-empty $K' \subseteq K$ such that \hat{x}_k is fractional for every $k \in K'$, and the following pair of points is feasible:

$$\begin{aligned} x_k &= \hat{x}_k + \epsilon I(k \in K'), \quad k \in K; \quad u_k^\omega = \hat{u}_k^\omega + \epsilon I(0 < \hat{u}_k^\omega < 1), \quad k \in K, \quad \omega \in \Omega \\ x_k &= \hat{x}_k - \epsilon I(k \in K'), \quad k \in K; \quad u_k^\omega = \hat{u}_k^\omega - \epsilon I(0 < \hat{u}_k^\omega < 1), \quad k \in K, \quad \omega \in \Omega, \end{aligned}$$

where $\epsilon = \min_{k \in K'}(\hat{x}_k, 1 - \hat{x}_k)$ and $I(\cdot)$ is the indicator function. Since (\hat{x}, \hat{u}) can be written as a strict convex combination of these two points, we have a contradiction. \square

It is also possible to prove Proposition 15 using the fact that the associated constraint matrix is totally unimodular since every element in the matrix is 0, 1, or -1, every row has at most two non-zero entries, and every row with two non-zero entries sums to 0. The significance of Proposition 15 is that we can solve an instance of BiSNIP in polynomial time if the solution to that instance corresponds to an extreme point of the concave envelope of the efficient frontier. We can find all such points by solving a parameterized LP. Alternatively, the Lagrangian relaxation can be cast as a parameterized min-cut problem [2] and solved efficiently using the push-relabel algorithm [15, 16] or the pseudoflow algorithm [17]. Note that in general there may exist points on the efficient frontier that lie below the concave envelope, so this does not contradict the fact that BiSNIP is NP-complete. We refer to such points as being *convex dominated* [31]. We formally describe this notion in the context of our problem as follows.

Definition 1. Let $K' \subseteq K$. We define the gain function $g : 2^K \rightarrow \mathbb{R}$ as:

$$g(K') = \sum_{(\omega, k): \bar{K}_k^\omega \subseteq K'} p^\omega s_k^\omega.$$

Definition 2. Let $K' \subseteq K$. We define the cost function $c : 2^K \rightarrow \mathbb{R}$ as:

$$c(K') = \sum_{k \in K'} c_k.$$

Definition 3. Let $\hat{K} \subseteq K$. Then \hat{K} is convex dominated if there exists $\hat{K}^-, \hat{K}^+ \subseteq K$ with $c(\hat{K}^-) \leq c(\hat{K}) \leq c(\hat{K}^+)$ and $c(\hat{K}^-) \neq c(\hat{K}^+)$ such that:

$$g(\hat{K}) < \alpha g(\hat{K}^-) + (1 - \alpha)g(\hat{K}^+),$$

where $\alpha = \frac{c(\hat{K}^+) - c(\hat{K})}{c(\hat{K}^+) - c(\hat{K}^-)}$.

Any solution on the concave envelope of the efficient frontier corresponds to a set of checkpoints \hat{K} which satisfies $g(\hat{K}) \geq \alpha g(\hat{K}^-) + (1 - \alpha)g(\hat{K}^+)$ for all \hat{K}^-, \hat{K}^+ satisfying the conditions of Definition 3, and for an extreme point of the efficient frontier's convex hull, this inequality is strict unless $c(\hat{K}) = c(\hat{K}^-)$ or $c(\hat{K}) = c(\hat{K}^+)$.

Before we prove the nestedness property, we present some useful properties of the gain function and the efficient frontier. First, we show that the gain function is supermodular (see, e.g., [34]).

Proposition 16. $g(\cdot)$ is a supermodular function. That is, if $A \subset B \subseteq K$ and $k' \in A$, then

$$g(A) - g(A \setminus \{k'\}) \leq g(B) - g(B \setminus \{k'\}).$$

Proof. For any $K' \subseteq K$ such that $k' \in K'$ we have:

$$\begin{aligned} g(K') - g(K' \setminus \{k'\}) &= \sum_{(\omega, k): \bar{K}_k^\omega \subseteq K'} p^\omega s_k^\omega - \sum_{(\omega, k): \bar{K}_k^\omega \subseteq K' \setminus \{k'\}} p^\omega s_k^\omega \\ &= \sum_{(\omega, k) \in G(K', k')} p^\omega s_k^\omega, \end{aligned}$$

where $G(K', k') = \{(\omega, k) : \bar{K}_k^\omega \subseteq K', k' \in \bar{K}_k^\omega\}$. But $G(A, k') \subseteq G(B, k')$ coupled with non-negativity of the gain function's summand yields the desired result. \square

Definition 4. Let $A, B \subseteq K$ satisfy $c(A) < c(B)$. We define the gain-to-cost ratio $m : 2^K \times 2^K \rightarrow \mathbb{R}$ as:

$$m(A, B) = \frac{g(B) - g(A)}{c(B) - c(A)}. \quad (2.32)$$

Next, we present some useful results regarding sets that are not convex dominated.

Lemma 17. Let $A, B, C \subseteq K$ satisfy $c(A) < c(B) < c(C)$. Then

- (a) $m(A, C) = \alpha m(A, B) + (1 - \alpha)m(B, C)$ for some $\alpha \in (0, 1)$;
 - (b) if A is not convex dominated then $m(A, B) = \min_{\{K' \subseteq K : c(K') \leq c(A)\}} m(K', B)$;
 - (c) if B is not convex dominated then $m(A, B) = \max_{\{K' \subseteq K : c(K') \geq c(B)\}} m(A, K')$;
- and,
- (d) if B is not convex dominated and $m(A, C) = m(A, B) = m(B, C)$, then neither A nor C is convex dominated.

Proof. With $\alpha = \frac{c(B) - c(A)}{c(C) - c(A)}$, part (a) holds immediately. We prove part (b) by contradiction. Suppose that $m(A, B) > \min_{\{K' \subseteq K : c(K') \leq c(A)\}} m(K', B)$. Then there exists $K^* \subseteq K$ with $c(K^*) \leq c(A)$ which satisfies $m(A, B) > m(K^*, B)$.

If we let $\alpha = \frac{c(B)-c(A)}{c(B)-c(K^*)}$, then A is convex dominated by K^* and B as follows:

$$\begin{aligned}\frac{g(B) - g(A)}{c(B) - c(A)} &> \frac{g(B) - g(K^*)}{c(B) - c(K^*)} \implies \\ g(B) - g(A) &> \alpha(g(B) - g(K^*)) \implies \\ \alpha g(K^*) + (1 - \alpha)g(B) &> g(A).\end{aligned}$$

Similarly we can prove part (c) by contradiction. Suppose $m(A, B) < m(A, K^*)$ for some K^* satisfying $c(K^*) \geq c(B)$. If we let $\alpha = \frac{c(K^*)-c(B)}{c(K^*)-c(A)}$, then B is convex dominated by A and K^* as follows:

$$\begin{aligned}\frac{g(B) - g(A)}{c(B) - c(A)} &< \frac{g(K^*) - g(A)}{c(K^*) - c(A)} \implies \\ g(B) - g(A) &< (1 - \alpha)(g(K^*) - g(A)) \implies \\ g(B) &< \alpha g(A) + (1 - \alpha)g(K^*).\end{aligned}$$

Finally, to prove (d) assume that C is convex dominated and satisfies $m(A, C) = m(A, B) = m(B, C)$. Then there exists $C^-, C^+ \subseteq K$ with $c(C^-) \leq c(C) \leq c(C^+)$ and $c(C^-) \neq c(C^+)$ which satisfy

$$g(C) < \alpha g(C^-) + (1 - \alpha)g(C^+),$$

where $\alpha = \frac{c(C^+)-c(C)}{c(C^+)-c(C^-)}$. But if $c(C) = c(C^-)$, then $g(C^-) > g(C)$ and $m(A, C^-) > m(A, C) = m(A, B)$, contradicting part (c). If $c(C) = c(C^+)$ we would have $m(A, C^+) > m(A, B)$ again contradicting part (c). So we can assume $c(C^-) < c(C) < c(C^+)$. Then we have

$$\begin{aligned}\alpha(g(C^+) - g(C^-)) &< g(C^+) - g(C) \implies \\ \frac{g(C^+) - g(C^-)}{c(C^+) - c(C^-)} &< \frac{g(C^+) - g(C)}{c(C^+) - c(C)}.\end{aligned}$$

So $m(C^-, C^+) < m(C, C^+)$ and, by part (a), $m(C^-, C) < m(C, C^+)$. But by part (c), since $c(C^+) > c(B)$ we must have $m(A, C^+) \leq m(A, B) = m(A, C)$, which by part (a) implies $m(C, C^+) \leq m(A, C)$. To summarize we have

$$m(C^-, C) < m(C, C^+) \leq m(A, C) = m(A, B) = m(B, C).$$

Now suppose $c(C^-) < c(B)$. Since B is not convex dominated and $m(C^-, C) < m(B, C)$, this contradicts part (b). So we must have $c(C^-) \geq c(B)$. Since $m(C^-, C) < m(A, C)$, by part (a) we have $m(A, C) < m(A, C^-)$. By hypothesis this implies $m(A, B) < m(A, C^-)$ which is a contradiction of part (c). Similar logic shows that A is also not convex dominated. \square

Of particular interest is part (c) of Lemma 17, which states that a set of checkpoints which is not convex dominated maximizes the gain-to-cost ratio between the set and any other set with smaller cost. In other words, the solutions that we can find in polynomial time conveniently happen to be those that maximize the decrease in evasion probability per unit installation cost. Additionally, we can show that solutions that correspond to corner points of the concave envelope are nested, that is, any pair of corner point solutions corresponds to a pair of sets of checkpoints one of whom is a subset of the other. This can be a desirable property if funds for installing detectors are made available over time [17, 23]. The nestedness property is a direct consequence of the following.

Theorem 18. *Let $A, B \subseteq K$. Then if neither A nor B is convex dominated, then neither AB nor $A \cup B$ is convex dominated.*

Proof. Assume without loss of generality that $c(A) \leq c(B)$. If $A \subseteq B$ then $AB = A$ and $A \cup B = B$ and the result holds by hypothesis. Suppose then that $A \not\subseteq B$. Then we have $c(AB) < c(A) < c(A \cup B)$. By Proposition 16 we have that $g(A \cup B) - g(B) \geq g(A) - g(AB)$, and since $c(A \cup B) - c(B) = c(A) - c(AB) = c(A \setminus B)$, this implies that $m(B, A \cup B) \geq m(AB, A)$. Also, by Lemma 17(b) we have $m(AB, A \cup B) \geq m(A, A \cup B)$ which implies $m(AB, A) \geq m(A, A \cup B)$ by Lemma 17(a). Combining these results we have:

$$m(B, A \cup B) \geq m(AB, A) \geq m(A, A \cup B). \quad (2.33)$$

But by Lemma 17(b), we have $m(B, A \cup B) \leq m(A, A \cup B)$ and so (2.33) must hold with equality throughout. So $m(AB, A) = m(A, A \cup B)$ and by Lemma 17(d), neither AB nor $A \cup B$ is convex dominated. \square

Corollary 19. *Let $K_1, K_2 \subseteq K$ correspond to points on the concave envelope of the efficient frontier with $c(K_1) < c(K_2)$. If either K_1 or K_2 corresponds to an extreme point, then $K_1 \subset K_2$. Moreover, if both K_1 and K_2 correspond to extreme points and $c(K_1) = c(K_2)$, then $K_1 = K_2$.*

Corollary 19 follows since $g(K_i) = \alpha_i g(K_1 K_2) + (1 - \alpha_i) g(K_1 \cup K_2)$, $i = 1, 2$, for $\alpha_i = \frac{c(K_1 \cup K_2) - c(K_i)}{c(K_1 \cup K_2) - c(K_1 K_2)}$, which implies that K_1 and K_2 do not correspond to extreme points unless $K_1 = K_1 K_2$ and $K_2 = K_1 \cup K_2$. The second part of the corollary implies that each extreme point corresponds to a unique set of checkpoints.

We conclude this section by proposing an easy-to-implement algorithm for generating all extreme points, based on the following observation.

Proposition 20. *Let (\hat{x}, \hat{u}) be an extreme point solution to the LP relaxation of (2.29). Then all fractional components of \hat{x} must be equal.*

Proof. We assume without loss of generality that $c_k \geq 1, \forall k$. Suppose there exists $k_1, k_2 \in K$ such that \hat{x}_{k_1} and \hat{x}_{k_2} are both fractional and $\hat{x}_{k_1} \neq \hat{x}_{k_2}$. Let $K_1 = \{k \in K : \hat{x}_k = \hat{x}_{k_1}\}$ and $K_2 = \{k \in K : \hat{x}_k = \hat{x}_{k_2}\}$. By Lemma 14 we have that if $\hat{u}_k^\omega > 0$ then $\hat{u}_k^\omega = \min_{k' \in \bar{K}_k^\omega} x_{k'}$. Let

$$\epsilon = \min \left\{ \frac{1}{2} \min_{k' \in K \setminus K_1} |\hat{x}_{k_1} - \hat{x}_{k'}|, \frac{1}{2} \min_{k' \in K \setminus K_2} |\hat{x}_{k_2} - \hat{x}_{k'}|, \hat{x}_{k_1}, 1 - \hat{x}_{k_1}, \hat{x}_{k_2}, 1 - \hat{x}_{k_2} \right\}.$$

The following two points are feasible to the LP relaxation of (2.29):

$$\begin{aligned} x_k &= \hat{x}_k + \frac{\epsilon}{c(K_1)} I(k \in K_1) - \frac{\epsilon}{c(K_2)} I(k \in K_2), \quad k \in K; \\ u_k^\omega &= \hat{u}_k^\omega + \frac{\epsilon}{c(K_1)} I(\hat{u}_k^\omega = \hat{x}_{k_1}) - \frac{\epsilon}{c(K_2)} I(\hat{u}_k^\omega = \hat{x}_{k_2}), \quad k \in K, \omega \in \Omega \\ x_k &= \hat{x}_k - \frac{\epsilon}{c(K_1)} I(k \in K_1) + \frac{\epsilon}{c(K_2)} I(k \in K_2), \quad k \in K; \\ u_k^\omega &= \hat{u}_k^\omega - \frac{\epsilon}{c(K_1)} I(\hat{u}_k^\omega = \hat{x}_{k_1}) + \frac{\epsilon}{c(K_2)} I(\hat{u}_k^\omega = \hat{x}_{k_2}), \quad k \in K, \omega \in \Omega. \end{aligned}$$

Since (\hat{x}, \hat{u}) can be written as a strict convex combination of these two points, it is not an extreme point. \square

The algorithm to generate the solutions corresponding to all extreme points of the efficient frontier's convex envelope proceeds as follows. Assume without loss of generality that $c_k \geq 1, \forall k$, and let (\hat{x}, \hat{u}) be a solution to the LP relaxation of (2.29) with $f = 1$. If $K^* \subseteq K$ indexes the positive components of \hat{x} and $K^* \neq \emptyset$, then $\sum_{k \in K^*} c_k \hat{x}_k = 1$ and $\hat{x}_k = I(k \in K^*)/c(K^*), \forall k \in K$ by Proposition 20. So the objective value is simply $g(K^*)/c(K^*)$ and therefore

K^* represents a Pareto-efficient solution which is not convex dominated as it must solve $\min_{K' \subseteq K} g(K')/c(K')$. By Corollary 19 and Lemma 17(c), we can then resolve the LP relaxation of (2.29) with x_k fixed to 1 for all $k \in K^*$ and with $f = c(K^*) + 1$ to generate the next extreme point. We can generate all extreme points by iterating in this fashion until we arrive at a solution in which all checkpoints the smuggler would consider traversing are interdicted. Note that this algorithm may also generate some points that are on the concave upper envelope but not extreme points. This algorithm is summarized in the pseudo-code of Algorithm 1.

Algorithm 1: *GetExtremePoints*(p, r, c)

Input: Scenario probabilities $p^\omega > 0$, evasion probabilities $r_k^\omega \geq 0$, detector installation costs $c_k \geq 1$

Output: Sets $K_1, \dots, K_n \subseteq K$ corresponding to all extreme points on the concave envelope of the efficient frontier

$K_1 \leftarrow \emptyset$

$n \leftarrow 1$

loop

Solve the LP relaxation of (2.29) with $b = \sum_{k \in K_n} c_k + 1$ and the added constraints $x_k = 1, k \in K_n$ and let (\hat{x}, \hat{u}) be the optimal solution

if $K_n = \{k \in K : \hat{x}_k > 0\}$ **then**

break

end if

Let $K_{n+1} = \{k \in K : \hat{x}_k > 0\}$

$n \leftarrow n + 1$

if $K_n == K$ **then**

break

end if

end loop

return K_1, \dots, K_n

2.7 Preprocessing

Model (2.29) has an additional $|K||\Omega|$ decision variables not present in (2.4). While the strengthened LP relaxation more than compensates for the extra variables, we can possibly reduce the size of model (2.29) in two ways. Our first approach uses the fact that the budget constraint implies that some of the auxiliary variables must be zero and the second is a generalization of the scenario aggregation scheme described previously.

Proposition 21. *Let $\bar{K}_k^\omega = \{k' : r_{k'}^\omega \geq r_k^\omega\}$ and suppose $\sum_{k' \in \bar{K}_k^\omega} c_{k'} > b$ for some $k \in K$ and $\omega \in \Omega$. Then $u_k^\omega = 0$ in any feasible solution to (2.29).*

Proof. Suppose $u_k^\omega > 0$ and $\sum_{k' \in \bar{K}_k^\omega} c_{k'} > b$ for some $k \in K$ and $\omega \in \Omega$. Then constraints (2.29c) and (2.29d), coupled with the integrality of x , imply that $x_{k'} = 1$ for all $k' \in \bar{K}_k^\omega$. This is inconsistent with constraint (2.29b) since $\sum_{k \in K} c_k x_k \geq \sum_{k' \in \bar{K}_k^\omega} c_{k'} x_{k'} > b$, and so u_k^ω must be 0 in any feasible solution. \square

Proposition 21 makes use of the fact that smuggler ω traverses a checkpoint with a lower evasion probability than that of k only if we interdict all checkpoints with an evasion probability at least as high as that of k . If interdicting all such checkpoints consumes more budget than we have, then we can fix the corresponding u_k^ω to zero. This simple observation can both reduce the number of variables and greatly strengthen the LP relaxation, especially if the budget is small relative to the number of checkpoints. To see why the LP relaxation may be tightened, consider the following example.

Example 1. Let $|\Omega| = 1$ and $|K| = 3$ with $c_1 = c_2 = c_3 = 1$ and $r_1 = 1$, $r_2 = 0.9$, $r_3 = 0$. If $b = 1$, then (2.29) is:

$$\begin{aligned}
\max_{x,u} \quad & 0.1u_1 + 0.9u_2 \\
\text{s.t.} \quad & u_1 \geq u_2 \\
& x_1 \geq u_1 \\
& x_2 \geq u_2 \\
& x_1 + x_2 + x_3 \leq 1 \\
& x_1, x_2, x_3 \in \{0, 1\} \\
& 0 \leq u \leq 1.
\end{aligned} \tag{2.35}$$

The optimal solution of (2.35) is $x^* = (1, 0, 0)$, $u^* = (1, 0, 0)$ which gives a decrease in the evasion probability of 0.1. But the optimal solution to the LP relaxation of (2.35) is $x^{LP} = (0.5, 0.5, 0)$, $u^{LP} = (0.5, 0.5, 0)$, giving a decrease in evasion probability of 0.5. However, $u_2 > 0$ forces $x_1 = x_2 = 1$, and so we can fix $u_2 = 0$ and the optimal solution to the LP relaxation is now integer feasible.

We now show that we can aggregate some of the auxiliary variables that remain if a pair of smugglers ranks their checkpoints similarly.

Proposition 22. Consider model (2.29) and suppose that $\bar{K}_k^{\omega_1} = \bar{K}_k^{\omega_2}$ for some $k \in K$ and $\omega_1, \omega_2 \in \Omega$. Then there is an optimal solution with $u_k^{\omega_1} = u_k^{\omega_2}$.

Proof. Constraints (2.29c) and (2.29d) imply that $u_k^\omega \leq \min_{k' \in \bar{K}_k^\omega} x_{k'}$. This inequality is tight for at least one optimal solution to (2.29) since the objective function coefficients of u_k^ω are all non-negative and there are no other constraints on u , aside from the simple bound constraints which are automatically satisfied. So if $\bar{K}_k^{\omega_1} = \bar{K}_k^{\omega_2}$ then $u_k^{\omega_1} = u_k^{\omega_2}$ in at least one optimal solution. \square

By Proposition 22, for any $\omega_1, \omega_2 \in \Omega$ and $k \in K$ with $\bar{K}_k^{\omega_1} = \bar{K}_k^{\omega_2}$, we can eliminate the variable $u_k^{\omega_2}$ by replacing $s_k^{\omega_1}$ with $s_k^{\omega_1} + s_k^{\omega_2}$. In addition, we can eliminate a constraint from both (2.29c) and (2.29d).

The hypothesis of Proposition 22 is satisfied if both smugglers ω_1 and ω_2 prefer the same set of checkpoints to checkpoint k . This may occur, for example, if a pair of smugglers shares an origin and have destinations in close proximity to each other, or vice versa. The pair of smugglers may, therefore, rank their top checkpoints differently but rank the rest of their checkpoints identically. This pair of smugglers does not satisfy the strict ordering condition of Proposition 6 but does, at least for those checkpoints that are identically ranked, satisfy the ordering condition of Proposition 22 and so at least some reduction in problem size is possible.

2.8 Branching

We now describe a branching scheme which makes use of the idea, presented in Proposition 21, that we can eliminate auxiliary variables that, if

positive, force the budget constraint to be violated. For each smuggler, we can identify and fix such variables by sorting the checkpoints in order of decreasing conditional evasion probability, then scanning the sorted list until we find a checkpoint k with $\sum_{k' \in \bar{K}_k^\omega} c_{k'} > b$. We can view this as applying a greedy algorithm to each smuggler in which we interdict the chosen smuggler's best uninterdicted checkpoint until the budget is depleted. This greedy algorithm gives an optimal policy if the interdictor could wait until the smuggler scenario was revealed and then deploy the detectors, that is, the greedy algorithm solves the wait-and-see problem. The solution to the wait-and-see problem provides a lower bound on the conditional evasion probability for each smuggler scenario, and the auxiliary variables that we fix to zero correspond to scenario-checkpoint pairs with evasion probabilities smaller than this lower bound.

The potential for the reformulated model to still have weak LP relaxation bounds lies with the fact that, if the smugglers do not rank the checkpoints similarly, it is impossible to achieve all of these lower bounds simultaneously if we must deploy detectors before the smuggler scenario is revealed. However, as we allocate detectors within a branch-and-bound tree, our remaining budget, after having fixed $x_k = 1$ for some subset of checkpoints, will decrease and likely tighten the lower bounds on some of the conditional evasion probabilities. We could resolve the wait-and-see problem and fix additional u_k^ω variables at each node in the tree, but this can be time consuming. This motivates a branching scheme that allocates detectors as quickly as possible. So while a standard branching scheme would be to pick a checkpoint k' with

$x_{k'}$ fractional in the LP relaxation and create a subproblem in which $x_{k'} = 0$ and another in which $x_{k'} = 1$, the scheme we describe here instead branches on whether an entire subset of checkpoints receives detectors. Then, for each subproblem generated, we recompute the wait-and-see bounds, fix additional u_k^ω variables, and finally decide whether to branch on another subset of checkpoints or to hand the subproblem off to a general purpose branch-and-bound solver.

To further motivate this scheme, consider the following. Suppose that, for a particular smuggler, there exists a subset of checkpoints such that each checkpoint, if not interdicted, gives the smuggler a path with high evasion probability. Also suppose that if the entire subset is interdicted that the smuggler is forced to traverse a path with a much lower evasion probability. This phenomenon may occur in practical problems if there exists a cluster of checkpoints in close proximity to each other and, say, the origin for a particular smuggler. If this subset of checkpoints also provides paths with high evasion probabilities to some other smugglers, we have strong reason to believe that the entire subset should be interdicted in an optimal solution. To check this hypothesis, we can generate two subproblems: one in which every checkpoint in the subset is interdicted and one in which at least one checkpoint in the subset is not interdicted. So if S is the subset of checkpoints we think should be interdicted, we create one subproblem in which $\sum_{k \in S} x_k = |S|$ and another in which $\sum_{k \in S} x_k \leq |S| - 1$. We refer to the creation of the former subproblem as “branching up” and the latter as “branching down.” The benefit of this

scheme is that it allows us, if the subset S is intelligently chosen, to fix several extra variables in the subproblems that could not be fixed at the root node. The following result describes how extra variables may be fixed.

Proposition 23. *Consider model (2.29) and let $S \subset K$. Then*

- (i) $\sum_{k \in S} x_k = |S|$ implies that $u_k^\omega = 0$ for all $k \in K$ and $\omega \in \Omega$ such that $\sum_{k' \in \bar{K}_k^\omega \cup S} c_{k'} > b$
- (ii) $\sum_{k \in S} x_k \leq |S| - 1$ implies that $u_k^\omega = 0$ for all $k \in K$ and $\omega \in \Omega$ such that $S \subseteq \bar{K}_k^\omega$.

Proof. Part (i) follows from the fact that, if $u_k^\omega > 0$, then $x_{k'} = 1$ for all $k' \in \bar{K}_k^\omega$. But $\sum_{k \in S} x_k = |S|$ implies that $x_k = 1$ for all $k \in S$. If interdicting all checkpoints that are either in \bar{K}_k^ω or S exceeds the budget, that is if $\sum_{k' \in \bar{K}_k^\omega \cup S} c_{k'} > b$, then (2.29) is infeasible. Part (ii) follows immediately by noting that $u_k^\omega > 0$ implies that $\sum_{k' \in \bar{K}_k^\omega} x_{k'} = |\bar{K}_k^\omega|$, which in turn implies that $\sum_{k \in S} x_k = |S|$ if $S \subseteq \bar{K}_k^\omega$. \square

We choose a subset S to branch on so that each of the subproblems generated makes significant progress towards feasibility. The size of set S provides a measure of how much progress is made branching up since a larger S leads to more variables being fixed in the resulting subproblem. To ensure significant progress is made branching down, S should be chosen such that some of the components of u which were positive in the solution to the LP relaxation of the parent problem are fixed to 0 in the subproblem. Since increasing the size of S typically leads to fewer components of u being fixed

in the branch down, there is an inherent tradeoff between the progress made in one branch and the progress made in the other.

We quantify the progress made in the branch down as follows. If we view (2.29) as a reward collecting problem in which a reward s_k^ω is earned when the variable u_k^ω is set to 1, then given a subset S and a feasible LP relaxation solution (\hat{x}, \hat{u}) , we can compute the total reward earned by the LP relaxation solution but that cannot be earned if S is not fully interdicted via the following function:

$$Loss(S) = \sum_{\omega \in \Omega} \sum_{k: S \subseteq \bar{K}_k^\omega} p^\omega s_k^\omega \hat{u}_k^\omega.$$

Ideally, we seek a subset that maximizes this loss function for each potential subset size t , then pick some subset along this efficient frontier of solutions. Since maximizing the loss function for a fixed subset size t is as hard as solving an instance of BiSNIP, we instead greedily approximate this efficient frontier as follows. Starting with $S_0 = \emptyset$ and for every possible subset size $t = 1, \dots, b$, we compute

$$Loss_t = \max_{k \in K \setminus S_{t-1}} Loss(S_{t-1} \cup \{k\}),$$

and let $S_t = S_{t-1} \cup \{k^*\}$, where k^* is the maximizer.

Once we have an approximation for the efficient frontier, we must choose amongst its members a subset on which to branch. In our computational experiments, we choose to branch on the subset S_{t^*} , where $t^* \in \operatorname{argmax}_{1 \leq t \leq b} t \cdot Loss_t$. The idea behind this choice is that it should promote solutions near the center of our approximated efficient frontier and guarantee

that each branch makes positive progress. In order to obtain an integer feasible solution as quickly as possible, we use a depth-first node selection strategy in which the branch up is evaluated before the branch down. A pseudo-code representation of the customized branch-and-bound algorithm can be found in Appendix A.

2.9 Computational Results

In this section, we discuss results from US model instances, restricting attention to land border crossings entering the continental US from Mexico and Canada. Using a North American road network, we model 7 origins in Mexico, 7 origins in Canada and 10 destinations in the US, giving a total of $|\Omega| = 140$ threat scenarios. Since all detectors are identical, we use a cardinality-constrained special case of the BiSNIP model, i.e., $c_k = 1$, for all $k \in K$, in constraint set X , and we solve the model for various budget values, b , representing the number of border crossings equipped with detectors. Each checkpoint has an indigenous evasion probability based on its perceived vulnerability, p_k , and this varies by checkpoint, k . However, facing the same threat specified by ω , we assume detectors in distinct locations behave identically and the probability a smuggler evades detection, by the detector, is q^ω , which does not depend on k . If a detector is installed at k we assume both the indigenous detection capability and the detector technology are independently employed so that q_k^ω used in (2.2d) is given by $q_k^\omega = q^\omega p_k$. If ω only specifies the origin-destination pair (e.g., because distinct shielding scenarios have already

been aggregated) then we can drop q 's dependence on ω so that $q_k^\omega = qp_k$ for some constant q . The indigenous evasion probabilities are based on a multi-attribute factor model described in detail in [40]. All of the MIPs associated with both sets of BiSNIP instances were solved via the commercially-available CPLEX software [10].

Figure 2.3 shows the 136 motor-crossing checkpoints we consider. The figure also indicates four clusters of checkpoints important in results we describe below. We again solve the associated BiSNIP instances for a range of values of b , the number of detectors we can install. These hedge against 140 origin-destination threat scenarios, with half originating in Canada and the other half in Mexico. In addition to ranging b we assume the effectiveness of the detection equipment is independent of the scenario and checkpoint, that is $q_k^\omega = qp_k$ for some constant q . We create multiple model instances by ranging the value of q .

Figure 2.4 shows the optimal evasion probability over all threat scenarios versus the budget for four values of the detector effectiveness, q . The evasion probability is reported as a fraction of that when no detectors are installed. Significant jumps in the graph occur when we are given just enough detectors to interdict an entire cluster of checkpoints. For example, we notice a large decrease in the evasion probability as the budget increases to $b = 34$ as such a budget allows us to interdict every checkpoint along the Mexican border. Smaller but still significant jumps occur when the budget increases

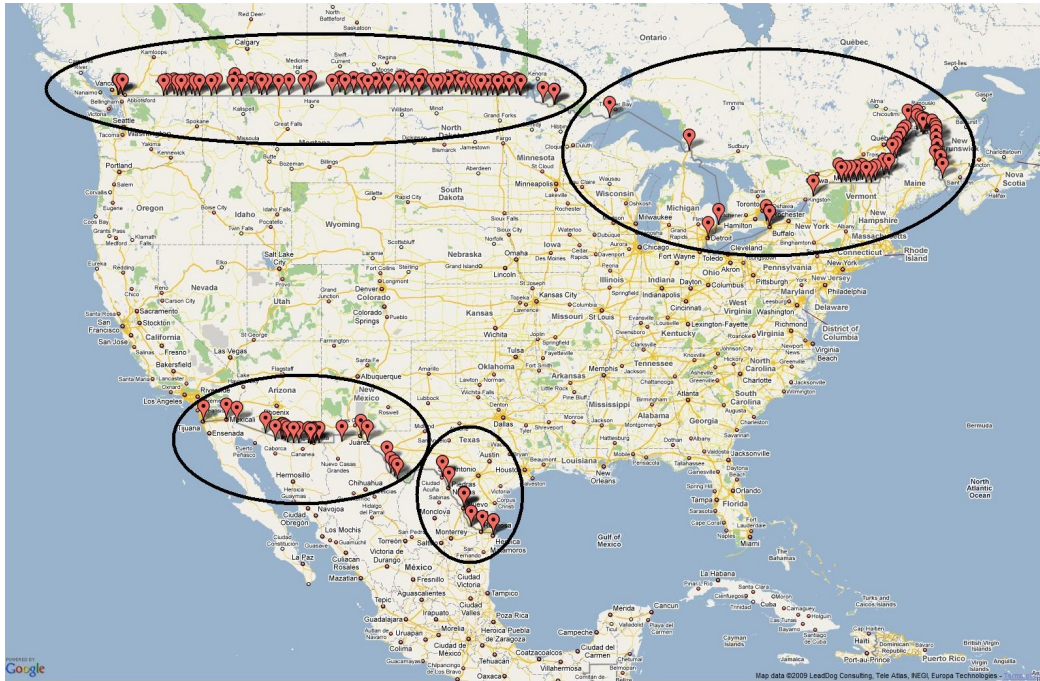


Figure 2.3: The figure shows 136 motor-crossing checkpoints from Canada and Mexico into the continental United States, and groups the checkpoints into four clusters.

to 11, allowing us to interdict all checkpoints in Mexico east of Big Bend (see Figure 2.5), and when the budget increases to $b = 97$, allowing us to interdict all checkpoints in Mexico and all checkpoints in Canada west of Lake Huron (see Figure 2.6a and 2.6b).

Also noteworthy is the fact that for small values of the budget ($b < 11$), the optimal solution interdicts checkpoints along the Great Lakes (see Figure 2.5a). Intuitively this is because there are more gaps between those checkpoints than there are anywhere else. Finally, we note that the solutions did vary as detectors become less effective. A notable example of this is that

with more effective detectors ($q = 0, 0.25, 0.5$), there is an incentive to shift all detectors from eastern Canada to western Canada when the budget increases from 96 to 97. This was not the case with the most ineffective detectors ($q = 0.75$) as such detectors could not convince smugglers with origins in western Canada to travel around the Great Lakes to traverse a detector-free checkpoint (see Figure 2.6c).




Figure 2.4: The figure shows the improvement factor as a function of the number of detectors installed for the US model. The four plots correspond to different levels of effectiveness of the detectors, specifically, with $q = 0.75, 0.50, 0.25$ and 0 in $q_k^\omega = qp_k$.

To show the value of using wait-and-see bounds and reformulation (2.29), we report optimality gaps for perfectly reliable detectors, $q = 0$, and various budget levels b under the following configurations: (1) Model (2.4); (2) Model (2.4) with (2.5); (3) Model (2.29); (4) Model (2.29) with the pre-processing suggested by Proposition 21. Table 2.1 reports optimality gaps as percentages of the optimal unconditional evasion probability given by the optimal value to (2.4).

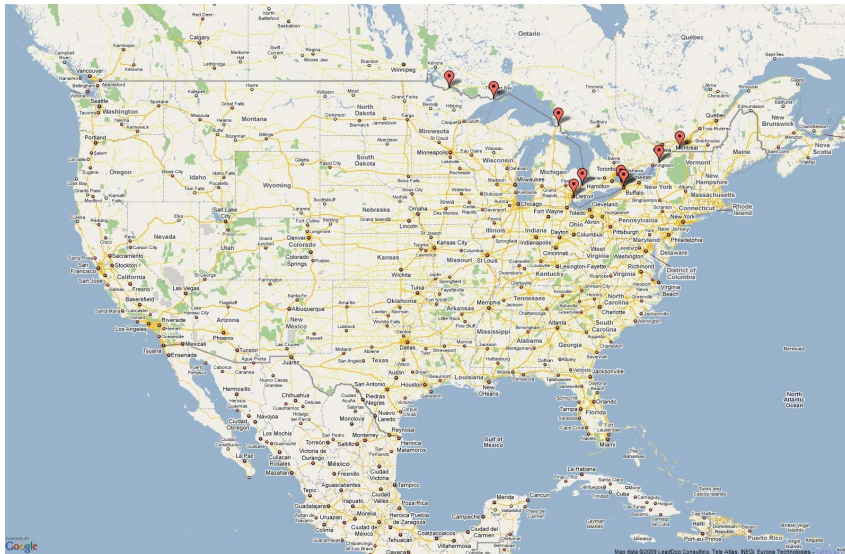
To show the value of the various computational enhancements proposed in this chapter, we report solution times for perfectly reliable detectors, $q = 0$, and various budget levels b under the following configurations: (1) Model (2.4) with (2.5) (BASE); (2) Model (2.29) using Proposition 21 (REF); (3) Model (2.29) using Propositions 21 and 22 (REF-AGG); (4) Model (2.29) using Propositions 21 and 22 solved by the customized branch-and-bound algorithm of Section 2.8 (REF-AGG-C). Table 2.2 reports solution times in seconds. The computation times reported were on a 3.73 GHz Dell Xeon dual-processor machine with 8 GB of memory, running CPLEX version 10.1 with an absolute tolerance of 10^{-4} .

b	(2.4)	(2.4) + (2.5)	(2.29)	Preprocessed (2.29)
10	17.5	0.53	14.3	0.116
20	31.2	2.01	28.1	0.327
30	53.7	2.89	50.9	1.46
40	11.3	3.62	3.49	0
50	22.0	6.14	11.8	0.964
60	32.4	6.08	20.8	0.456
70	49.4	14.2	36.3	6.67
80	69.0	21.0	54.2	13.1
90	101	31.3	83.4	23.8
100	139	40.9	117	33.8

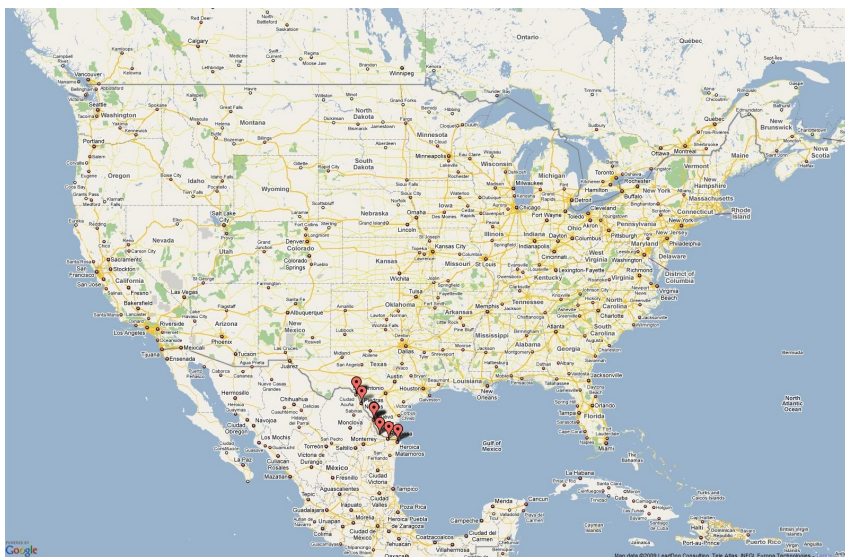
Table 2.1: Optimality gaps as percentages of optimal unconditional evasion probabilities for US model instances with perfectly reliable detectors, $q = 0$.

b	BASE	REF	REF-AGG	REF-AGG-C
10	0.4	0.7	0.3	0.06
20	1.9	2.5	1.3	0.15
30	3.4	7.3	1.2	0.17
40	2.3	1.56	0.4	0.16
50	4.9	12	2.6	0.32
60	14	29	6.5	0.39
70	194	162	53	1.3
80	1332	311	70	1.6
90	×	1235	113	2.6
100	1133	1904	80	2.9

Table 2.2: Solution times in seconds for US model instances with perfectly reliable detectors, $q = 0$. × indicates that the solution time exceeded 2 hours.

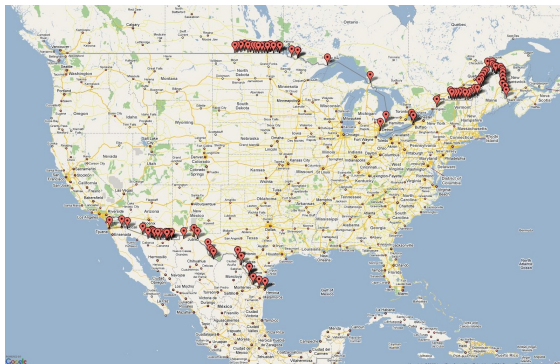


(a) $b = 10$

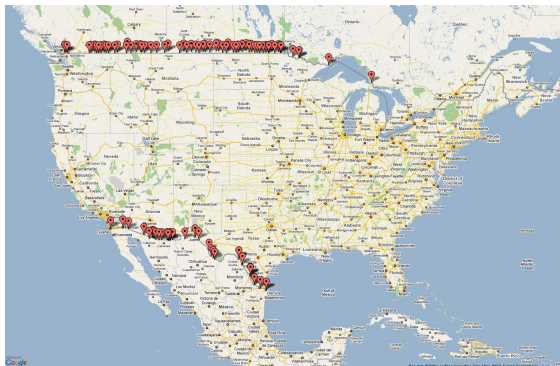


(b) $b = 11$

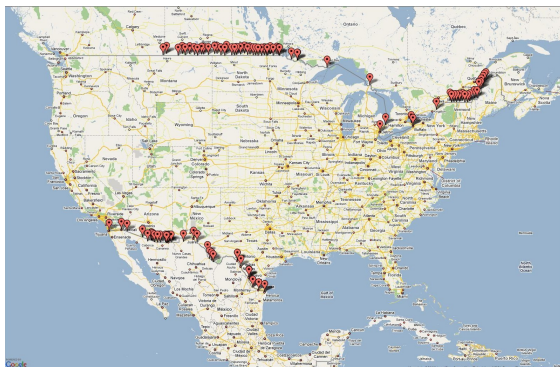
Figure 2.5: Part (a) of the figure shows the optimal solution to the US model instance with perfectly reliable detectors, $q = 0$, and with a budget to install detectors at $b = 10$ locations. Part (b) of the figure is identical but for $b = 11$. Note that the full number of checkpoints are not visible in the map due to their close proximity.



(a) $b = 96$ and $q = 0$



(b) $b = 97$ and $q = 0$



(c) $b = 97$ and $q = 0.75$

Figure 2.6: Part (a) of the figure shows the optimal solution to the US model instance with perfectly reliable detectors, $q = 0$, and with a budget to install detectors at $b = 96$ locations. Part (b) of the figure is identical but for $f = 97$. Part (c) of the figure is for $b = 97$ and $q = 0.75$.

Chapter 3

Two-Person Zero-Sum Games for Network Interdiction

3.1 Introduction

This chapter considers game-theoretic models for nuclear smuggling detection. In the previous chapter, we consider a stochastic network interdiction problem in which a smuggler of nuclear material chooses a maximum-reliability path through a transportation network with full knowledge of the locations of the locations of radiation detectors installed by an interdictor. In Section 3.2 we consider a variant of this problem in which we assume that the smuggler knows the locations of the detectors only via a probability distribution. The resulting model is a two-person zero-sum Cournot game. We give some complexity results and discuss solution techniques. Section 3.3 describes two-person zero-sum game models in which one player, subject to a budget constraint, may add strategies to his own strategy set or remove strategies from his opponent's strategy set. We conclude in Section 3.4 by describing a model in which only a subset of detector locations are revealed to the smuggler.

3.2 TPZSCGs with Exponentially Many Strategies

3.2.1 Motivation

We consider a zero-sum Cournot-game with two opposing parties, an interdictor and an evader. Suppose that the interdictor, in an attempt to thwart smuggling attempts, has a fixed budget with which he can deploy detectors to a set of border checkpoints indexed by $j \in J$. Then suppose that an evader wishing to cross the border must choose a checkpoint from J . If the evader attempts to cross checkpoint j and no detector is installed then he evades detection with probability p_j . If checkpoint j has a detector installed, the evasion probability is instead $q_j < p_j$. Since the interdictor seeks to minimize the evasion probability and the evader seeks to maximize it, this problem can be modeled as a two-person zero-sum game.

If the deployment of detectors is transparent, that is, if the evader knows which checkpoints received detectors, then this problem can be modeled as a Stackelberg game. This case has been discussed extensively in [25, 26]. We examine the case where the evader knows the problem parameters (p_j , q_j , and the interdiction budget and costs) but cannot see where detectors have been deployed. In this case since the evader is unaware of the actions taken by the interdictor, the two parties effectively make their respective decisions simultaneously and so a Cournot model is appropriate. We may prefer a Cournot model if, for example, all checkpoints are equipped with a “black box” that is indistinguishable from a real detector and the interdictor then places a real detector in some subset of the black boxes. Alternatively, suppose

that all checkpoints already have detectors installed and that the interdicator can “upgrade” some subset of these detectors but the upgraded detectors are indistinguishable from the original detectors.

3.2.2 Single-Evader Model

In the following TPZSCG, the interdicator is the row player and the evader the column player. The model concerns a geographic region in which the evader begins. The interdicator’s goal is to contain the evader in this region by detecting any attempt to cross the region’s border. Each of the interdicator’s pure strategies represents a feasible deployment of detectors across the border checkpoints, and each of the evader’s pure strategies represents a checkpoint to cross. So our payoff matrix has a row for each subset of checkpoints on which the interdicator can install detectors, without exceeding the budget, and a column for each checkpoint. The evader’s goal is to find a mixed strategy, here a probability distribution over the checkpoints, which maximizes the probability that he crosses the border undetected, while the interdicator’s goal is to find a probability distribution over all feasible detector deployments which minimizes this probability. We can formulate this problem

as follows:

Indices and Sets

$i \in I$ feasible detector deployments the interdicator can choose

$j \in J$ checkpoints the evader can choose

Data

c_j cost of installing a detector at checkpoint j

b total budget for installing detectors

p_j probability evader can traverse j undetected when no detector is installed

q_j probability evader can traverse j undetected when a detector is installed

A_{ij} game's payoff if evader selects checkpoint j to cross and interdicator selects detector deployment i , i.e., $A_{ij} = q_j$ if deployment i places a detector at checkpoint j and otherwise $A_{ij} = p_j$. Here, for each row $i \in I$ we have $\sum_{j \in J} c_j I(A_{ij} = q_j) \leq b$, where $I(\cdot)$ is the indicator function

Decision Variables

x_i probability that the interdicator chooses detector deployment i

y_j probability that the evader chooses checkpoint j to cross the border

Formulation

$$v^* = \max_{y, \lambda} \lambda \tag{3.1a}$$

$$\text{s.t.} \quad \lambda \leq \sum_{j \in J} A_{ij} y_j \quad : x_i, \quad i \in I \tag{3.1b}$$

$$\sum_{j \in J} y_j = 1 \tag{3.1c}$$

$$y_j \geq 0, \quad j \in J. \tag{3.1d}$$

The right-hand side of constraint (3.1b) is the evasion probability associated with detector deployment $i \in I$, and the evader's checkpoint-selection strategy

y . The constraint has the effect of selecting the minimum of these evasion probabilities over all detector deployments $i \in I$, and the evader seeks to maximize that value. The optimal dual variables $x_i, i \in I$, on constraints (3.1b) define an optimal mixed strategy over all feasible detector deployments.

It may seem that in this model the evader must place a probability distribution across the checkpoints first and then the interdicator picks a detector deployment that hedges optimally against the evader's distribution. That is, it seems that we are forcing the evader to act first. However, if we were to reverse the order of the decisions and have the interdicator pick a distribution first, we would arrive at a formulation which is the dual of (3.1). So although one might view the decisions in this model as taking place sequentially, the allowance of mixed strategies implies simultaneous decisions by the players.

Model (3.1) is an LP with $|I| + 1$ structural constraints, and since $|I|$ is the number of feasible detector deployments, the number of constraints is of exponential size. For example, if $c_j = 1, j \in J, b = 20$, and $|J| = 100$, then the set I will have cardinality $|I| = \binom{|J|}{b} = \binom{100}{20}$. However, the bulk of these constraints are irrelevant, that is, they are slack at an optimal solution. More specifically, in our setting, J is of modest size, and we know that at an optimal extreme point, at most $|J| - 1$ of the constraints (3.1b) have positive dual variables $x_i, i \in I$. In later sections we describe how the weighted majority algorithm can be used to find approximate solutions to (3.1). Here, we describe a row-generation scheme which can solve model (3.1) exactly and gives insights into its complexity. We first solve a relaxation of model (3.1) with constraints

(3.1b) defined over $I' \subset I$. Solving the associated LP, we obtain \hat{y} and $\hat{\lambda}$ as well as dual multipliers \hat{x} . Then, we determine whether \hat{y} and $\hat{\lambda}$ are feasible to the original problem, i.e., with (3.1b) defined on set I . If not, we find the most violated constraint (3.1b). This assessment and identification is carried out via solving a so-called separation problem.

This separation problem is equivalent to finding a detector deployment that minimizes the evasion probability given that the evader's mixed strategy is \hat{y} . Since λ appears on the left-hand side of all the constraints in (3.1b), we need to find the smallest right-hand side over all deployments $i \in I$, that is, we solve:

$$\min_{i \in I} \sum_{j \in J} A_{ij} \hat{y}_j. \quad (3.2)$$

We can express $A_{ij} = p_j - (p_j - q_j)z_j$, where z_j equals 1 one if a detector is installed on checkpoint j and z_j equals 0 otherwise, and the set $Z = \{z : \sum_{j \in J} c_j z_j \leq b, z_j \in \{0, 1\}, j \in J\}$ enumerates all feasible detector deployments $i \in I$. As a result, the separation problem for model (3.2) can be rewritten as:

$$\max_{z \in Z} \sum_{j \in J} (p_j - q_j) \hat{y}_j z_j. \quad (3.3)$$

Model (3.3) is a knapsack problem and is NP-hard. It can be solved, however, in pseudo-polynomial time using dynamic programming. Given an optimal solution z^* to (3.3) we check whether

$$\hat{\lambda} \leq \sum_{j \in J} [p_j - (p_j - q_j)z_j^*] \hat{y}_j. \quad (3.4)$$

If so, then the current solution (\hat{x}, \hat{y}) solves (3.1). If not then z^* yields a new row i^* with $A_{i^*} = (p_j - (p_j - q_j)z_j^*)_{j \in J}$. We replace I' with $I' \cup \{i^*\}$, re-solve the associated relaxation of (3.1), and repeat until (3.4) is satisfied. Alternatively, we know the optimal solution to (3.1) defined over a subset $I' \subset I$ yields an optimal value $\bar{v} \geq v^*$. And, the value $\sum_{j \in J} [p_j - (p_j - q_j)z_j^*] \hat{y}_j \leq v^*$ because it corresponds to a feasible strategy of the evader coupled with an optimal response of the interdicator. If these upper and lower bounds on v^* are within ϵ we may terminate with an ϵ -optimal solution. In this case the interdicator cannot decrease the evasion probability by more than ϵ by deviating from the mixed strategy suggested by this near-optimal solution to the model.

3.2.3 Multiple-Evader Model

Now say that instead of a single evader, a random evader $\omega \in \Omega$ is chosen according to a probability distribution p^ω known to the interdicator. Each evader may have different evasion probabilities p_j^ω and q_j^ω . Let y_j^ω be evader ω 's mixed strategy and let $A_{i_j}^\omega$ be defined for each evader in a fashion analogous to that for the single-evader model. Then the multiple-evader model

may be formulated as follows:

$$v^* = \max_y \min_x \sum_{i \in I} \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega A_{ij}^\omega x_i y_j^\omega \quad (3.5a)$$

$$\text{s.t.} \quad \sum_{i \in I} x_i = 1 \quad : \lambda \quad (3.5b)$$

$$\sum_{j \in J} y_j^\omega = 1, \quad \omega \in \Omega \quad (3.5c)$$

$$x_i \geq 0, \quad i \in I \quad (3.5d)$$

$$y_j^\omega \geq 0, \quad j \in J, \omega \in \Omega. \quad (3.5e)$$

Here we take the view that the interdicator is playing a game against multiple evaders, one of whom is selected according to p^ω . Each evader has a different payoff matrix A^ω as evasion probabilities may vary across evaders. Again, since the objective (3.5a) is concave in $y = [y^\omega]_{\omega \in \Omega}$ for fixed x and convex in x for fixed y and the feasible regions for x and y are both convex, we may interchange the “max” and the “min” and obtain the same optimal value, v^* . So, we can view the formulation as having an interdicator and multiple evaders decide on strategies simultaneously, but then only one evader (selected by p^ω) is realized.

If we (a) fix y to create an LP with variables x , (b) define λ as the dual variable for the single structural constraint in that LP, (c) take the dual of the

LP, and (d) then free y , we obtain the following equivalent problem:

$$v^* = \max_{y, \lambda} \quad \lambda \quad (3.6a)$$

$$\text{s.t.} \quad \lambda \leq \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega A_{ij}^\omega y_j^\omega \quad : x_i, \quad i \in I \quad (3.6b)$$

$$\sum_{j \in J} y_j^\omega = 1, \quad \omega \in \Omega \quad (3.6c)$$

$$y_j^\omega \geq 0, \quad j \in J, \quad \omega \in \Omega. \quad (3.6d)$$

Constraint set (3.6b) contains an exponential number of constraints, most of which are slack at a basic feasible solution. So we can solve relaxations of (3.6) with constraints (3.6b) only defined over a subset, $I' \subset I$, of feasible detector deployments to obtain \hat{y} and $\hat{\lambda}$ as well as dual variables \hat{x} , then identify the most violated constraint in (3.6b) and add the associated detector deployment to I' . To identify the most violated constraint in (3.6b) for some \hat{y} and $\hat{\lambda}$ we find the constraint with the smallest right-hand side. Mimicking the development from the single-evader case, this problem can be formulated as follows:

$$\max_{z \in Z} \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega (p_j^\omega - q_j^\omega) \hat{y}_j^\omega z_j. \quad (3.7)$$

As before, the optimal z^* indicates that (\hat{x}, \hat{y}) is optimal to (3.5) if

$$\hat{\lambda} \leq \sum_{\omega \in \Omega} \sum_{j \in J} p^\omega [p_j^\omega - (p_j^\omega - q_j^\omega) z_j^*] \hat{y}_j^\omega, \quad (3.8)$$

and identifies a violated row i^* otherwise. We can view model (3.7) as finding the best response to the mixed strategies of all the evaders in Ω . If the row i^* corresponding to that best response is not already in I' , then it is added to I' and the relaxation is resolved.

3.2.4 Cardinality-Constrained Case

Here we adapt an idea from [39] to simplify model (3.5) in the case where the detector installation cost is constant over all border checkpoints. Letting θ^ω be the dual variables of constraints (3.6c) and taking the dual of (3.6) we arrive at the following formulation:

$$v^* = \min_{x, \theta} \sum_{\omega \in \Omega} \theta^\omega \quad (3.9a)$$

$$\text{s.t.} \quad \theta^\omega \geq p^\omega \sum_{i \in I} A_{ij}^\omega x_i \quad : y_j^\omega, \quad j \in J, \quad \omega \in \Omega \quad (3.9b)$$

$$\sum_{i \in I} x_i = 1 \quad : \lambda \quad (3.9c)$$

$$0 \leq x_i, \quad i \in I. \quad (3.9d)$$

Model (3.9) has exponentially many variables but can be reduced to a model with $|\Omega| + |J|$ variables if we assume that $c_j = 1$ for all $j \in J$. We define the strategy-checkpoint incidence matrix D by

$$D_{ij} = \begin{cases} 1 & \text{if strategy } i \text{ places a detector on checkpoint } j \\ 0 & \text{otherwise.} \end{cases}$$

If we define $\hat{x}_j = \sum_{i \in I} D_{ij} x_i, j \in J$, and express $A_{ij}^\omega = p_j^\omega + (q_j^\omega - p_j^\omega) D_{ij}$, then constraints (3.9b) can be written as:

$$\begin{aligned} \theta^\omega &\geq p^\omega \sum_{i \in I} (p_j^\omega + (q_j^\omega - p_j^\omega) D_{ij}) x_i \\ &= p^\omega (p_j^\omega \sum_{i \in I} x_i + (q_j^\omega - p_j^\omega) \sum_{i \in I} D_{ij} x_i) \\ &= p^\omega (p_j^\omega + (q_j^\omega - p_j^\omega) \hat{x}_j). \end{aligned}$$

Since $\sum_{i \in I} x_i = 1$, we know that $\hat{x}_j = \sum_{i \in I} D_{ij} x_i \leq \max_i D_{ij} = 1$ and $\sum_{j \in J} \hat{x}_j = \sum_{i \in I} x_i \sum_{j \in J} D_{ij} = b$. This suggests the following relaxation to

model (3.9):

$$\min_{\hat{x}, \theta} \sum_{\omega \in \Omega} \theta^\omega \quad (3.10a)$$

$$\text{s.t.} \quad \theta^\omega \geq p^\omega (p_j^\omega + (q_j^\omega - p_j^\omega) \hat{x}_j), \quad j \in J, \omega \in \Omega \quad (3.10b)$$

$$\sum_{j \in J} \hat{x}_j = b \quad (3.10c)$$

$$0 \leq \hat{x}_j \leq 1, \quad j \in J. \quad (3.10d)$$

This is a linear program with $|J| + |\Omega|$ variables and $|J||\Omega| + 1$ structural constraints. For any (x, θ) pair that is feasible to (3.9), the pair $(D^T x, \theta)$ is feasible to (3.10), and so (3.10) is clearly a relaxation to (3.9). The following result shows that we can also map \hat{x} back to x and so the optimal values of (3.9) and (3.10) are in fact equal.

Proposition 24. *For every (\hat{x}, θ) which is feasible to (3.10), there exists an (x, θ) with $D^T x = \hat{x}$ which is feasible to (3.9).*

Proof. We must show that the system $D^T x = \hat{x}$, $e^T x = 1$, $x \geq 0$ always has a solution, where we again use e to denote the vector of all 1s. Summing the constraints of $D^T x = \hat{x}$ and dividing by b yields $e^T x = 1$, so that constraint is redundant. It suffices to show, therefore, that there exists a solution to $D^T x = \hat{x}$, $x \geq 0$. By Farkas' lemma that system has a solution if and only if the system $\pi D^T \leq 0$, $\pi \hat{x} > 0$ does not have a solution. $\pi D^T \leq 0$ implies that $\sum_{j \in J} D_{ij} \pi_j \leq 0, i \in I$. Since vector $(D_{ij})_{j \in J}$ has exactly b components equal to 1 and the rest are zero, these inequalities can be written as $\sum_{j \in J} \pi_j \leq 0$

for all $J' \subset J, |J'| = b$. But we know that

$$\pi \hat{x} = \sum_{j \in J} \pi_j \hat{x}_j \leq \max_{J' \subset J, |J'|=b} \sum_{j \in J'} \pi_j \leq 0$$

since $\sum_{j \in J} \hat{x}_j = b$ and \hat{x} is bounded above by 1 componentwise. Therefore $\pi D^T \leq 0, \pi \hat{x} > 0$ has no solution and $D^T x = \hat{x}, x \geq 0$ must have a solution. \square

Given \hat{x} which solves (3.10) we can find a corresponding x which solves (3.9) by solving the phase 1 linear program:

$$\begin{aligned} \min_{x,s} \quad & \sum_{j \in J} s_j \\ \text{s.t.} \quad & \sum_{i \in I} D_{ij} x_i + s_j = \hat{x}_j \quad : \pi_j, \quad j \in J \\ & 0 \leq x_i, \quad i \in I \\ & 0 \leq s_j, \quad j \in J. \end{aligned}$$

The dual of this LP is:

$$\max_{\pi} \quad \sum_{j \in J} \hat{x}_j \pi_j \tag{3.12a}$$

$$\text{s.t.} \quad \sum_{j \in J} D_{ij} \pi_j \leq 0 \quad : x_i, \quad i \in I \tag{3.12b}$$

$$\pi_j \leq 1 \quad : s_j, j \in J. \tag{3.12c}$$

Since constraint set (3.12b) is exponentially sized, we use row generation to solve (3.12). Given an optimal solution $\hat{\pi}$ to a relaxation of (3.12) with constraint set (3.12b) only defined over $I' \subset I$, we can identify the most violated of the relaxed constraints by maximizing $\sum_{j \in J} D_{ij} \hat{\pi}_j$ over all $i \in I$.

For any given $i \in I$, the vector $(D_{ij})_{j \in J}$ has exactly b components equal to 1, so this maximization can be done by finding the b largest values of $\hat{\pi}_j$ over $j \in J$. If the sum of these b largest values is non-positive then the current solution $\hat{\pi}$ solves (3.12) and the dual variables of (3.12b) are a solution to the system $Dx = \hat{x}$, $x \geq 0$. Otherwise we add the corresponding row to I' and repeat. Alternatively, as we describe in the following section, we may use the weighted majority algorithm to find an approximate solution to the game. We conclude by noting that the ability to compute the value of the game v^* for this special case by solving a polynomially-sized linear program facilitates the development of an extension of model (3.5) in which some subset of detectors are visible to the smuggler. We return to this idea in Section 3.4.

3.2.5 Weighted Majority

We now describe how the weighted majority algorithm can be used to generate a near-optimal solution to (3.5). Every iteration, $t = 1, \dots, T$, of the weighted majority algorithm can be viewed as a fictitious play of the game in the following sense. We take the perspective of the smugglers $\omega \in \Omega$, for whom we maintain weights $w_j^{t,\omega}$ for every checkpoint $j \in J$ and iteration $t = 1, \dots, T$ which are updated via:

$$w_j^{t+1,\omega} = w_j^{t,\omega} \beta(\omega)^{l_j^{t,\omega}},$$

where $\beta(\omega) \in (0, 1)$ is an appropriately chosen constant and $l_j^{t,\omega}$ is the loss if smuggler ω traverses checkpoint j in iteration t . We describe how both are chosen shortly. Every iteration t , the smugglers choose mixed strategies $y^{t,\omega}$

according to:

$$y_j^{t,\omega} = \frac{w_j^{t,\omega}}{\sum_{j' \in J} w_{j'}^{t,\omega}}, \quad j \in J, \quad t = 1, \dots, T,$$

and the interdicator responds with the best response x^t given by:

$$x^t \in \operatorname{argmin}_{x \in X} \sum_{i \in I} \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega A_{ij}^\omega x_i y_j^{t,\omega},$$

where $X = \{x \in \mathbb{R}_+^{|I|} : \sum_{i \in I} x_i = 1\}$. There exists at least one pure-strategy best response, which can be computed by solving a problem of the form (3.7). Each loss $l_j^{t,\omega}$ is chosen to be the detection probability for checkpoint j and smuggler ω given the interdicator's best response, that is, $l_j^{t,\omega} = \sum_{i \in I} (1 - A_{ij}^\omega) x_i^t$. This gives $l_j^{t,\omega} = 1 - p_j^\omega$ if the interdicator's best response is to allocate a detector to checkpoint j , and $l_j^{t,\omega} = 1 - q_j^\omega$ otherwise. We choose $\beta(\omega)$ according to:

$$\beta(\omega) = \frac{1}{1 + \sqrt{\frac{2 \ln |J|}{\tilde{L}(\omega)}}}, \quad \omega \in \Omega.$$

where $\tilde{L}(\omega)$ is an upper bound on the detection probability for smuggler ω . Nominally we may choose $\tilde{L}(\omega) = T \min_{j \in J} (1 - q_j^\omega)$. A possibly tighter choice is $T(1 - v^*(\omega))$ where $v^*(\omega)$ is the optimal value of

$$\min_{\hat{x}, \theta^\omega} \theta^\omega \tag{3.13}$$

$$\text{s.t.} \quad \theta^\omega \geq p_j^\omega + (q_j^\omega - p_j^\omega) \hat{x}_j, \quad j \in J, \tag{3.14}$$

$$\sum_{j \in J} c_j \hat{x}_j = b \tag{3.15}$$

$$0 \leq \hat{x}_j \leq 1, \quad j \in J. \tag{3.16}$$

We note that (3.13) can be solved efficiently via a greedy algorithm.

By Theorem 3, after T iterations, we have for each $\omega \in \Omega$

$$\frac{L_{\beta(\omega)}^\omega}{T} \leq \min_{j \in J} \frac{L_j^\omega}{T} + \Delta_T^\omega, \quad (3.17)$$

where

$$\begin{aligned} L_{\beta(\omega)}^\omega &= \sum_{t=1}^T \sum_{i \in I} \sum_{j \in J} (1 - A_{ij}^\omega) x_i^t y_j^{t,\omega} \\ L_j^\omega &= \sum_{t=1}^T \sum_{i \in I} (1 - A_{ij}^\omega) x_i^t \\ \Delta_T^\omega &= \frac{\sqrt{2\bar{L}(\omega) \ln |J|}}{T} + \frac{\ln |J|}{T}. \end{aligned}$$

Now let $\bar{x}_i = \frac{1}{T} \sum_{t=1}^T x_i^t$ and $\bar{y}_j^\omega = \frac{1}{T} \sum_{t=1}^T y_j^{t,\omega}$. We have that

$$\begin{aligned} 1 - v^* &\leq \max_{x \in X} \sum_{i \in I} \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega (1 - A_{ij}^\omega) x_i \bar{y}_j^\omega \\ &= \max_{x \in X} \sum_{\omega \in \Omega} \sum_{i \in I} \sum_{j \in J} p^\omega (1 - A_{ij}^\omega) x_i \left(\frac{1}{T} \sum_{t=1}^T y_j^{t,\omega} \right) \\ &\leq \frac{1}{T} \sum_{t=1}^T \max_{x \in X} \sum_{i \in I} \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega (1 - A_{ij}^\omega) x_i y_j^{t,\omega} \\ &= \sum_{\omega \in \Omega} p^\omega \frac{1}{T} \sum_{t=1}^T \sum_{i \in I} \sum_{j \in J} (1 - A_{ij}^\omega) x_i^t y_j^{t,\omega} \\ &= \sum_{\omega \in \Omega} p^\omega \frac{L_{\beta(\omega)}^\omega}{T}, \end{aligned} \quad (3.18a)$$

where v^* is the optimal evasion probability. We also have

$$\begin{aligned} \min_{j \in J} \frac{L_j^\omega}{T} &= \min_{j \in J} \frac{1}{T} \sum_{t=1}^T \sum_{i \in I} (1 - A_{ij}^\omega) x_i^t \\ &= \min_{y^\omega \in Y} \sum_{i \in I} \sum_{j \in J} (1 - A_{ij}^\omega) \bar{x}_i y_j^\omega, \end{aligned}$$

where $Y = \{y^\omega \in \mathbb{R}_+^{|J|} : \sum_{j \in J} y_j^\omega = 1\}$. This implies

$$1 - v^* \geq \sum_{\omega \in \Omega} p^\omega \min_{j \in J} \frac{L_j^\omega}{T}.$$

Summing (3.17) over all $\omega \in \Omega$ weighted by p^ω shows that \bar{x} and \bar{y} approximate (3.5) within $\sum_{\omega \in \Omega} p^\omega \Delta_T^\omega$.

3.3 Optimal Design of a Two-Person Zero-Sum Cournot Game

In a two-person game, the players' actions are modeled as taking place either simultaneously or sequentially. Game theoreticians refer to the former case as a Cournot game and the latter as a Stackelberg game. Since in reality players usually do not act at the exact same moment in time, the modeling choice between a Cournot and a Stackelberg game typically depends upon whether one player can acquire knowledge of the other player's action. If the player who acts second (the second-mover) is not aware of the first-mover's action, then the players are effectively choosing their strategies simultaneously and a Cournot model is appropriate.

These two models can be viewed as extreme since they assume that the first-mover's actions are either all transparent or all non-transparent to the second-mover. In this section we consider models that are "hybrids" of the Cournot and Stackelberg models in that some of the first-mover's actions are revealed to the second-mover and others are not. To accomplish this, we start with a standard Cournot game but allow the first-mover to make transparent

“design” decisions which alter the payoff matrix of the game. In particular, we discuss two models: one in which the first-mover can purchase additional pure strategies for his own use during the game, and another in which the first-mover can remove some of the second-mover’s pure strategies. In both cases the first-mover is subject to a budget constraint which limits the number of pure strategies he can add or remove.

3.3.1 Investing in Premium Strategies

We consider a variant of a TPZSCG in which there are two versions of each column strategy: a “free” version which can be played at no cost and a “premium” version which can only be played if the column player invests in that strategy in the design stage. The column player has a limited budget and must make these investments via a binary variable t_j , $j \in J$, before the TPZSCG is played. If the premium strategy $j \in J$ has been selected in this design stage, and then the row player chooses $i \in I$ and the column player chooses $j \in J$, the payoff is A_{ij} . If the respective players choose i and j when premium strategy j has not been selected the payoff is B_{ij} where $A_{ij} \geq B_{ij}$. The investments t_j , $j \in J$, are discrete in nature and are transparent to the row player. After t has been selected, each player places a probability distribution on the available strategies in the induced TPZSCG. Through these discrete premium choices we optimally design a TPZSCG, from the perspective of the column player. This problem can be formulated as follows:

Indices and Sets

$i \in I$ indexes the row player's pure strategies

$j \in J$ indexes the column player's pure strategies

Data

c_j cost of investing in premium strategy j

b investment budget

A_{ij} game's payoff if the row player plays strategy i and the column player plays the premium version of strategy j

B_{ij} game's payoff if the row player plays strategy i and the column player plays the free version of strategy j

u_j upper bound on the probability that the column player plays the premium version of strategy j

Row Player's Decision Variables

x_i probability that the row player plays strategy i

Column Player's Decision Variables

t_j takes value 1 if the column player invests in strategy j and 0 otherwise

y_j probability that the column player plays the premium version of strategy j

z_j probability that the column player plays the free version of strategy j

Formulation

$$v^* = \max_t \max_{y,z} \min_x \sum_{i \in I} \sum_{j \in J} x_i (A_{ij} y_j + B_{ij} z_j) \quad (3.20a)$$

$$\text{s.t.} \quad \sum_{i \in I} x_i = 1 \quad : \theta \quad (3.20b)$$

$$\sum_{j \in J} (y_j + z_j) = 1 \quad (3.20c)$$

$$\sum_{j \in J} c_j t_j \leq b \quad (3.20d)$$

$$0 \leq x_i, \quad i \in I \quad (3.20e)$$

$$0 \leq y_j \leq t_j, \quad j \in J \quad (3.20f)$$

$$0 \leq z_j, \quad j \in J \quad (3.20g)$$

$$t_j \in \{0, 1\}, \quad j \in J. \quad (3.20h)$$

The outer maximization with respect to t selects premium strategies, subject to a knapsack constraint, (3.20d) and (3.20h). The order of the inner $\max_{y,z} \min_x$ can be equivalently written $\min_x \max_{y,z}$. The decisions x and (y, z) may be viewed as being made simultaneously and are governed by the respective convexity constraints (3.20b), (3.20e) and (3.20c), (3.20f), (3.20g). Constraint (3.20f) also disallows playing premium strategy y_j if it has not been selected via t_j . When $t_j = 1$ we are allowed to select y_j and nominally, that is limited only by 1.

3.3.2 Tightening the Formulation

In what follows, we replace constraint (3.20f) by $0 \leq y_j \leq u_j t_j, j \in J$, and seek values of $u_j, j \in J$, so that the new model is equivalent to model

(3.20), i.e., with $u_j = 1$. As we show, tightening the values of u_j , $j \in J$, improves our ability to solve model (3.20) and this plays a key role in our solution strategy described below. With θ denoting the dual variable on constraint (3.20b), we fix the decisions in the outer maximization, t, y, z , and take the dual of the inner minimization with respect to x , to arrive at the following mixed-integer program:

$$v^* = \max_{t,y,z,\theta} \theta \quad (3.21a)$$

$$\text{s.t.} \quad \theta \leq \sum_{j \in J} (A_{ij}y_j + B_{ij}z_j), \quad i \in I \quad (3.21b)$$

$$\sum_{j \in J} (y_j + z_j) = 1 \quad (3.21c)$$

$$\sum_{j \in J} c_j t_j \leq b \quad (3.21d)$$

$$0 \leq y_j \leq u_j t_j, \quad j \in J \quad (3.21e)$$

$$0 \leq z_j, \quad j \in J \quad (3.21f)$$

$$t_j \in \{0, 1\}, \quad j \in J. \quad (3.21g)$$

The parameters $u_j, j \in J$, are upper bounds on the probability the column player plays the premium version of strategy j . If $u_j = 1$ for all $j \in J$, model (3.21) tends to have a very weak linear programming relaxation: Constraint (3.21d) is redundant if we allow t to be continuous and $b \geq \max_{j \in J} c_j$. Unless we can obtain tighter upper bounds on $y_j, j \in J$, relaxing t 's integrality constraints is equivalent to allowing use of all premium strategies without investment. Therefore, naive application of a branch-and-bound algorithm is computationally ineffective. It is crucial, therefore, to make $u_j, j \in J$, as small

as possible without eliminating (all) optimal solutions to model (3.20). We refer to such u -values as being *valid*.

To find tighter valid upper bounds on y , suppose we have valid values for $u_j, j \in J$, (e.g., initially $u_j = 1, j \in J$) and a feasible solution to (3.21), and let \underline{v} be the associated objective function value. We know $\underline{v} \leq v^*$ and that if we wish to improve upon this lower bound we must achieve a payoff greater than \underline{v} for all row strategies. Therefore for each $j' \in J$ we let:

$$u_{j'} = \max_{t,y,z} y_{j'} \quad (3.22a)$$

$$\text{s.t.} \quad \underline{v} \leq \sum_{j \in J} (A_{ij}y_j + B_{ij}z_j), \quad i \in I \quad (3.22b)$$

$$(3.21c) - (3.21g). \quad (3.22c)$$

The linear programming relaxation of model (3.22) allows use of all premium strategies without prior investment. Ignoring this potential concern allows (3.22) to be solved quickly but may result in loose values of u_j . Still, our preliminary computational experience has shown that the bounds generated by (3.22) tend to be significantly less than 1 and that using these bounds to tighten model (3.21) can significantly improve its solution time when using a branch-and-bound solver.

Solving model (3.22) as a mixed-integer program generates tighter bounds on y but could result in $|J|$ auxiliary problems that are essentially as difficult to solve as model (3.21). However, we can generate an instance of model (3.22) for each $j' \in J$ and solve those instances in parallel with an instance of model (3.21). While the tightest upper bounds come from solving model (3.22) to

optimality, relaxations also produce valid upper bounds on y . So, these instances of (3.22) can periodically report their progress, specifically their best linear-programming relaxation upper bound, so far. These are then reported both to model (3.21) and to other instances of model (3.22). As using smaller u_j can greatly speed the solution time of model (3.21), all problem instances can benefit from the improved bounds. In simplest form, these computations can run on $|J|+1$ computing nodes. For practical problems the number of column strategies $|J|$ can be large, e.g., in the hundreds or much larger. Clearly for larger models, multiple instances of (3.22) are solved on each computing node.

3.3.3 Removing an Opponent's Premium Strategies

Suppose we alter the previous model by giving control of the decision variables $t_j, j \in J$, to the row player and inverting their role so that they now forbid, instead of allow, the use of the premium strategies. That is, the column player can use premium strategy j if and only if the associated $t_j = 0$. Also, let c_j be the row player's cost of removing premium strategy j from the column player's set of pure strategies. This modified problem can be formulated as

follows:

$$v^* = \min_t \min_x \max_{y,z} \sum_{i \in I} \sum_{j \in J} x_i (A_{ij} y_j + B_{ij} z_j) - \sum_{j \in J} \pi_j t_j y_j \quad (3.23a)$$

$$\text{s.t.} \quad \sum_{i \in I} x_i = 1 \quad (3.23b)$$

$$\sum_{j \in J} (y_j + z_j) = 1 \quad : \lambda \quad (3.23c)$$

$$\sum_{j \in J} c_j t_j \leq b \quad (3.23d)$$

$$0 \leq x_i, \quad i \in I \quad (3.23e)$$

$$0 \leq y_j, j \in J \quad (3.23f)$$

$$0 \leq z_j, \quad j \in J \quad (3.23g)$$

$$t_j \in \{0, 1\}, \quad j \in J. \quad (3.23h)$$

The term $-\sum_{j \in J} \pi_j t_j y_j$ in the objective guarantees that $y_j = 0$ if the corresponding $t_j = 1$, provided that π_j is sufficiently large. We fix the decisions in the outer minimizations, t , x , and take the dual of the inner maximization

with respect to y, z , to arrive at the following mixed-integer program:

$$v^* = \min_{t,x,\lambda} \lambda \quad (3.24a)$$

$$\text{s.t.} \quad \lambda \geq \sum_{i \in I} A_{ij} x_i - \pi_j t_j \quad : y_j, \quad j \in J \quad (3.24b)$$

$$\lambda \geq \sum_{i \in I} B_{ij} x_i \quad : z_j, \quad j \in J \quad (3.24c)$$

$$\sum_{i \in I} x_i = 1 \quad (3.24d)$$

$$\sum_{j \in J} c_j t_j \leq b \quad (3.24e)$$

$$0 \leq x_i, \quad i \in I \quad (3.24f)$$

$$t_j \in \{0, 1\}, \quad j \in J. \quad (3.24g)$$

The decision variables y_j are now the dual variables of constraints (3.24b). For sufficiently large π_j , setting some t_j to 1 will force the corresponding constraint in (3.24b) to be slack and the corresponding dual variable y_j to be 0. Excessively large π_j , however, can result in a weak linear programming relaxation and so we focus our efforts on making π_j as small as possible and still valid. A valid π_j has the property that $\pi_j > \sum_{i \in I} A_{ij} x_i^* - \lambda^*$ for at least one (x^*, λ^*) optimal to (3.24). To find valid values for the π_j we can solve the following problem for each $j' \in J$:

$$\pi_{j'} = \max_{t,x,\lambda} \sum_{i \in I} A_{ij'} x_i - \lambda \quad (3.25a)$$

$$\text{s.t.} \quad (3.24b) - (3.24g). \quad (3.25b)$$

Of course, any upper bound on the optimal solution to (3.25) also provides a valid choice for π_j , so we may relax the integrality constraints and solve

(3.25) as an LP or use the best remaining upper bound from an incomplete branch-and-bound tree. We may solve (3.25) without an upper bound \bar{v} on v^* , although if such a bound is available we can add the constraint $\lambda \leq \bar{v}$ to the formulation.

3.3.4 Complexity

Call the model of Sections 2.1-2.2 in which the column player invests in premium strategies, Add-TPZSCG, and call the model of Section 2.3 in which the row player can invest to forbid use of such a premium strategy Remove-TPZSCG. We now show that we can reduce the VERTEX-COVER problem to the decision versions of both Add-TPZSCG and Remove-TPZSCG and thus both problems are strongly NP-complete. The following defines the VERTEX-COVER problem:

VERTEX-COVER:

INSTANCE: Graph $G(V, E)$, a positive integer $k \leq |V|$.

QUESTION: Does there exist a subset $V' \subset V$ such that $|V'| \leq k$ and that every edge in E is adjacent to at least one vertex in V' ?

For the following, we assume unit costs to add or remove a premium strategy, $c_j = 1, j \in J$, and a payoff of zero under all free strategies, $B = 0$.

Add-TPZSCG-DECISION:

INSTANCE: Payoff matrix $A_{ij}, i \in I, j \in J$, positive integer $b \leq |J|$, and a real α .

QUESTION: Does there exist a subset $J' \subset J$ of columns in A with $|J'| = b$ and non-negative weights $y_j, j \in J'$, with $\sum_{j \in J'} y_j = 1$ such that

$$\sum_{j \in J'} A_{ij} y_j \geq \alpha,$$

for all $i \in I$?

Theorem 25. *Add-TPZSCG-DECISION is strongly NP-complete.*

Proof. We first show that Add-TPZSCG-DECISION is in NP and then reduce VERTEX-COVER, which is known to be strongly NP-complete, to Add-TPZSCG-DECISION. A polynomial-length guess for an instance of Add-TPZSCG-DECISION consists of $J' \subset J$ with $|J'| = b$ and probabilities $y_j, j \in J'$. Such a guess verifies a yes-instance of Add-TPZSCG-DECISION if the column player can guarantee a payoff of at least α by playing pure strategy j with probability y_j . To check this, we simply compare $\sum_{j \in J'} A_{ij} y_j$ to α for each $i \in I$. The number of steps required to do these comparisons is bounded by $O(|I||J|)$ so Add-TPZSCG-DECISION is in NP.

Next, we give a polynomial time reduction from VERTEX-COVER to Add-TPZSCG-DECISION. For each $e \in E$, create a row $e \in I$ in the payoff matrix A and for each $v \in V$ create a column $v \in J$. Let $A_{ev} = 1$ if edge $e \in E$ is adjacent to $v \in V$ and $A_{ev} = 0$ otherwise. Finally let $b = k$ and $\alpha = \frac{1}{k}$.

Suppose that the instance of VERTEX-COVER is a yes-instance. Say we invest in the premium strategies corresponding to the vertex cover V' (so

$J' = V'$) and play each of those strategies with probability $y_v = 1/k, v \in V'$.

Then

$$\sum_{v \in V'} A_{ev} y_v = \frac{1}{k} \sum_{v \in V'} A_{ev} \geq \frac{1}{k},$$

for each $e \in E$ since A_{ev} must be 1 for at least one $v \in V'$ since V' is a vertex cover. Therefore the corresponding instance of Add-TPZSCG-DECISION is also a yes-instance.

Conversely, if the instance of Add-TPZSCG-DECISION is a yes-instance then there exists a subset $J' \subset J$ such that at least one $v \in J'$ is positive for each $e \in I$ (otherwise we would achieve a payoff of zero under row strategy $e \in I$). Therefore $V' = J'$ forms a cardinality k vertex cover of graph G and the instance of VERTEX-COVER is also a yes-instance.

Since the above transformation can be done in polynomial time, Add-TPZSCG-DECISION is strongly NP-complete. \square

We now define a decision version of Remove-TPZSCG and show that it is also strongly NP-complete.

Remove-TPZSCG-DECISION:

INSTANCE: Payoff matrix $A_{ij}, i \in I, j \in J$, positive integer $b \leq |J|$, and a real α .

QUESTION: Does there exist a subset $J' \subset J$ of columns in A with $|J'| = b$ and non-negative weights $x_i, i \in I$ with $\sum_{i \in I} x_i = 1$ such that

$$\sum_{i \in I} A_{ij} x_i \leq \alpha,$$

for all $j \in J \setminus J'$?

Theorem 26. *Remove-TPZSCG-DECISION is strongly NP-complete.*

Proof. We show this again by a reduction from VERTEX-COVER. Confirming that a guess consisting of a subset $J' \subset J$ and probabilities $x_i, i \in I$, verifies a yes-instance of Remove-TPZSCG-DECISION involves comparing $\sum_{i \in I} A_{ij} x_i$ to α for each $j \in J \setminus J'$, which can be done in $O(|I||J|)$ time. Therefore Remove-TPZSCG-DECISION is in NP.

We can reduce VERTEX-COVER to Remove-TPZSCG-DECISION as follows. First, in the payoff matrix A create a row $v \in I$ and column $v \in J$ for each vertex $v \in V$ and a column $e \in J$ for each edge $e \in E$. Let $A_{ve} = -1$ if edge $e \in E$ is adjacent to $v \in V$ and $A_{ve} = 0$ otherwise. Also let $A_{vv'} = -1$ for $v = v'$ and $A_{vv'} = 0$ otherwise. That is, $-A$ will be the vertex-edge adjacency matrix of G concatenated with an identity matrix. Finally let $b = |V| - k$ and $\alpha = -\frac{1}{k}$.

If the instance of Remove-TPZSCG-DECISION is a yes-instance, then we can find probabilities $x_i, i \in I$, such that we achieve a payoff of at most $-1/k$ for each column strategy not in J' . For columns $v' \in J \setminus J'$ this amounts to:

$$\sum_{v \in I} A_{vv'} x_v = -x_{v'} \leq -\frac{1}{k}.$$

Clearly if $\sum_{v \in I} x_v = 1$, this inequality can only hold for k vertices, so the $|V| - k$ column strategies that we are allowed to remove must come from the strategies corresponding to the vertices of G . Then $x_v = 1/k$ for all $v \in J \setminus J'$.

With no more column strategies to remove, we must also guarantee a payoff of at most $-1/k$ for each $e \in J$. This amounts to:

$$\sum_{v \in J \setminus J'} A_{ve} x_v = \frac{1}{k} \sum_{v \in J \setminus J'} A_{ve} \leq -\frac{1}{k},$$

for all $e \in J$. This is only possible if $A_{ve} = -1$ for at least one $v \in J \setminus J'$, which implies that $J \setminus J'$ is a vertex cover with cardinality k . Therefore the instance of VERTEX-COVER is also a yes-instance.

Suppose the instance of VERTEX-COVER is a yes-instance and let $V' \subset V$ be a cardinality- k vertex cover. If for each $v \in V \setminus V'$ we remove column strategy $v \in J$ and for each $v' \in V'$ we play row strategy $v' \in I$ with probability $1/k$, then we achieve a payoff of at most $-1/k$ for each of the remaining column strategies. Since we did so by only removing $|V| - k$ column strategies, the instance of Remove-TPZSCG-DECISION must also be a yes-instance.

This transformation can be done in polynomial time so Remove-TPZSCG-DECISION is strongly NP-complete. □

3.4 Transparent and Non-transparent Assets

3.4.1 Knapsack-Constrained Non-transparent Assets

The previous section assumes that all the detectors that the interdictor deploys are non-transparent, that is, that the evader cannot determine where they are. Suppose now that the interdictor has one type of asset which is non-transparent as before but also has another type of asset which, when deployed,

is transparent (i.e., visible) to the evader. In our initial model, each checkpoint can have none, one, or both types of assets. For this we propose a two-stage model in which the interdicator deploys the transparent assets in the first stage and the interdicator and smuggler play the previously described two-person zero-sum Cournot game in the second stage.

Suppose checkpoint j does not receive a transparent asset. In this case, if it receives a non-transparent asset the evasion probability for evader ω is q_j^ω ; otherwise the evasion probability is p_j^ω . If checkpoint j does receive a transparent asset then the evasion probability for evader ω is r_j^ω if no non-transparent asset is present and s_j^ω if a non-transparent asset is present. We assume $p_j^\omega > q_j^\omega$, $p_j^\omega > r_j^\omega$, $r_j^\omega > s_j^\omega$, and $q_j^\omega > s_j^\omega$. No ordering between r_j^ω and q_j^ω is assumed. To incorporate the transparent assets into the model, we create two columns in the payoff matrix for each checkpoint $j \in J$ and evader $\omega \in \Omega$; a “premium” column which contains the p_j^ω and q_j^ω evasion probabilities and a “free” column which contains the r_j^ω and s_j^ω probabilities. Deploying a transparent asset to a checkpoint simply prevents all evaders from using the

premium column for that checkpoint. We can formulate this as follows:

Indices and Sets

$i \in I$ indexes feasible deployments of the non-transparent assets

$j \in J$ indexes checkpoints the evader can choose

$\omega \in \Omega$ indexes evader scenarios

Data

c_j^1 cost of installing a transparent asset at checkpoint j

c_j^2 cost of installing a non-transparent asset at checkpoint j

b_1 budget for installing transparent assets

b_2 budget for installing non-transparent assets

p_j^ω probability evader ω can traverse j undetected when neither type of asset is present

q_j^ω probability evader ω can traverse j undetected when there is no transparent asset present but a non-transparent asset is present

A_{ij}^ω game's payoff if evader ω selects checkpoint j and the interdictor selects deployment i with no transparent asset installed at j , i.e., $A_{ij}^\omega = q_j^\omega$ if deployment i places a non-transparent asset on checkpoint j and otherwise $A_{ij}^\omega = p_j^\omega$. $\sum_{j \in J} c_j^2 I(A_{ij}^\omega = q_j^\omega) \leq b^2, i \in I$

r_j^ω probability evader ω can traverse j undetected when there is a transparent asset present but no non-transparent asset is present

s_j^ω probability evader ω can traverse j undetected when both types of assets are present

B_{ij}^ω game's payoff if evader ω selects checkpoint j and the interdictor selects deployment i with a transparent asset installed at j , i.e., $B_{ij}^\omega = s_j^\omega$ if deployment i places a non-transparent asset on checkpoint j and otherwise $B_{ij}^\omega = r_j^\omega$. $\sum_{j \in J} c_j^2 I(B_{ij}^\omega = s_j^\omega) \leq b^2, i \in I$

Interdictor's Decision Variables

t_j indicates if the interdictor installs a transparent asset at checkpoint j

x_i probability that the interdictor chooses deployment plan i for installing non-transparent assets

Evader's Decision Variables

y_j^ω probability that evader ω chooses checkpoint j when there is no transparent asset present

z_j^ω probability that evader ω chooses checkpoint j when there is a transparent asset present

Formulation

$$v^* = \min_{t \in T} h(t), \quad (3.26)$$

where $T = \{t : \sum_{j \in J} c_j^1 t_j \leq b_1, t_j \in \{0, 1\}, j \in J\}$ and $h(t)$ is the optimal value of:

$$\max_{y, z} \min_x \sum_{i \in I} \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega x_i (A_{ij}^\omega y_j^\omega + B_{ij}^\omega z_j^\omega) \quad (3.27a)$$

$$\text{s.t.} \quad \sum_{i \in I} x_i = 1 \quad : \lambda \quad (3.27b)$$

$$\sum_{j \in J} (y_j^\omega + z_j^\omega) = 1 \quad : \theta^\omega, \quad \omega \in \Omega \quad (3.27c)$$

$$0 \leq x_i, \quad i \in I \quad (3.27d)$$

$$0 \leq y_j^\omega \leq 1 - t_j, \quad j \in J, \quad \omega \in \Omega \quad (3.27e)$$

$$0 \leq z_j^\omega, \quad j \in J, \quad \omega \in \Omega. \quad (3.27f)$$

Fixing y and z and taking the dual of the inner linear program with respect

to x , where λ is the dual variable on (3.27b), gives us the following:

$$h(t) = \max_{y,z,\lambda} \lambda \quad (3.28a)$$

$$\text{s.t.} \quad \lambda \leq \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega (A_{ij}^\omega y_j^\omega + B_{ij}^\omega z_j^\omega) \quad : x_i \quad i \in I \quad (3.28b)$$

$$\sum_{j \in J} (y_j^\omega + z_j^\omega) = 1, \quad \omega \in \Omega \quad (3.28c)$$

$$0 \leq y_j^\omega \leq 1 - t_j \quad : \pi_j^\omega(t), \quad j \in J, \omega \in \Omega \quad (3.28d)$$

$$0 \leq z_j^\omega, \quad j \in J, \omega \in \Omega. \quad (3.28e)$$

If $t_j = 0$ for some j , an evader would prefer to use the y_j^ω decision variable over z_j^ω since the associated evasion probabilities are larger. If $t_j = 1$ for some j , an evader wishing to traverse checkpoint j is forced to use the z_j^ω decision variable due to (3.28d). Unfortunately the optimal value of this optimization problem is a concave function of t since t appears on the right-hand side of the constraints. This does not make h amenable to minimization. We reformulate (3.28) with the goal of effectively moving t into the objective function. To do so, we find values $\bar{\pi}_j^\omega$ such that $\bar{\pi}_j^\omega \geq \pi_j^\omega(t)$, that is, we find upper bounds on the optimal dual variables of (3.28d). Then, applying Proposition 4 and following the same procedure we used to reformulate model (1.13) as model (1.20), we assess a penalty $\bar{\pi}_j^\omega$ if $t_j = 1$ and $y_j^\omega > 0$ and relax the upper bounds

on y in constraints (3.28d). This results in the following problem:

$$\bar{h}(t) = \max_{y,z,\lambda} \quad \lambda - \sum_{j \in J} \sum_{\omega \in \Omega} \bar{\pi}_j^\omega t_j y_j^\omega \quad (3.29a)$$

$$\text{s.t.} \quad \lambda \leq \sum_{j \in J} \sum_{\omega \in \Omega} p^\omega (A_{ij}^\omega y_j^\omega + B_{ij}^\omega z_j^\omega) \quad : x_i, \quad i \in I \quad (3.29b)$$

$$\sum_{j \in J} (y_j^\omega + z_j^\omega) = 1 \quad : \theta^\omega, \quad \omega \in \Omega \quad (3.29c)$$

$$0 \leq y_j^\omega, \quad j \in J, \omega \in \Omega \quad (3.29d)$$

$$0 \leq z_j^\omega, \quad j \in J, \omega \in \Omega. \quad (3.29e)$$

By Proposition 4, it can be shown that $h(t) = \bar{h}(t)$ for all $t \in T$ but \bar{h} is convex over the convex hull of T while h is concave over the convex hull of T . To make our model amenable to decomposition, we use \bar{h} instead of h . If we let L index the extreme points of the feasible region of model (3.29) and $(y_j^{\omega(l)}, z_j^{\omega(l)}, \lambda^{(l)})_{l \in L}$ be those extreme points then model (3.26) can be reformulated as:

$$\min_{t,\gamma} \quad \gamma \quad (3.30a)$$

$$\text{s.t.} \quad \gamma \geq \lambda^{(l)} - \sum_{j \in J} \sum_{\omega \in \Omega} \bar{\pi}_j^\omega y_j^{\omega(l)} t_j, \quad l \in L \quad (3.30b)$$

$$t \in T. \quad (3.30c)$$

To generate constraints (3.30b), we first solve a relaxation of (3.30) with a subset of the possible constraints in (3.30b) to obtain a feasible t , then substitute that t into model (3.29) and solve to generate a new extreme point $(\hat{y}, \hat{z}, \hat{\lambda})$. We then add the constraint associated with that extreme point and repeat. Optimal solutions to the relaxations of (3.30) give us lower bounds on the optimal value of (3.26), and optimal solutions to (3.29) for some fixed

$t \in T$ give upper bounds. If these bounds are within ϵ of each other we may terminate with an ϵ -optimal solution.

When we solve (3.29) for a fixed value of t , we generate the constraints in (3.29b) on an as-needed basis by solving a relaxation with a small subset of constraints to generate $(\hat{y}, \hat{z}, \hat{\lambda})$ and then solving a separation problem of the form (3.7). The solution to the separation problem either identifies a violated constraint in (3.29b) or proves that all constraints are satisfied by the current $(\hat{y}, \hat{z}, \hat{\lambda})$.

3.4.2 Cardinality-Constrained Non-transparent Assets

If $c_j^2 = 1$ for all $j \in J$, then we can compute the optimal value of the second-stage problem by solving a polynomially-sized LP as shown in Section 3.2.4. This allows us to formulate the two-stage model as a single large-scale MIP. To show this, we first write the dual of model (3.29) as:

$$\bar{h}(t) = \min_{x, \theta} \sum_{\omega \in \Omega} \theta^\omega \quad (3.31a)$$

$$\text{s.t.} \quad \theta^\omega \geq \sum_{i \in I} p^\omega A_{ij}^\omega x_i - \bar{\pi}_j^\omega t_j \quad : y_j^\omega, \quad j \in J, \omega \in \Omega \quad (3.31b)$$

$$\theta^\omega \geq \sum_{i \in I} p^\omega B_{ij}^\omega x_i \quad : z_j^\omega, \quad j \in J, \omega \in \Omega \quad (3.31c)$$

$$\sum_{i \in I} x_i = 1 \quad : \lambda \quad (3.31d)$$

$$0 \leq x_i, \quad i \in I. \quad (3.31e)$$

If we let D be the strategy-checkpoint incidence matrix as defined in Section 3.2.4, then we can express our payoff matrices as $A_{ij}^\omega = p_j^\omega - (q_j^\omega - p_j^\omega)D_{ij}$

and $B_{ij}^\omega = r_j^\omega - (s_j^\omega - r_j^\omega)D_{ij}$. Defining $\hat{x}_j = \sum_{i \in I} D_{ij}x_i$ as the probability that checkpoint j receives a non-transparent asset, we can write model (3.31) as:

$$\bar{h}(t) = \min_{\hat{x}, \theta} \sum_{\omega \in \Omega} \theta^\omega \quad (3.32a)$$

$$\text{s.t.} \quad \theta^\omega \geq p^\omega(p_j^\omega - (q_j^\omega + p_j^\omega)\hat{x}_j) - \bar{\pi}_j^\omega t_j, \quad j \in J, \omega \in \Omega \quad (3.32b)$$

$$\theta^\omega \geq p^\omega(r_j^\omega - (s_j^\omega - r_j^\omega)\hat{x}_j), \quad j \in J, \omega \in \Omega \quad (3.32c)$$

$$\sum_{j \in J} \hat{x}_j = b^2 \quad (3.32d)$$

$$0 \leq \hat{x}_j \leq 1, \quad j \in J. \quad (3.32e)$$

Dividing p^ω out of θ^ω so that θ^ω now represents the conditional evasion probability of smuggler ω and suppressing any “bar” or “hat” notation, we can write the full two-stage model as follows:

$$\min_{t, x, \theta} \sum_{\omega \in \Omega} p^\omega \theta^\omega \quad (3.33a)$$

$$\text{s.t.} \quad \theta^\omega \geq p_j^\omega + (q_j^\omega - p_j^\omega)x_j - \pi_j^\omega t_j, \quad j \in J, \omega \in \Omega \quad (3.33b)$$

$$\theta^\omega \geq r_j^\omega + (s_j^\omega - r_j^\omega)x_j, \quad j \in J, \omega \in \Omega \quad (3.33c)$$

$$\sum_{j \in J} c_j^1 t_j \leq b^1 \quad (3.33d)$$

$$\sum_{j \in J} x_j = b^2 \quad (3.33e)$$

$$t_j \in \{0, 1\}, \quad j \in J \quad (3.33f)$$

$$0 \leq x_j \leq 1, \quad j \in J. \quad (3.33g)$$

We must choose π_j^ω to be sufficiently large such that we are guaranteed that constraint (3.33b) is not binding for any j with $t_j = 1$. One possibility is to let $\pi_j^\omega = p_j^\omega$, but this can lead to a loose LP relaxation value. In the next section

we discuss a reformulation which both tightens the LP relaxation value and obviates the need for tight bounds for the parameter π_j^ω .

3.4.3 Reformulation for $p = q$

We now consider a reformulation of model (3.33) which can significantly tighten its LP relaxation. We first consider the simplified case in which the non-transparent asset does not decrease the evasion probability at a checkpoint unless a transparent asset is also installed at that checkpoint. This situation could arise, for example, if the transparent asset is a detector and the non-transparent asset is an upgrade to that detector. In this case we have $p_j^\omega = q_j^\omega$ for all $j \in J$ and $\omega \in \Omega$ and constraint (3.33b) becomes:

$$\theta^\omega \geq p_j^\omega(1 - t_j), \quad j \in J, \omega \in \Omega, \quad (3.34)$$

if we choose $\pi_j^\omega = p_j^\omega$. We can now state the following result regarding the convex hull of the polyhedron induced by constraints (3.34) and (3.33c) for a particular $\omega \in \Omega$.

Proposition 27. *Let $\Theta^\omega = \{(t, x, \theta^\omega) : \theta^\omega \geq p_j^\omega(1 - t_j), j \in J, \theta^\omega \geq r_j^\omega + (s_j^\omega - r_j^\omega)x_j, j \in J, t \in \mathbb{Z}_+^{|J|}, x \in \mathbb{R}_+^{|J|}, \theta^\omega \in \mathbb{R}_+\}$ where $0 \leq s_j^\omega \leq r_j^\omega \leq p_j^\omega \leq 1$ and let $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ be an extreme point of the convex hull of Θ^ω . Then*

- (i) $\hat{\theta}^\omega$ must equal either p_j^ω or r_j^ω for some $j \in J$ or 0;
- (ii) $\hat{t}_j = 1$ if $p_j^\omega > \hat{\theta}^\omega$ and $\hat{t}_j = 0$ otherwise; and,
- (iii) $\hat{x}_j = \frac{(\hat{\theta}^\omega - r_j^\omega)^+}{s_j^\omega - r_j^\omega}$.

Proof. Suppose that $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ is an extreme point of the convex hull of Θ^ω .

Then

$$\hat{\theta}^\omega \geq \theta_{\min} \equiv \max(\max_{j \in J} p_j^\omega (1 - \hat{t}_j), \max_{j \in J} r_j^\omega + (s_j^\omega - r_j^\omega) \hat{x}_j, 0).$$

If $\hat{\theta}^\omega > \theta_{\min}$ then the points $(\hat{t}, \hat{x}, \hat{\theta}^\omega + \epsilon)$ and $(\hat{t}, \hat{x}, \hat{\theta}^\omega - \epsilon)$ where $\epsilon = \hat{\theta}^\omega - \theta_{\min}$ are both in Θ^ω , contradicting the assumption that $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ is an extreme point. Therefore $\hat{\theta}^\omega = \theta_{\min}$. Now $\max_{j \in J} p_j^\omega (1 - \hat{t}_j)$ can only equal p_j^ω for some $j \in J$ or 0 since $p_j^\omega \leq 1$ and \hat{t}_j must be integer. To prove (i) it only remains to show that if $\hat{\theta}^\omega \neq p_j^\omega$ for all $j \in J$ and $\hat{\theta}^\omega \neq 0$, then $\hat{\theta}^\omega = r_j^\omega$ for some $j \in J$. Suppose not. Then there exists a non-empty subset $J' \subseteq J$ such that $\hat{\theta}^\omega = r_j^\omega + (s_j^\omega - r_j^\omega) \hat{x}_j > 0$ and $\hat{x}_j > 0$ for all $j \in J'$. Let $\epsilon = \min(\epsilon_1, \epsilon_2)$ where $\epsilon_1 = \hat{\theta}^\omega - \max(\max_{j \in J} p_j^\omega (1 - \hat{t}_j), \max_{j \in J \setminus J'} r_j^\omega + (s_j^\omega - r_j^\omega) \hat{x}_j, 0)$ and $\epsilon_2 = \min_{j \in J'} (r_j^\omega - s_j^\omega) \hat{x}_j$. If $e_{J'} = \sum_{j \in J'} \frac{e_j}{r_j^\omega - s_j^\omega}$, where e_j is the unit vector with the j th component equal to 1, then the points $(\hat{t}, \hat{x} + \epsilon e_{J'}, \hat{\theta}^\omega - \epsilon)$ and $(\hat{t}, \hat{x} - \epsilon e_{J'}, \hat{\theta}^\omega + \epsilon)$ are both in Θ^ω . This proves (i).

To prove (ii), note that we must have $\hat{t}_j \geq 1$ for all j such that $p_j^\omega > \hat{\theta}^\omega$ and $\hat{t}_j \geq 0$ otherwise. Suppose that for some j' , either $\hat{t}_{j'} \geq 2$ and $p_{j'}^\omega > \hat{\theta}^\omega$ or $\hat{t}_{j'} \geq 1$ and $p_{j'}^\omega \leq \hat{\theta}^\omega$. Then $(\hat{t} + e_{j'}, \hat{x}, \hat{\theta}^\omega)$ and $(\hat{t} - e_{j'}, \hat{x}, \hat{\theta}^\omega)$ are both in Θ^ω giving a contradiction. Therefore $\hat{t}_j = 1$ for all j such that $p_j^\omega > \hat{\theta}^\omega$ and $\hat{t}_j = 0$ otherwise. Similarly to prove (iii), note that $\hat{x}_j \geq \frac{(\hat{\theta}^\omega - r_j^\omega)^+}{s_j^\omega - r_j^\omega}$ for all $j \in J$ and suppose that this inequality is strict for some $j' \in J$. Then $(\hat{t}, \hat{x} + \epsilon e_{j'}, \hat{\theta}^\omega)$ and $(\hat{t}, \hat{x} - \epsilon e_{j'}, \hat{\theta}^\omega)$, where $\epsilon = \hat{x}_{j'} - \frac{(\hat{\theta}^\omega - r_{j'}^\omega)^+}{s_{j'}^\omega - r_{j'}^\omega}$, are both in Θ^ω and $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ is not an extreme point of the convex hull of Θ^ω . \square

The significance of Proposition 27 is that at an extreme point of the convex hull of Θ^ω , θ^ω can only take on at most $2|J| + 1$ distinct values at an extreme point and, given a value of θ^ω , both t and x take on unique values. Thus, the convex hull of Θ^ω has at most $2|J| + 1$ extreme points. This, coupled with the fact that the extreme directions are $(e_j, 0, 0), j \in J$, $(0, e_j, 0), j \in J$, and $(0, 0, 1)$, means that we can form a polynomially-sized representation of the convex hull of Θ^ω . Specifically, if we append a “dummy” checkpoint j_0 to J with $p_{j_0}^\omega = 0$ and let auxiliary variables u_j^ω correspond to the extreme point at which $\theta^\omega = p_j^\omega$ and v_j^ω correspond to the extreme point at which $\theta^\omega = r_j^\omega$, we can replace constraints (3.34) and (3.33c) with the following system:

$$\theta^\omega \geq \sum_{j \in J} (p_j^\omega u_j^\omega + r_j^\omega v_j^\omega), \quad \omega \in \Omega \quad (3.35a)$$

$$t_j \geq \sum_{j': p_{j'}^\omega < p_j^\omega} u_{j'}^\omega + \sum_{j': r_{j'}^\omega < p_j^\omega} v_{j'}^\omega, \quad j \in J, \omega \in \Omega \quad (3.35b)$$

$$x_j \geq \sum_{j' \in J} \left[\frac{(r_j^\omega - p_{j'}^\omega)^+}{r_j^\omega - s_j^\omega} u_{j'}^\omega + \frac{(r_j^\omega - r_{j'}^\omega)^+}{r_j^\omega - s_j^\omega} v_{j'}^\omega \right], \quad j \in J, \omega \in \Omega \quad (3.35c)$$

$$\sum_{j \in J} (u_j^\omega + v_j^\omega) = 1, \quad \omega \in \Omega \quad (3.35d)$$

$$0 \leq u_j^\omega \leq 1, \quad j \in J, \omega \in \Omega \quad (3.35e)$$

$$0 \leq v_j^\omega \leq 1, \quad j \in J, \omega \in \Omega. \quad (3.35f)$$

The auxiliary variables u_j^ω and v_j^ω choose a target evasion probability and constraints (3.35b) and (3.35c) require that the appropriate resources be deployed to meet that target. More specifically, (3.35b) requires that a transparent asset be deployed at checkpoint j if the target evasion probability is lower than p_j^ω , and (3.35c) requires that the probability that a non-transparent asset is

deployed at checkpoint j be sufficiently large if the target evasion probability is lower than r_j^ω . Finally, constraint (3.35a) chooses the appropriate evasion probability and (3.35d)-(3.35f) are standard convexity constraints. Taking into account the fact that constraint (3.35a) is tight at an optimal solution when we minimize a positively weighted sum of θ^ω , we arrive at the following reformulation of (3.33) for $p_j^\omega = q_j^\omega$:

$$\begin{aligned} \min_{u,v,t,x} \quad & \sum_{\omega \in \Omega} p^\omega \sum_{j \in J} (p_j^\omega u_j^\omega + r_j^\omega v_j^\omega) & (3.36) \\ \text{s.t.} \quad & (3.35b) - (3.35f) \\ & (3.33d) - (3.33g). \end{aligned}$$

3.4.4 Reformulation for $p \neq q$

We now present a polyhedral analysis and the resulting reformulation for (3.33) when $p \neq q$. First, we present a result regarding the convex hull of the polyhedron induced by constraints (3.33b) and (3.33c) for a particular $\omega \in \Omega$.

Proposition 28. *Let $\Theta^\omega = \{(t, x, \theta^\omega) : \theta^\omega \geq p_j^\omega(1 - t_j) + (q_j^\omega - p_j^\omega)x_j, j \in J, \theta^\omega \geq r_j^\omega + (s_j^\omega - r_j^\omega)x_j, j \in J, t \in \mathbb{Z}_+^{|J|}, x \in \mathbb{R}_+^{|J|}, \theta^\omega \in \mathbb{R}_+\}$, where $0 \leq s_j^\omega \leq r_j^\omega \leq p_j^\omega \leq 1$ and $q_j^\omega \leq p_j^\omega$, and let $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ be an extreme point of the convex hull of Θ^ω . Then*

- (i) $\hat{\theta}^\omega$ must equal either p_j^ω or r_j^ω for some $j \in J$ or 0, and
- (ii) for all $j \in J$, either $\hat{t}_j = 1$ and $\hat{x}_j = \frac{(r_j^\omega - \hat{\theta}^\omega)^+}{r_j^\omega - s_j^\omega}$ or $\hat{t}_j = 0$ and $\hat{x}_j = \frac{(p_j^\omega - \hat{\theta}^\omega)^+}{p_j^\omega - q_j^\omega}$.

Proof. Suppose that $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ is an extreme point of the convex hull of Θ^ω .

Then

$$\hat{\theta}^\omega = \max(\max_{j \in J} p_j^\omega (1 - \hat{t}_j) + (q_j^\omega - p_j^\omega) \hat{x}_j, \max_{j \in J} r_j^\omega + (s_j^\omega - r_j^\omega) \hat{x}_j, 0).$$

But $\max_{j \in J} p_j^\omega (1 - \hat{t}_j) + (q_j^\omega - p_j^\omega) \hat{x}_j$ is either no bigger than 0 or equals $p_{j'}^\omega$ for some $j' \in J$. Suppose not. Then there exists a non-empty set $J' = \{j \in J : \hat{t}_j = 0, \hat{\theta}^\omega = p_j^\omega + (q_j^\omega - p_j^\omega) \hat{x}_j, \hat{x}_j > 0\}$. Then the points $(\hat{t}, \hat{x} + \epsilon e_{J'}, \hat{\theta}^\omega - \epsilon)$ and $(\hat{t}, \hat{x} - \epsilon e_{J'}, \hat{\theta}^\omega + \epsilon)$, where $e_{J'} = \sum_{j \in J'} \frac{e_j}{p_j^\omega - q_j^\omega}$ and ϵ is sufficiently small, are both in Θ^ω . By similar logic, $\max_{j \in J} r_j^\omega + (s_j^\omega - r_j^\omega) \hat{x}_j$ is either no bigger than 0 or will equal $r_{j'}^\omega$ for some $j \in J$. This proves (i).

To prove (ii), we first note that if $\hat{t}_j \geq 2$ for any $j \in J$, then the points $(\hat{t} + e_j, \hat{x}, \hat{\theta}^\omega)$ and $(\hat{t} - e_j, \hat{x}, \hat{\theta}^\omega)$ are both in Θ^ω . Thus, for each $j \in J$, it suffices only to consider the cases $\hat{t}_j = 1$ and $\hat{t}_j = 0$. If $\hat{t}_j = 1$, then $\hat{x}_j \geq \frac{(r_j^\omega - \hat{\theta}^\omega)^+}{r_j^\omega - s_j^\omega}$, and if $\hat{t}_j = 0$, then $\hat{x}_j \geq \frac{(p_j^\omega - \hat{\theta}^\omega)^+}{p_j^\omega - q_j^\omega}$. In either case, if the inequality on \hat{x}_j is strict then the points $(\hat{t}, \hat{x} + \epsilon e_j, \hat{\theta}^\omega)$ and $(\hat{t}, \hat{x} - \epsilon e_j, \hat{\theta}^\omega)$ are both in Θ^ω and $(\hat{t}, \hat{x}, \hat{\theta}^\omega)$ cannot be an extreme point of the convex hull of Θ^ω . So the inequality on \hat{x}_j must hold with equality which proves (ii). \square

We can still write down a polynomially-sized representation of the convex hull of Θ^ω even though the set now has exponentially many extreme points. In addition to u_j^ω and v_j^ω which have the same interpretation as before, we introduce auxiliary variables $\alpha_{jj'}^\omega$ and $\beta_{jj'}^\omega$. The new variable $\alpha_{jj'}^\omega$ ($\beta_{jj'}^\omega$) will equal 1 if $u_{j'}^\omega = 1$ ($v_{j'}^\omega = 1$) and $t_{j'} = 0$. We can view these new variables as selecting from amongst the two cases referred to in part (ii) of Proposition 28. Then

the convex hull of Θ^ω can be represented as follows:

$$\theta^\omega \geq \sum_{j \in J} p_j^\omega u_j^\omega + r_j^\omega v_j^\omega, \quad \omega \in \Omega \quad (3.37a)$$

$$t_j \geq \sum_{j': p_{j'}^\omega < p_j^\omega} (u_{j'}^\omega - \alpha_{jj'}^\omega) + \sum_{j': r_{j'}^\omega < r_j^\omega} (v_{j'}^\omega - \beta_{jj'}^\omega), \quad j \in J, \omega \in \Omega \quad (3.37b)$$

$$x_j \geq \sum_{j' \in J} \left[\frac{(r_j^\omega - p_{j'}^\omega)^+}{r_j^\omega - s_j^\omega} u_{j'}^\omega + \frac{(r_j^\omega - r_{j'}^\omega)^+}{r_j^\omega - s_j^\omega} v_{j'}^\omega \right], \quad j \in J, \omega \in \Omega \quad (3.37c)$$

$$x_j \geq \sum_{j' \in J} \left[\frac{(p_j^\omega - p_{j'}^\omega)^+}{p_j^\omega - q_j^\omega} \alpha_{jj'}^\omega + \frac{(p_j^\omega - r_{j'}^\omega)^+}{p_j^\omega - q_j^\omega} \beta_{jj'}^\omega \right], \quad j \in J, \omega \in \Omega \quad (3.37d)$$

$$\alpha_{jj'}^\omega \leq u_{j'}^\omega, \quad j \in J, j' \in J, \omega \in \Omega \quad (3.37e)$$

$$\beta_{jj'}^\omega \leq v_{j'}^\omega, \quad j \in J, j' \in J, \omega \in \Omega \quad (3.37f)$$

$$\sum_{j \in J} (u_j^\omega + v_j^\omega) = 1, \quad \omega \in \Omega \quad (3.37g)$$

$$\alpha, \beta, u, v \geq 0. \quad (3.37h)$$

Then model (3.33) can be reformulated as:

$$\begin{aligned} \min_{\alpha, \beta, u, v, t, x} \quad & \sum_{\omega \in \Omega} p^\omega \sum_{j \in J} (p_j^\omega u_j^\omega + r_j^\omega v_j^\omega) \\ \text{s.t.} \quad & (3.37b) - (3.37h) \\ & (3.33d) - (3.33g). \end{aligned} \quad (3.38)$$

Chapter 4

Two-Country Network Interdiction

4.1 Introduction

Thus far, we have restricted our attention to network interdiction problems in which an interdictor installs radiation detectors at border crossings of a single country with the goal of minimizing the evasion probability of a smuggler. If the interdictor and smuggler agree on the values of the evasion probabilities and the smuggler is aware of the detector locations, then the problem can be solved in polynomial time if the smuggler's origin-destination pair is known ahead of time but is NP-complete if the origin-destination pair is known only via a probability distribution. Thus, while the stochastic version of the problem is hard to solve, we can easily compute a lower bound on the smuggler's evasion probability conditional on his origin-destination pair. If we solve a mixed-integer programming formulation of the problem via a branch-and-bound algorithm, such a lower bound can significantly tighten the LP relaxation and decrease the overall solution time.

In the one-country problem, the country in question could be either the country in which the nuclear material originates or the country to which it is being smuggled. A natural extension to this problem is to allow the installation

of detectors at border crossings of both the origin and destination countries. In Section 4.2 we show that this two-country model also has the property that it is solvable in polynomial time given the smuggler's origin-destination pair but is NP-complete if the origin-destination pair is stochastic. In Section 4.3, we formulate the stochastic version of the two-country model and present solution techniques which perform well in practice. In Section 4.4 we show that the three-country model is NP-complete even in the deterministic case, and that the two-country stochastic model remains NP-complete even when the budget constraint is dualized.

4.2 Two-Country Deterministic Network Interdiction Problem

4.2.1 Formulation

We first show that the deterministic version of the two-country problem can be solved in polynomial time by solving a series of vertex cover problems on a bipartite network. We are given a transportation network with node set N and arc set A , an origin $o \in N$, and a destination $d \in N$. For each arc $a \in A$, the probability that the smuggler can traverse the arc undetected is p_a if there is no detector installed on the arc and q_a otherwise. We let $i \in I$ index all arcs corresponding to outbound border crossings for the origin country, $j \in J$ index all arcs corresponding to inbound border crossings for the destination country, and assume that every o - d path includes exactly one arc in I and one arc in J and that only arcs in I and J can receive detectors.

Since the smuggler encounters no detectors along any path from o to the tail of an arc $i \in I$, we can precompute the maximum probability that the smuggler can reach checkpoint $i \in I$ undetected by solving a maximum-reliability path problem. By the same reasoning, we can precompute the maximum probability that the smuggler can travel undetected from the head of an arc $i \in I$ to the tail of an arc $j \in J$ and from the head of an arc $j \in J$ to the destination d . We define parameter γ_k as the product of these three probabilities where $k \in K$ indexes all possible (i, j) pairs. Figure 4.1 shows the network topology after preprocessing for the stochastic version of this problem considered in Section 4.3. The topology of the deterministic version differs only in that it includes a single origin-destination pair.




Figure 4.1: Topology of the preprocessed network for the two-country stochastic network interdiction problem. Only border crossing arcs can receive detectors.

In what follows, we refer to the elements $k \in K$ as paths since, assuming that the aforementioned maximum-reliability paths are unique and that the smuggler has perfect information, there is a one-to-one correspondence between

the elements of K and the paths the smuggler would potentially traverse. If $i(k) \in I$ and $j(k) \in J$ are the checkpoint arcs traversed by path k , then the probability the smuggler can traverse path k undetected is the product of γ_k and one of $p_{i(k)}p_{j(k)}$, $p_{i(k)}q_{j(k)}$, $q_{i(k)}p_{j(k)}$, and $q_{i(k)}q_{j(k)}$. We can formulate the two-country deterministic network interdiction problem as follows.

Indices and sets:

I set of outbound border checkpoints for the origin country
 J set of inbound border checkpoints for the destination country,
 $I \cap J = \emptyset$

$K = I \times J$ set of paths the smuggler may traverse

$i(k), j(k)$ checkpoint arcs traversed by path k

Data:

b total budget for installing detectors
 c_i, c_j cost of installing detector at border checkpoint i, j
 γ_k probability smuggler can traverse paths from o to tail of $i(k)$,
from the head of $i(k)$ to tail of $j(k)$, and from the head of $j(k)$
to d undetected
 p_i, p_j probability smuggler can traverse checkpoint i, j undetected
with no detector installed
 q_i, q_j probability smuggler can traverse checkpoint i, j undetected
with a detector installed; $q_i < p_i$ and $q_j < p_j$

Decision variables:

x_i, x_j 1 if a detector is installed at checkpoint i, j and 0 otherwise
 θ probability that the smuggler evades detection

Formulation:

$$\min_{x, \theta} \quad \theta \quad (4.1a)$$

$$\text{s.t.} \quad x \in X \quad (4.1b)$$

$$\theta \geq \gamma_k p_{i(k)} p_{j(k)} (1 - x_{i(k)} - x_{j(k)}), \quad k \in K \quad (4.1c)$$

$$\theta \geq \gamma_k q_{i(k)} p_{j(k)} (1 - x_{j(k)}), \quad k \in K \quad (4.1d)$$

$$\theta \geq \gamma_k p_{i(k)} q_{j(k)} (1 - x_{i(k)}), \quad k \in K \quad (4.1e)$$

$$\theta \geq \gamma_k q_{i(k)} q_{j(k)}, \quad k \in K, \quad (4.1f)$$

where $X = \{x \in \mathbb{B}^{|I|+|J|} : \sum_{i \in I} c_i x_i + \sum_{j \in J} c_j x_j \leq b\}$. Constraint (4.1c) states that if neither checkpoint along a path k is interdicted, then the smuggler achieves an evasion probability of at least $\gamma_k p_{i(k)} p_{j(k)}$. Constraints (4.1d) and (4.1e) handle the case in which one checkpoint along path k is interdicted and the other is not, and (4.1f) handles the case in which both checkpoints along path k are interdicted.

4.2.2 Simplifying the Model

Model (4.1) is a mixed-integer linear program and could be solved using a standard branch-and-bound solver. We do not recommend this, and instead outline a procedure for solving (4.1) in polynomial time. Before doing so, we describe some ways to reduce the size of the model and simplify notation. This is also useful for our subsequent discussion of the stochastic variant of the model. First, constraint (4.1f) is nothing more than a simple lower bound

on the variable θ and can be replaced by the single dominating constraint

$$\theta \geq \max_{k \in K} \gamma_k q_{i(k)} q_{j(k)}. \quad (4.2)$$

We can also reduce the number of constraints in (4.1d) and (4.1e). If we define $K_j = \{k \in K : j(k) = j\}$, then we can rewrite (4.1d) as follows

$$\theta \geq \gamma_k q_{i(k)} p_j (1 - x_j), \quad j \in J, k \in K_j. \quad (4.3)$$

For a fixed $j \in J$, these constraints are identical apart from the $\gamma_k q_{i(k)} p_j$ coefficient on the right-hand side. So, for each $j \in J$, we only need to include the constraint with the largest coefficient and therefore (4.1d) is dominated by

$$\theta \geq \max_{k \in K_j} \gamma_k q_{i(k)} p_j (1 - x_j), \quad j \in J. \quad (4.4)$$

Similarly, if we define $K_i = \{k \in K : i(k) = i\}$, then (4.1e) is dominated by

$$\theta \geq \max_{k \in K_i} \gamma_k p_i q_{j(k)} (1 - x_i), \quad i \in I. \quad (4.5)$$

Replacing (4.1d) and (4.1e) with (4.4) and (4.5) reduces the size of the model by replacing $2|I||J|$ constraints with $|I| + |J|$ constraints.

Now, we can write model (4.1) with a single set of constraints linking θ to x by noting that both (4.4) and (4.5) are special cases of (4.1c) with either $x_{i(k)}$ or $x_{j(k)}$ fixed to 0. To handle these special cases, we append “dummy” checkpoints i_0 and j_0 to sets I and J , respectively, and create new decision variables x_{i_0} and x_{j_0} which are both fixed to 0. For each $i \in I$, we create a path $k(i)$ with $i(k(i)) = i$ and $j(k(i)) = j_0$. We similarly create a path $k(j)$ for each

$j \in J$. Also, we define an augmented set of paths $\bar{K} = K \cup_{i \in I} \{k(i)\} \cup_{j \in J} \{k(j)\}$ and a parameter $(r_k)_{k \in \bar{K}}$ as follows:

$$r_k = \begin{cases} \gamma_k p_{i(k)} p_{j(k)} & k \in K \\ \max_{k' \in K_i} \gamma_{k'} p_i q_{j(k')} & k = k(i), i \in I \\ \max_{k' \in K_j} \gamma_{k'} q_{i(k')} p_j & k = k(j), j \in J. \end{cases} \quad (4.6)$$

Then the constraint

$$\theta \geq r_k (1 - x_{i(k)} - x_{j(k)}), \quad k \in \bar{K}, \quad (4.7)$$

includes constraint (4.1c) as well as both constraints (4.4) and (4.5). Finally, we use simple lower bound on θ , $\underline{\theta} \equiv \max_{k \in K} \gamma_k q_{i(k)} q_{j(k)}$, to tighten the coefficients of (4.7) and obtain the following model:

$$\min_{x, \theta} \quad \theta \quad (4.8a)$$

$$\text{s.t.} \quad x \in X \quad (4.8b)$$

$$\theta \geq r_k - (r_k - \underline{\theta})^+ (x_{i(k)} + x_{j(k)}), \quad k \in \bar{K} \quad (4.8c)$$

$$\theta \geq \underline{\theta}, \quad (4.8d)$$

where we redefine $X = \{x \in \mathbb{B}^{|I|+|J|} : \sum_{i \in I} c_i x_i + \sum_{j \in J} c_j x_j \leq b, x_{i_0} = x_{j_0} = 0\}$.

4.2.3 Solution Techniques

In this section, we prove that (4.8) can be solved in polynomial time. This is a direct result of the fact that the optimal value of (4.8) can only take on a modest number of values, coupled with the fact that, for any fixed $\hat{\theta}$, we

can efficiently assess the feasibility of (4.8) with the added constraint $\theta \leq \hat{\theta}$ by solving a vertex cover problem on a bipartite graph. We also propose a bisection search procedure for solving (4.8) which uses this feasibility test as a subroutine.

Theorem 29. *Model (4.8) can be solved in polynomial time.*

Proof. At an optimal solution (x^*, θ^*) to (4.8), either $\theta^* = \underline{\theta}$ or at least one constraint in (4.8c) is tight. Since the components of x^* are binary, it must hold that either $\theta^* = r_k$ for some $k \in \bar{K}$ or $\theta^* = \underline{\theta}$. Since θ^* can take on at most $|I||J| + |I| + |J| + 1$ values, we can solve (4.8) in polynomial time if, for a fixed target evasion probability $\hat{\theta} \geq \underline{\theta}$, we can either find a feasible solution $(\hat{x}, \hat{\theta})$ to (4.8) or show that (4.8) with the added constraint $\theta \leq \hat{\theta}$ is infeasible. Equivalently, we must either find an $x \in X$ satisfying

$$x_{i(k)} + x_{j(k)} \geq 1 \quad \forall k \in \bar{K} : r_k > \hat{\theta}. \quad (4.9)$$

or prove that no such $x \in X$ exists. This equivalence is due to the fact that if $x_{i(k)} + x_{j(k)} = 0$ for any $k \in \bar{K}$ with $r_k > \hat{\theta}$, then we have $\theta \geq r_k > \hat{\theta}$ by (4.8c).

In the (weighted) vertex cover problem we are given an undirected graph $G(V, E)$, weights w_v , $v \in V$, and a positive real α , and must determine whether there exists a subset $V' \subseteq V$ with $\sum_{v \in V'} w_v \leq \alpha$ such that for every edge $(i, j) \in E$ at least one of i and j belongs to V' . We can formulate a vertex cover problem that finds an $x \in X$ satisfying (4.9) if it exists as follows. We create a vertex $v_i \in V$ with weight c_i for every $i \in I$ and a vertex $v_j \in V$ with

weight c_j for every $j \in J$. Then for every $k \in \bar{K}$ such that $r_k > \hat{\theta}$, we create an edge $e_k \in E$ such that $e_k = (v_{i(k)}, v_{j(k)})$ if $i(k) \neq i_0$ and $j(k) \neq j_0$. If for some such k we have $i(k) = i_0$, then to satisfy (4.9) we must choose $x_{j(k)} = 1$. In this case, we delete the vertex $v_{j(k)}$ from V and all edges adjacent to the deleted vertex from E . Similarly, if $j(k) = j_0$ for $e_k \in E$ then we must have $x_{i(k)} = 1$ and we delete vertex $v_{i(k)}$ and all adjacent edges. Deleting a vertex is equivalent to forcing interdiction of the corresponding checkpoint. Finally, we let α be the difference between the budget b and the sum of the weights of the vertices that were deleted. If $\alpha < 0$ then model (4.8) with the additional constraint $\theta \leq \hat{\theta}$ is infeasible. Otherwise, we determine feasibility by solving the vertex cover problem. Since every $e \in E$ can be written as $e = (v_i, v_j)$ where $i \in I$ and $j \in J$, the graph we have defined is bipartite. If $c_i = c_j = 1$ for all $i \in I$ and $j \in J$, by König's theorem the vertex cover problem can be solved in polynomial time by solving the associated maximum cardinality edge matching problem. For general interdiction costs, we can solve the vertex cover problem by finding a minimum cut on a directed graph G' defined identically to G aside from the following modifications. Add a source vertex s and a sink vertex t . For all $i \in I$ add a directed edge (s, v_i) with capacity c_i , and for all $j \in J$ add a directed edge (v_j, t) with capacity c_j . Let every $e = (v_i, v_j) \in E$ be directed from v_i to v_j and have infinite capacity. Every finite-capacity s - t cut in G' corresponds to a feasible vertex cover of G with total weight equal to the value of the cut, and vice versa [1]. Thus, we can solve (4.8) in polynomial time. □

Several computational enhancements we develop for the stochastic variant of (4.8) use the construction and solution of the bipartite vertex cover problem described in Theorem 29 as a subroutine. To facilitate the development of both those computational enhancements and a bisection search for solving (4.8), we include the following pseudo-code description.

Algorithm 2: *CanCoverPaths*(S, c, b)

Input: Subset of paths $S \subseteq \bar{K}$, installation costs c_i, c_j , interdiction budget b

Output: Return true if there exists $x \in X$ satisfying $x_{i(k)} + x_{j(k)} \geq 1 \forall k \in S$, return false otherwise

For every $i \in I$ ($j \in J$) add vertex v_i (v_j) with weight c_i (c_j) to V
 $E \leftarrow \emptyset$

for all $k \in S$ **do**

if $i(k) = i_0$ **then**

$V \leftarrow V \setminus \{j(k)\}$

$b \leftarrow b - c_{j(k)}$

end if

if $j(k) = j_0$ **then**

$V \leftarrow V \setminus \{i(k)\}$

$b \leftarrow b - c_{i(k)}$

end if

end for

for all $k \in S$ **do**

if $i(k) \neq i_0$ and $j(k) \neq j_0$ and $i(k) \in V$ and $j(k) \in V$ **then**

$E \leftarrow (v_{i(k)}, v_{j(k)})$

end if

end for

if $b < 0$ **then**

return false

end if

if there exists a vertex cover of weight b or less for $G(V, E)$ **then**

return true

else

return false

end if

We conclude this section by describing a simple bisection search procedure for solving instances of (4.8). First, let $K^+ = \{k \in \bar{K} : r_k > \underline{\theta}\}$ and index the paths $k \in K^+$, $k_1, \dots, k_{|K^+|}$, such that $r_{k_1} \geq r_{k_2} \geq \dots \geq r_{k_{|K^+|}}$. Let

$\underline{i} = 1$ and $\bar{i} = |K^+|$. Then, repeat the following until $\underline{i} = \bar{i}$. For $\hat{\theta} = r_{k_i}$ where $i = \lceil (\underline{i} + \bar{i})/2 \rceil$, determine if there exists $x \in X$ that satisfies (4.9) by running Algorithm 2 with $S = \{k_1, \dots, k_{i-1}\}$. If so, then $\theta^* \leq r_{k_i}$ and let $\underline{i} = i$. If not, then $\theta^* > r_{k_i}$ and let $\bar{i} = i - 1$. If $\underline{i} = \bar{i}$ and $\underline{i} \neq |K^+|$, then $\theta^* = r_{k_{\underline{i}}}$. If $\underline{i} = \bar{i} = |K^+|$ determine if there exists $x \in X$ satisfying (4.9) for $\hat{\theta} = \underline{\theta}$ by running Algorithm 2 with $S = K^+$. If so, then $\theta^* = \underline{\theta}$. Otherwise $\theta^* = r_{k_{|K^+|}}$.

4.3 Two-Country Stochastic Network Interdiction Problem

In this section we consider a stochastic variant of the two-country network interdiction problem in which the arc evasion probabilities and the smuggler's origin-destination pair are random. We are given a finite number of smuggler scenarios $\omega \in \Omega$, each of which is realized with a known probability p^ω . Each scenario specifies an origin-destination pair (o^ω, d^ω) and arc evasion probabilities p_a^ω and q_a^ω , which are defined in the same manner as their deterministic counterparts. As such, we define parameters γ_k^ω and r_k^ω which are scenario dependent but are otherwise defined identically to γ_k and r_k , respectively. If we let decision variable θ^ω be the smuggler's evasion probability conditional on the realization of scenario ω and seek to minimize the unconditional evasion probability, then we obtain the following stochastic extension of model (4.8):

$$\min_{x, \theta} \sum_{\omega \in \Omega} p^\omega \theta^\omega \quad (4.10a)$$

$$\text{s.t. } x \in X \quad (4.10b)$$

$$\theta^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)^+(x_{i(k)} + x_{j(k)}), \quad k \in \bar{K}, \omega \in \Omega \quad (4.10c)$$

$$\theta^\omega \geq \underline{\theta}^\omega, \quad \omega \in \Omega, \quad (4.10d)$$

where $\underline{\theta}^\omega$ is a lower bound on θ^ω . Nominally we may choose

$$\underline{\theta}^\omega = \max_{k \in K} \gamma_k^\omega q_{i(k)}^\omega q_{j(k)}^\omega.$$

Alternatively, for each $\omega \in \Omega$ in turn, we can let $\underline{\theta}^\omega$ be the minimum of θ^ω subject to constraints (4.10b)-(4.10d) defined only over the current ω . We can compute these lower bounds in polynomial time by solving a single scenario problem in the form of (4.8) for each $\omega \in \Omega$. Doing so is equivalent to solving the wait-and-see problem associated with model (4.10). That is, we are computing the minimum evasion probability for each smuggler as if we know the smuggler's identity before installing the detectors. For larger instances, especially those with a small budget relative to the number of checkpoints, the computational effort spent computing tighter values of $\underline{\theta}^\omega$ is more than made up for by decreased computational effort when running a branch-and-bound algorithm to solve model (4.10).

Preliminary computational experiments revealed two main obstacles to effectively solving (4.10). First, since the number of paths can grow quadratically in the number of checkpoints, the size of constraint set (4.10c) can be

large even for modest instances. Tighter lower bounds $\underline{\theta}^\omega$ on the conditional evasion probabilities help to address this issue since constraint (4.10c) can be removed for any $k \in \bar{K}$ and $\omega \in \Omega$ such that $r_k^\omega \leq \underline{\theta}^\omega$. In Section 4.3.1, we show that by solving a sequence of bipartite vertex cover problems we can identify additional constraints in (4.10c) that cannot be binding. Doing so is especially worthwhile when solving instances with larger budget values. Second, the LP relaxation of (4.10) typically produces very weak lower bounds. In Section 4.3.2 we describe a stronger reformulation of (4.10) which is similar to the reformulation of BiSNIP from Section 2.5. Additionally, we recognize that the bounds obtained by solving the wait-and-see problem are typically loose since we do not actually know the smuggler’s identity ahead of time. In Section 4.3.3 we describe a customized branching scheme which helps to alleviate this problem.

4.3.1 Identifying Additional Non-binding Constraints

Recall that constraint (4.10c) need not be generated for any $k \in \bar{K}$ and $\omega \in \Omega$ satisfying $r_k^\omega \leq \underline{\theta}^\omega$, where $\underline{\theta}^\omega$ is any valid lower bound on θ^ω . This occurs when, for example, we have insufficient budget to interdict at least one checkpoint along every path that gives smuggler ω a higher evasion probability than that of path k . In this case, smuggler ω never traverses path k and the corresponding constraint in (4.10c) cannot be binding. The same holds true if it is impossible to force smuggler ω to traverse path k without interdicting one of the checkpoints used by path k . We state this formally as follows.

Proposition 30. *For a particular $k' \in \bar{K}$ and $\omega' \in \Omega$, constraint (4.10c) cannot be binding in a feasible solution to (4.10) unless there exists $x \in X$ satisfying*

$$x_{i(k)} + x_{j(k)} \geq 1, \quad \forall k \in \bar{K} : r_k^{\omega'} > r_{k'}^{\omega'} \quad (4.11)$$

and

$$x_{i(k')} = x_{j(k')} = 0. \quad (4.12)$$

Proof. Consider constraint (4.10c) for a fixed $k' \in \bar{K}$ and $\omega' \in \Omega$. If (4.11) does not hold, then there exists $k'' \in \bar{K}$ with $r_{k''}^{\omega'} > r_{k'}^{\omega'}$ such that $x_{i(k'')} + x_{j(k'')} = 0$. Then $\theta^{\omega'} \geq r_{k''}^{\omega'} > r_{k'}^{\omega'}$ and the constraint cannot be binding. If (4.12) does not hold, then the constraint cannot be binding since its right-hand side is at most $\underline{\theta}^{\omega'}$ and $\theta^{\omega'} \geq \underline{\theta}^{\omega'}$. \square

We can run Algorithm 2 with $S = \{k \in \bar{K} : r_k^{\omega'} > r_{k'}^{\omega'}\}$ to determine whether there exists $x \in X$ satisfying (4.11). To determine whether there exists $x \in X$ satisfying both (4.11) and (4.12), we treat checkpoints $i(k')$ and $j(k')$ as non-interdictable checkpoints, i.e., for any path in S which traverses $i(k')$ or $j(k')$, replace $i(k')$ with i_0 and $j(k')$ with j_0 , and run Algorithm 2 using the modified set S as input. We repeat this for all $k' \in \bar{K}$ and $\omega' \in \Omega$ with $r_{k'}^{\omega'} > \underline{\theta}^{\omega'}$ and delete the instance of (4.10c) corresponding to k' and ω' if Algorithm 2 returns false.

4.3.2 Reformulation

We now consider a reformulation of (4.10) which has a tighter LP relaxation. This reformulation is analogous to that for BiSNIP, which we described in Chapter 2. To make the connection between (4.10) and BiSNIP, we define a new decision variable v_k , $k \in \bar{K}$, which can equal 1 if either checkpoint along path k is interdicted and must equal 0 otherwise. That is, we have

$$\begin{aligned} v_k &\leq x_{i(k)} + x_{j(k)}, \quad k \in \bar{K} \\ 0 &\leq v_k \leq 1, \quad k \in \bar{K}, \end{aligned}$$

and constraint (4.10c) can be written as

$$\theta^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)^+ v_k, \quad k \in \bar{K}, \quad \omega \in \Omega.$$

Consider the mixed-integer set $\Theta^\omega = \{(v, \theta^\omega) : \theta^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)^+ v_k, k \in \bar{K}, \theta^\omega \geq \underline{\theta}^\omega, v \in \mathbb{Z}_+^{|\bar{K}|}\}$. We know from the analysis of an equivalent set in Section 2.5 that the convex hull of Θ^ω has at most $|\bar{K}| + 1$ extreme points and $|\bar{K}| + 1$ extreme directions. We also know how to construct a polynomially-sized description of this convex hull. We apply the same analysis to (4.10) as follows. Let $k(l, \omega) \in \bar{K}$ be an index mapping of the paths satisfying $r_{k(1, \omega)}^\omega \geq r_{k(2, \omega)}^\omega \geq \dots \geq r_{k(|\bar{K}|, \omega)}^\omega$ and $\cup_{l=1}^{|\bar{K}|} \{k(l, \omega)\} = \bar{K}$ for all $\omega \in \Omega$. Also define auxiliary variables u_k^ω which equal 1 if smuggler ω is forced to traverse a path with an evasion probability lower than that of path k and equal 0 otherwise. Then we obtain the following reformulation of (4.10):

$$\min_{x,u,\theta} \sum_{\omega \in \Omega} p^\omega \theta^\omega \quad (4.14a)$$

$$\text{s.t. } x \in X \quad (4.14b)$$

$$\theta^\omega \geq \sum_{k(l,\omega) \in \bar{K}} r_{k(l,\omega)}^\omega (u_{k(l-1,\omega)}^\omega - u_{k(l,\omega)}^\omega), \quad \omega \in \Omega \quad (4.14c)$$

$$x_{i(k)} + x_{j(k)} \geq u_k^\omega, \quad k \in \bar{K}, \quad \omega \in \Omega \quad (4.14d)$$

$$u_{k(l-1,\omega)}^\omega \geq u_{k(l,\omega)}^\omega, \quad l = 2, \dots, |\bar{K}|, \quad \omega \in \Omega \quad (4.14e)$$

$$0 \leq u_k^\omega \leq I(r_k^\omega > \underline{\theta}^\omega), \quad k \in \bar{K}, \quad \omega \in \Omega \quad (4.14f)$$

$$\theta^\omega \geq \underline{\theta}^\omega, \quad \omega \in \Omega, \quad (4.14g)$$

where $u_{k(0,\omega)}^\omega \equiv 1$ and v_k is replaced by $x_{i(k)} + x_{j(k)}$ in constraint (4.14d). A shortcoming of (4.14) is that it contains $|\bar{K}||\Omega|$ variables not present in (4.10), which only has $|I| + |J| + |\Omega|$ variables. Our computational experience has shown that the LP relaxation of (4.14) can be considerably more challenging to solve than that of (4.10). We attribute this to the fact that even an instance of (4.14) with a modest number of checkpoints and scenarios can have a large number of variables since the number of paths is quadratic in the number of checkpoints. On the other hand, the number of variables in (4.10) is linear in the number of checkpoints and the number of scenarios. For large-scale instances, we were unable to solve the LP relaxation of (4.14) in a reasonable amount of time. For the corresponding instances of (4.10), the LP relaxation solved quickly but produced lower bounds too weak to prove optimality.

We obtain a formulation with a tighter LP relaxation than (4.10) but

with fewer variables than (4.14) by including fewer constraints in the definitions of the sets Θ^ω . Specifically, we define $\Theta^\omega = \{(v, \theta^\omega) : \theta^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)^+ v_k, k \in K^\omega, \theta^\omega \geq \underline{\theta}^\omega, v \in \mathbb{Z}_+^{|\bar{K}|}\}$ where $K^\omega \subseteq \bar{K}, \omega \in \Omega$. Let $k'(l, \omega) \in K^\omega$ be an index mapping of the paths satisfying $r_{k'(1, \omega)}^\omega \geq r_{k'(2, \omega)}^\omega \geq \dots \geq r_{k'(|K^\omega|, \omega)}^\omega$ and $\cup_{l=1}^{|K^\omega|} \{k'(l, \omega)\} = K^\omega$ for all $\omega \in \Omega$. Then the following is a valid reformulation of (4.10):

$$\min_{x, u, \theta} \sum_{\omega \in \Omega} p^\omega \theta^\omega \quad (4.15a)$$

$$\text{s.t. } x \in X \quad (4.15b)$$

$$\theta^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)^+ (x_{i(k)} + x_{j(k)}), \quad k \in \bar{K} \setminus K^\omega, \quad \omega \in \Omega \quad (4.15c)$$

$$\theta^\omega \geq \sum_{k'(l, \omega) \in K^\omega} r_{k'(l, \omega)}^\omega (u_{k'(l-1, \omega)}^\omega - u_{k'(l, \omega)}^\omega), \quad \omega \in \Omega \quad (4.15d)$$

$$x_{i(k)} + x_{j(k)} \geq u_k^\omega, \quad k \in K^\omega, \quad \omega \in \Omega \quad (4.15e)$$

$$u_{k'(l-1, \omega)}^\omega \geq u_{k'(l, \omega)}^\omega, \quad l = 2, \dots, |K^\omega|, \quad \omega \in \Omega \quad (4.15f)$$

$$0 \leq u_k^\omega \leq I(r_k^\omega > \underline{\theta}^\omega), \quad k \in K^\omega, \quad \omega \in \Omega \quad (4.15g)$$

$$\theta^\omega \geq \underline{\theta}^\omega, \quad \omega \in \Omega, \quad (4.15h)$$

where $u_{k'(0, \omega)}^\omega \equiv 1$. For each $\omega \in \Omega$, constraint (4.10c) for paths $k \in K^\omega$ are included in the definition of Θ^ω and are enforced by constraints (4.15d) - (4.15g), which describe the convex hull of Θ^ω . For paths $k \in \bar{K} \setminus K^\omega$ constraint (4.10c) remains in its original form as (4.15c).

We determine which paths to include in each K^ω based on the following observation. Let $(\hat{x}, \hat{u}, \hat{\theta})$ be an optimal solution to the LP relaxation of (4.14), and let $\hat{v}_k = \min(\hat{x}_{i(k)} + \hat{x}_{j(k)}, 1)$ for all $k \in \bar{K}$. By induction on constraints

(4.14d) and (4.14e), we have that

$$\hat{u}_{k(l,\omega)}^\omega \leq \min_{1 \leq l' \leq l} \hat{v}_{k(l',\omega)}, \quad l = 1, \dots, |\bar{K}|, \quad \omega \in \Omega. \quad (4.16)$$

Since increasing a component of u can only decrease the objective value, we assume (4.16) holds with equality. If $\hat{v}_{k(l,\omega)} \geq \min_{1 \leq l' \leq l-1} \hat{v}_{k(l',\omega)}$ for some l and ω , then $\hat{u}_{k(l,\omega)}^\omega = \hat{u}_{k(l-1,\omega)}^\omega$, and constraint (4.14c) is equivalent to (4.15d) with $K^\omega = \bar{K} \setminus \{k(l,\omega)\}$. This implies that the inclusion of $k(l,\omega)$ in K^ω would not strengthen the LP relaxation of (4.15), and suggests the following routine for populating K^ω , $\omega \in \Omega$. Solve the LP relaxation of model (4.15) with $K^\omega = \emptyset$, $\omega \in \Omega$, and let $(\hat{x}, \hat{u}, \hat{\theta})$ be the optimal solution. Compute \hat{v} , and for every $\omega \in \Omega$ and $l = 1, \dots, |\bar{K}|$, add $k(l,\omega)$ to K^ω only if $\hat{v}_{k(l,\omega)} < \min_{1 \leq l' \leq l-1} \hat{v}_{k(l',\omega)}$. Resolve (4.15) and augment K^ω until $\hat{v}_{k(l,\omega)} \geq \min_{1 \leq l' \leq l-1} \hat{v}_{k(l',\omega)}$ for all $\omega \in \Omega$ and $k(l,\omega) \in \bar{K} \setminus K^\omega$ or until some other termination criteria is met. In practice, this results in only a modest number of paths being included in each K^ω , and so (4.15) typically solves faster than either (4.10) or (4.14) for challenging instances.

4.3.3 Branching

We note from our computational experience that instances of (4.10) typically become more challenging to solve as the budget for installing detectors increases. With unit interdiction costs, some of the most challenging instances are those with a budget, b , greater than half the checkpoints, $|I| + |J|$, which is counterintuitive since the number of feasible solutions, $|X|$, decreases when the budget exceeds that threshold. We reconcile this apparent inconsistency

with the fact that the lower bounds $\underline{\theta}^\omega$ obtained by solving the wait-and-see problem decrease as the budget increases. Since we calculate $\underline{\theta}^\omega$ by assuming that we dedicate all of our detectors to minimizing smuggler ω 's evasion probability, these bounds become particularly weak for instances with large budgets in which the smugglers do not rank the paths in a similar fashion. In this section, we present a customized branching scheme which addresses this issue. This branching scheme is based on the following result.

Proposition 31. *Consider model (4.10) and let $S \subset \bar{K}$. Then either*

$$x_{i(k)} + x_{j(k)} \geq 1 \quad \forall k \in S, \quad (4.17)$$

or

$$\theta^\omega \geq \min_{k \in S} r_k^\omega \quad \forall \omega \in \Omega. \quad (4.18)$$

Proof. Suppose $x_{i(k')} + x_{j(k')} = 0$ for some $k' \in S$. Then by (4.10c) we have $\theta^\omega \geq r_{k'}^\omega \geq \min_{k \in S} r_k^\omega$ for all $\omega \in \Omega$. \square

Proposition 31 states that either we interdict at least one checkpoint along every path in S , as in (4.17), or every smuggler can freely traverse at least one path in S , as in (4.18). This suggests a branching scheme in which we branch on whether (4.17) or (4.18) holds. The advantage to branching in this fashion is that the lower bounds $\underline{\theta}^\omega$ can be tightened in both subproblems if we intelligently choose the subset S . This is clear for the subproblem in which (4.18) holds as we can choose $\underline{\theta}^\omega = \min_{k \in S} r_k^\omega$. For the subproblem in which (4.17) holds, we must allocate at least one detector to every path in S ,

and so we can compute a tighter $\underline{\theta}^\omega$ for any smuggler ω who does not prefer to traverse any paths in S .

To elaborate further, for every $\omega \in \Omega$ let $\underline{\theta}^\omega(S)$ be the optimal value of the following problem:

$$\min_{x, \theta^\omega} \theta^\omega \quad (4.19a)$$

$$\text{s.t. } x \in X \quad (4.19b)$$

$$\theta^\omega \geq r_k^\omega - (r_k^\omega - \underline{\theta}^\omega)^+(x_{i(k)} + x_{j(k)}), \quad k \in \bar{K} \quad (4.19c)$$

$$\theta^\omega \geq \underline{\theta}^\omega \quad (4.19d)$$

$$x_{i(k)} + x_{j(k)} \geq 1, \quad k \in S. \quad (4.19e)$$

Problem (4.19) is a single scenario problem of the form (4.8), but with the added constraint (4.19e). Just as we discuss for (4.8) in Section 4.2.3, it holds that at an optimal solution $(x^*, \theta^{\omega*})$ to (4.19), $\theta^{\omega*}$ can take on at most $|\bar{K}| + 1$ distinct values. Therefore given S we can compute $\underline{\theta}^\omega(S)$ in polynomial time if we can determine whether problem (4.19) with the added constraint $\theta^\omega \leq \hat{\theta}^\omega$ is feasible in polynomial time. But (4.19) $\theta^\omega \leq \hat{\theta}^\omega$ is feasible if and only if there exists $x \in X$ satisfying

$$x_{i(k)} + x_{j(k)} \geq 1, \quad \forall k \in \{k' \in \bar{K} : r_{k'}^\omega > \hat{\theta}^\omega\} \cup S. \quad (4.20)$$

Condition (4.20) is equivalent to condition (4.9), and so we can determine feasibility by running Algorithm 2. If (4.17) holds then we can let $\underline{\theta}^\omega = \underline{\theta}^\omega(S)$ for every $\omega \in \Omega$. Of course, we may also add (4.17) to the subproblem as constraints, which may further tighten the relaxation. Constraints of the form

(4.17) are similar in spirit to the so-called *supervalid inequalities* developed by Israeli and Wood [18] for the deterministic network interdiction problem on a general network. In that setting the supervalid inequalities were shown to significantly improve computational efficiency.

We give a high-level description of the branching scheme as follows. First, branch on the disjunction (4.17) versus (4.18), and, for each subproblem generated, recompute the lower bounds $\underline{\theta}^\omega$, $\omega \in \Omega$. Then, choose the subproblem with the minimum $\sum_{\omega \in \Omega} p^\omega \underline{\theta}^\omega$ and branch recursively unless that subproblem satisfies some termination criteria. For our computational experiments, we set a maximum depth and size of the tree. If a subproblem exceeds the depth threshold, we do not perform customized branching on that subproblem, and if the number of outstanding subproblems in the tree exceeds a threshold, we terminate customized branching and solve each outstanding subproblem with a commercial branch-and-bound solver.

When branching on the disjunction (4.17) versus (4.18), the choice of S is critical in order to guarantee that each subproblem has tighter lower bounds than its parent. The choice of a larger S tends to strengthen the bounds for the subproblem associated with (4.17), and a smaller S tends to strengthen those of the subproblem associated with (4.18). We resolve this tradeoff in the following way. Let $\bar{S}_1, \dots, \bar{S}_M$ be the subsets for which (4.17) is enforced and $\underline{S}_1, \dots, \underline{S}_N$ be the subsets for which (4.18) is enforced at the

current subproblem. Then as the lower bounds on θ^ω we can choose

$$\underline{\theta}^\omega = \max(\underline{\theta}^\omega(\bar{S}), \max_{1 \leq n \leq N} \min_{k \in \underline{S}_n} r_k^\omega), \quad \omega \in \Omega,$$

where $\bar{S} = \cup_{m=1}^M \bar{S}_m$. To assist in determining a subset of paths to branch on, we define a function of S for each subproblem which estimates the progress made in that subproblem. For the subproblem in which (4.17) holds we use the size of S because, as indicated, a larger S tends to strengthen this subproblem. For the subproblem in which (4.18) holds, we use the following function:

$$Value(S) = \sum_{\omega \in \Omega} p^\omega (\min_{k \in S} r_k^\omega - \underline{\theta}^\omega)^+. \quad (4.21)$$

We can view this function as the sum of increases in the lower bound on θ^ω , weighted by scenario probability. Ideally, we would branch on a subset that solves

$$\max_{S \subset \bar{K} \setminus \bar{S}: |S|=t} Value(S), \quad (4.22)$$

for some $t \in \{1, \dots, |\bar{K} \setminus \bar{S}|\}$. Doing so would guarantee that we cannot increase $Value(S)$ without decreasing the size of S . Solving (4.22) exactly is out of the question since it is at least as hard as BiSNIP, which is known to be NP-complete. Instead, for each $t = 1, \dots, |\bar{K} \setminus \bar{S}|$ we greedily approximate the solution to (4.22) by computing

$$Value_t = \max_{k \in \bar{K} \setminus (S_{t-1} \cup \bar{S})} Value(S_{t-1} \cup \{k\}), \quad (4.23)$$

where $S_0 = \emptyset$, and letting $S_t = S_{t-1} \cup \{k^*\}$, where k^* is a maximizer in (4.23).

In order to promote progress in both subproblems, we choose to branch on the

subset S_{t^*} where

$$t^* \in \operatorname{argmax}_{1 \leq t \leq |\bar{K} \setminus \bar{S}|} t \cdot \text{Value}_t.$$

A pseudo-code representation of the customized branch-and-bound algorithm can be found in Appendix B.

4.3.4 Computational Results

In this section, we discuss results from two-country model instances with origins in Russia and destinations in the US. We consider 7 origins in Russia and 3 destinations in the US, giving a total of $|\Omega| = 21$ threat scenarios. We assume that the smuggler uses a motor crossing to leave Russia, travels to either Mexico or Canada via sea or air, and then uses a motor crossing to enter the US. Figure 4.2a shows the 303 Russian checkpoints and Figure 4.2b shows the 143 US checkpoints we consider. We assume that detectors are perfectly reliable and have unit installation costs.

We test four approaches: (1) Solve (4.10) (BASE); (2) Solve (4.15) (REF); (3) Solve (4.15) eliminating non-binding constraints as described in Section 4.3.1 (REF-PRUNE); (4) Solve (4.15) via the customized branch-and-bound scheme described in Section 4.3.3, eliminating non-binding constraints (REF-PRUNE-C). For each approach, we use wait-and-see bounds to tighten coefficients and remove constraints. Table 4.1 reports solution times in seconds for various budget values b . All MIPs were solved with CPLEX 10.1 with a relative tolerance of 10^{-4} .

To determine the value of solving the two-country problem, for budget



(a) *Russia*



(b) *U.S.*

Figure 4.2: Motor-crossing checkpoints for a two-country instance

values $b = 1, \dots, 100$, we solve two one-country problems, one in which we restrict ourselves to installing detectors only at Russian checkpoints and another in which we restrict ourselves to installing detectors only at US checkpoints.

b	BASE	REF	REF-PRUNE	REF-PRUNE-C
10	2	4	1	37
20	37	19	7	54
30	94	44	16	64
40	2340	413	285	152
50	1732	222	87	70
60	×	×	5935	170
70	×	×	×	968
80	×	×	×	2304
90	×	×	×	1852
100	×	×	×	2849

Table 4.1: Solution times in seconds for Russia-US model instances with perfectly reliable detectors. × indicates that the solution time exceeded 2 hours.

Figure 4.3 plots the objective value of better of the two one-country solutions and the objective value of the two-country solution versus the budget. The objective function values are scaled to one if no detectors are installed, that is, the y -axis is the ratio of the evasion probability when installing a number of detectors to that when no detectors are installed. Each one-country problems solved in a matter of seconds. While the computational effort required to solve a two-country problem is typically much greater than that of a one-country problem, for some budget levels we observe two-country solutions with objective values that are upwards of 2% less than that of the best one-country solution.

We obtain the largest drops in evasion probability by interdicting checkpoints in areas where checkpoints are sparse. For small budgets it is optimal to interdict checkpoints around the Great Lakes and the checkpoints entering

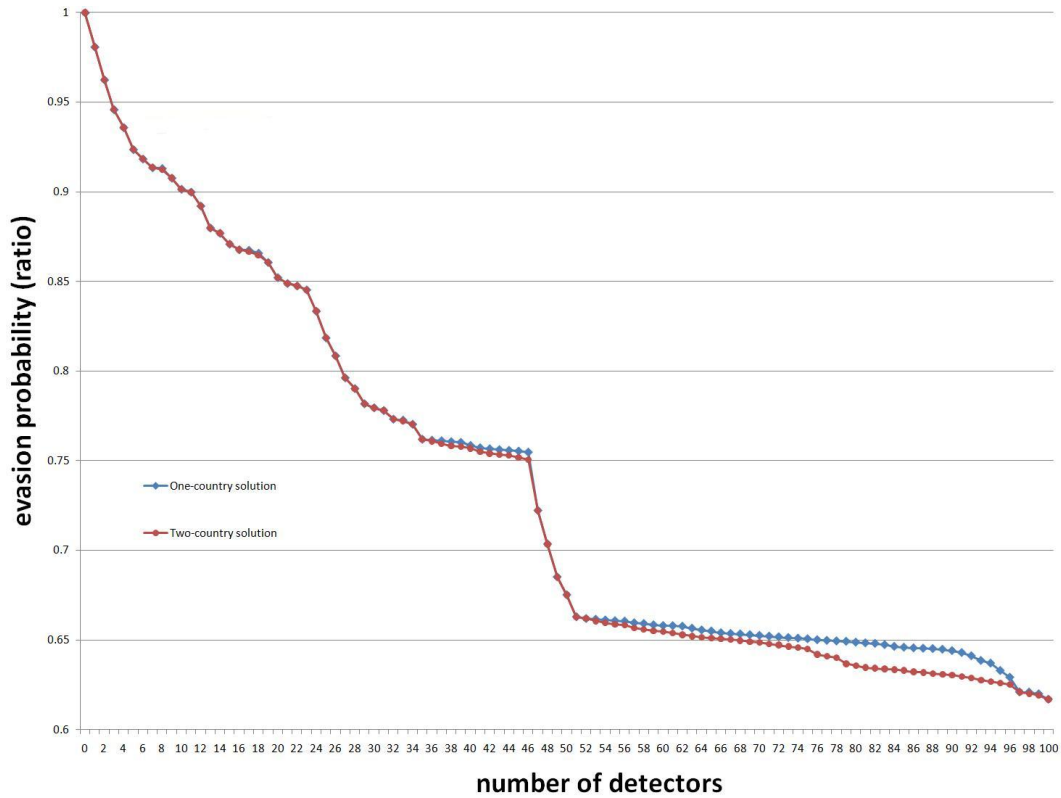
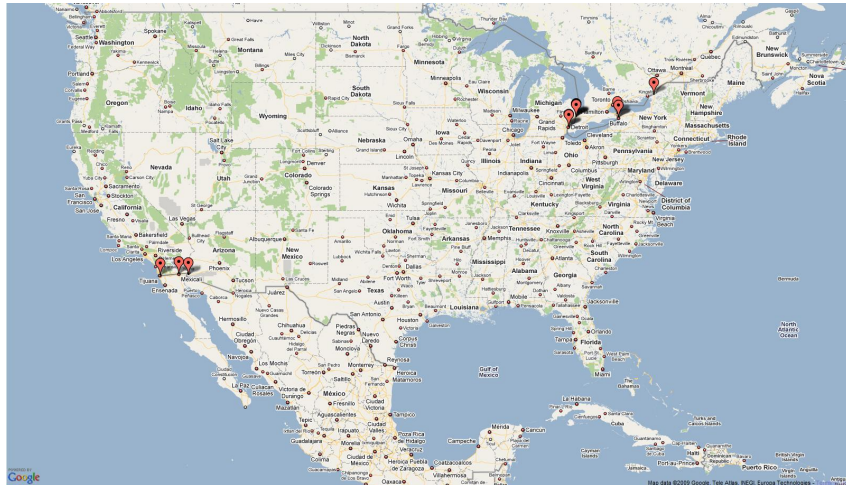


Figure 4.3: The improvement factor as a function of the number of detectors installed for the better of the two one-country solutions and the two-country solution for the Russia-US model.

California from Mexico, see for example the $b = 13$ solution in Figure 4.4a. We see a large drop in evasion probability from $b = 46$ to $b = 51$, when we obtain a large enough budget to interdict all checkpoints in Canada west of Lake Huron and all checkpoints entering California from Mexico as in the $b = 51$ solution in Figure 4.4b. The largest gap between the evasion probability of the best one-country solution and that of the best two-country solution occurs when $b = 88$, for which it is optimal to interdict several checkpoints along the US

border and to interdict every checkpoint along Russia's border with Finland (see Figure 4.5).



(a) $b = 13$



(b) $b = 51$

Figure 4.4: Optimal solutions to the Russia-US model instance with perfectly reliable detectors.



Figure 4.5: Optimal solution to the Russia-US model for $b = 88$.

4.4 Complexity

While most of the problems considered in this dissertation are NP-complete, some special cases, and variants, can be solved in polynomial time. For example, BiSNIP is known to be NP-complete, but as discussed in Section 2.6 can be solved efficiently if we dualize the budget constraint. Also, the two-country problem considered in this chapter is NP-complete in the stochastic setting but as we saw in Section 4.2.3 polynomially solvable in the deterministic setting. In general, the computational complexity of our class of network interdiction problems depends strongly on the number of countries whose borders we may interdict. A natural question, therefore, is whether we can extend the efficient algorithms that exist for the two special cases mentioned here to problems with additional countries. In this section, we answer this question in the negative. That is, we show that the two-country stochastic network interdiction problem is NP-complete even with a dualized budget constraint, and the three-country deterministic network interdiction problem is NP-complete.

4.4.1 Two-Country Stochastic Network Interdiction Problem with Dualized Budget Constraint

We show that the Lagrangian relaxation of model (4.10) with the budget constraint $\sum_{i \in I} c_i x_i + \sum_{j \in J} c_j x_j \leq b$ dualized is NP-complete, even if $c_i = c_j = 1$ for all $i \in I$ and $j \in J$, $r_k^\omega \in \{0, 1\}$, and $r_k^\omega = 0 \quad \forall k \in \bar{K} \setminus K$. That is to say, we assume detectors have unit costs are perfectly reliable and

that the probability the smuggler is caught by indigenous law enforcement is zero. Let $K_+^\omega = \{k \in K : r_k^\omega = 1\}$ denote the set of paths that smuggler ω has access to and consider the following model.

$$\min_{x, \theta} \quad z = \sum_{\omega \in \Omega} p^\omega \theta^\omega + \lambda \left(\sum_{i \in I} x_i + \sum_{j \in J} x_j \right) \quad (4.24a)$$

$$\text{s.t.} \quad \theta^\omega \geq 1 - x_{i(k)} - x_{j(k)}, \quad k \in K_+^\omega, \quad \omega \in \Omega \quad (4.24b)$$

$$\theta^\omega \geq 0, \quad \omega \in \Omega \quad (4.24c)$$

$$x \in \{0, 1\}^{|I|+|J|}. \quad (4.24d)$$

The decision problem of (4.24) is to determine if there exists (x, θ) satisfying (4.24b)-(4.24d) and $z \leq \alpha$ for some target α . We show that the decision problem is NP-complete via a reduction from (unweighted) vertex cover. That is, we show: Given a graph $G(V, E)$ and a positive integer $n \leq |V|$, there exists $V' \subseteq V$ with $|V'| \leq n$ such that every edge in E is adjacent to at least one vertex in V' if and only if there exists a solution to a transformed instance of the decision problem of (4.24).

We transform an instance of vertex cover to an instance of the decision problem of (4.24) as follows. For every $v \in V$, create checkpoints $\bar{i}(v), \underline{i}(v), \bar{i}'(v), \underline{i}'(v) \in I$ and $\bar{j}(v), \underline{j}(v), \bar{j}'(v), \underline{j}'(v) \in J$. Create a scenario $\omega' \in \Omega$ with $p^{\omega'} = \epsilon_1$ and for every $v \in V$ let $(\underline{i}(v), \bar{j}(v)), (\bar{i}(v), \underline{j}(v)) \in K_+^{\omega'}$ and for every $(v_1, v_2) \in E$ let $(\bar{i}(v_2), \bar{j}(v_1)) \in K_+^{\omega'}$. For every $v \in V$, create scenarios $\bar{\omega}(v), \underline{\omega}(v) \in \Omega$ with $p^{\bar{\omega}(v)} = \epsilon_3$ and $p^{\underline{\omega}(v)} = \epsilon_3 + \epsilon_4$ and let $(\bar{i}(v), \bar{j}'(v)), (\bar{i}'(v), \bar{j}(v)) \in K_+^{\bar{\omega}(v)}$ and $(\underline{i}(v), \underline{j}'(v)), (\underline{i}'(v), \underline{j}(v)) \in K_+^{\underline{\omega}(v)}$. Finally,

let $\lambda = \epsilon_2$ and

$$\alpha = 2|V|\epsilon_2 + |V|\epsilon_3 + n\epsilon_4, \quad (4.25)$$

where $\epsilon_1, \dots, \epsilon_4 > 0$ satisfy:

$$\epsilon_1 + 2|V|\epsilon_3 + |V|\epsilon_4 = 1 \quad (4.26a)$$

$$\epsilon_1 > 2|V|\epsilon_2 + |V|\epsilon_3 + n\epsilon_4 = \alpha \quad (4.26b)$$

$$\epsilon_2 > |V|\epsilon_3 + n\epsilon_4 \quad (4.26c)$$

$$\epsilon_3 > n\epsilon_4. \quad (4.26d)$$

Equation (4.26a) guarantees that p^ω , $\omega \in \Omega$, is a valid probability measure. The necessity of inequalities (4.26b) - (4.26d) becomes apparent in the following result. The transformed problem is a yes-instance if and only if there exists (x, θ) satisfying:

$$\alpha \geq \epsilon_1 \theta^{\omega'} + \epsilon_3 \sum_{v \in V} \theta^{\bar{\omega}(v)} + (\epsilon_3 + \epsilon_4) \sum_{v \in V} \theta^{\omega(v)} + \epsilon_2 \left(\sum_{i \in I} x_i + \sum_{j \in J} x_j \right) \quad (4.27a)$$

$$\theta^{\omega'} \geq 1 - x_{\underline{i}(v)} - x_{\bar{j}(v)}, \quad v \in V \quad (4.27b)$$

$$\theta^{\omega'} \geq 1 - x_{\bar{i}(v)} - x_{\underline{j}(v)}, \quad v \in V \quad (4.27c)$$

$$\theta^{\omega'} \geq 1 - x_{\bar{i}(v_2)} - x_{\bar{j}(v_1)}, \quad (v_1, v_2) \in E \quad (4.27d)$$

$$\theta^{\bar{\omega}(v)} \geq 1 - x_{\bar{i}(v)} - x_{\bar{j}'(v)}, \quad v \in V \quad (4.27e)$$

$$\theta^{\bar{\omega}(v)} \geq 1 - x_{\bar{j}'(v)} - x_{\bar{j}(v)}, \quad v \in V \quad (4.27f)$$

$$\theta^{\omega(v)} \geq 1 - x_{\underline{i}(v)} - x_{\underline{j}'(v)}, \quad v \in V \quad (4.27g)$$

$$\theta^{\omega(v)} \geq 1 - x_{\underline{i}'(v)} - x_{\underline{j}(v)}, \quad v \in V \quad (4.27h)$$

$$x \in \{0, 1\}^{|I|+|J|}. \quad (4.27i)$$

Figure 4.6 shows the paths available to each smuggler in the transformed instance of (4.24) for $V = \{1, 2, 3\}$ and $E = \{(1, 2), (2, 3)\}$.

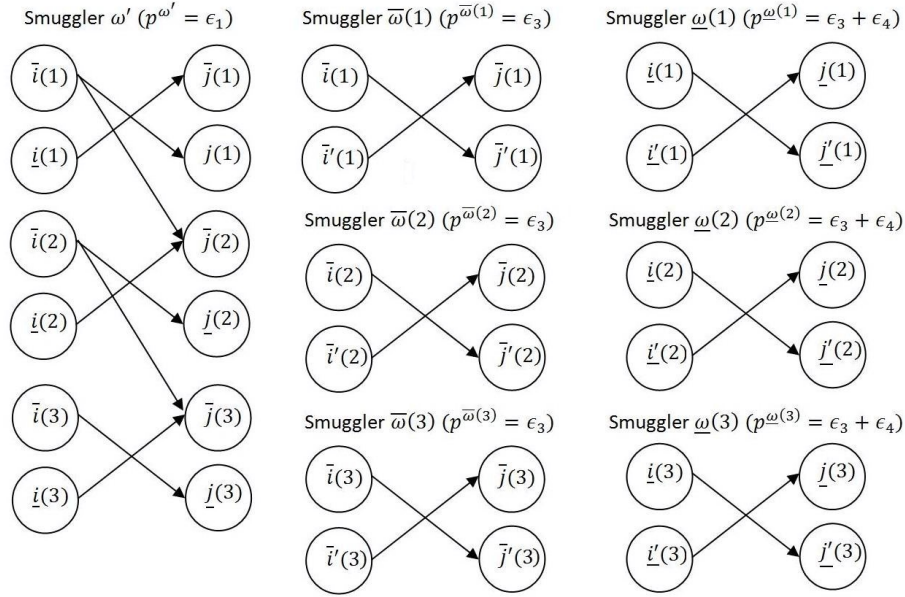


Figure 4.6: Transformed instance of (4.24) for $V = \{1, 2, 3\}$ and $E = \{(1, 2), (2, 3)\}$. Each arrow indicates a path that the smuggler has access to under the given scenario. All available paths have nominal evasion probability 1.

Informally, we connect a solution to (4.27) and a solution to the vertex cover problem in the following way. The checkpoints $\bar{i}(v)$ and $\bar{j}(v)$ are interdicted if and only if vertex v is included in V' . Similarly, the checkpoints $\underline{i}(v)$ and $\underline{j}(v)$ are interdicted if and only if vertex v is not included in V' . In the former case, smuggler $\bar{\omega}(v)$ is always detected and smuggler $\underline{\omega}(v)$ is never detected, and in the later case the opposite is true. The $|V|\epsilon_3$ term in (4.25) ensures that we detect either smuggler $\bar{\omega}(v)$ or smuggler $\underline{\omega}(v)$ for each $v \in V$, the $k\epsilon_4$

term ensures that at most k of the smugglers $\underline{\omega}(v)$, $v \in V$, evade detection, and the $2|V|\epsilon_2$ term ensures that we interdict no more than $2|V|$ checkpoints. This guarantees that every vertex is either included or not included in V' , and that V' includes at most k vertices. Finally, we choose $p^{\omega'} = \epsilon_1 > \alpha$ in order to force detection of smuggler ω' . This forces $\theta^{\omega'} = 0$, which guarantees V' is a valid vertex cover by constraint (4.27d). The following result formalizes this connection.

Lemma 32. *Let (x, θ) satisfy (4.27). Then*

- (a) $x_{\bar{i}(v)} = x_{\underline{i}(v)} = x_{\bar{j}(v)} = x_{\underline{j}(v)} = 0$ for all $v \in V$;
- (b) either $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 1$ and $x_{\underline{i}(v)} = x_{\underline{j}(v)} = 0$ or $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 0$ and $x_{\underline{i}(v)} = x_{\underline{j}(v)} = 1$ for all $v \in V$;
- (c) if $V' \subseteq V$ indexes all v satisfying $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 1$, then $|V'| \leq n$;
- (d) for every $(v_1, v_2) \in E$, either $x_{\bar{i}(v_1)} = x_{\bar{j}(v_1)} = 1$ or $x_{\bar{i}(v_2)} = x_{\bar{j}(v_2)} = 1$.

Proof. We first show that we must have

$$x_{\underline{i}(v)} + x_{\bar{j}(v)} \geq 1 \quad \forall v \in V \tag{4.28a}$$

$$x_{\bar{i}(v)} + x_{\underline{j}(v)} \geq 1 \quad \forall v \in V. \tag{4.28b}$$

If not, then $\theta^{\omega'} \geq 1$ by (4.27b) - (4.27c) and $z \geq \epsilon_1 > \alpha$ by (4.26b). Summing the inequalities in (4.28) gives us

$$\sum_{v \in V} (x_{\bar{i}(v)} + x_{\underline{i}(v)} + x_{\bar{j}(v)} + x_{\underline{j}(v)}) \geq 2|V|, \tag{4.29}$$

which, coupled with the fact that $\cup_{v \in V} \{\underline{i}(v), \bar{i}(v)\} \subset I$ and $\cup_{v \in V} \{\underline{j}(v), \bar{j}(v)\} \subset J$, implies

$$\sum_{i \in I} x_i + \sum_{j \in J} x_j \geq 2|V|. \quad (4.30)$$

If inequality (4.30) is strict, then we have $z \geq (2|V| + 1)\epsilon_2 > 2|V|\epsilon_2 + |V|\epsilon_3 + n\epsilon_4 = \alpha$ by (4.26c). Therefore

$$\sum_{i \in I} x_i + \sum_{j \in J} x_j = 2|V|. \quad (4.31)$$

Inequality (4.29) implies that if any of $x_{\bar{i}(v)}, x_{\underline{i}(v)}, x_{\bar{j}(v)}, x_{\underline{j}(v)}$ were equal to 1 then (4.31) would be violated. Therefore part (a) holds.

To prove part (b), we show that for all $v \in V$, exactly one of $\theta^{\bar{\omega}(v)}$ and $\theta^{\omega(v)}$ equals 0. From (4.28) we have that $x_{\bar{i}(v)} + x_{\underline{i}(v)} + x_{\bar{j}(v)} + x_{\underline{j}(v)} \geq 2 \quad \forall v \in V$. If this inequality were strict for some $v \in V$, then (4.31) would be violated. Therefore

$$x_{\bar{i}(v)} + x_{\underline{i}(v)} + x_{\bar{j}(v)} + x_{\underline{j}(v)} = 2 \quad \forall v \in V. \quad (4.32)$$

Suppose that for some $v \in V$ both $\theta^{\bar{\omega}(v)} < 1$ and $\theta^{\omega(v)} < 1$. Inequalities (4.27e) - (4.27f), coupled with the fact that $x_{\bar{j}(v)} = x_{\underline{j}(v)} = 0$, imply $x_{\bar{i}(v)} = x_{\underline{i}(v)} = 1$, and inequalities (4.27g) - (4.27h), coupled with the fact that $x_{\bar{i}(v)} = x_{\underline{i}(v)} = 0$, imply $x_{\bar{j}(v)} = x_{\underline{j}(v)} = 1$. This contradicts (4.32) so we must have either $\theta^{\bar{\omega}(v)} \geq 1$ or $\theta^{\omega(v)} \geq 1$ for all $v \in V$. But we cannot have both $\theta^{\bar{\omega}(v)} \geq 1$ and $\theta^{\omega(v)} \geq 1$ for any $v \in V$, since this implies $z \geq 2|V|\epsilon_2 + (|V| + 1)\epsilon_3 + \epsilon_4 > 2|V|\epsilon_2 + |V|\epsilon_3 + (n + 1)\epsilon_4 > \alpha$ by (4.26d). So either $\theta^{\bar{\omega}(v)} = 0$ which forces $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 1$ and $x_{\underline{i}(v)} = x_{\underline{j}(v)} = 0$, or $\theta^{\omega(v)} = 0$ which forces $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 0$ and $x_{\underline{i}(v)} = x_{\underline{j}(v)} = 1$. This proves part (b).

To prove part (c), note that part (b) implies that $\theta^{\bar{\omega}(v)} \geq 1$ for all $v \in V \setminus \bar{V}$ and $\theta^{\omega(v)} \geq 1$ for all $v \in V'$. Therefore, $z \geq 2|V|\epsilon_2 + |V|\epsilon_3 + |V'|\epsilon_4 > \alpha$ unless $|V'| \leq k$. Finally, we have already shown that we cannot have $\theta^{\omega'} \geq 1$, so by (4.27d) we must have either $x_{\bar{i}(v_1)} = 1$ or $x_{\bar{j}(v_2)} = 1$. By part (b), this implies that either $x_{\bar{i}(v_1)} = x_{\bar{j}(v_1)} = 1$ or $x_{\bar{i}(v_2)} = x_{\bar{j}(v_2)} = 1$, which proves part (d). \square

Theorem 33. *The decision version of (4.24) is NP-complete.*

Proof. We first establish that the decision version of (4.24) belongs to the class NP. Note that a polynomial-length guess consists of a subset $S \subseteq I \cup J$ of interdicted checkpoints. We can determine in polynomial time whether or not such a guess verifies an instance of the decision version of (4.24) as a yes-instance provided that we can efficiently evaluate the evasion probability for each smuggler given that every checkpoint in subset S is interdicted. Computing the evasion probability for a particular smuggler can be done in polynomial time by a complete enumeration of all paths.

We must show that the original vertex cover instance is a yes-instance if and only if the transformed instance of the decision problem of (4.24) is a yes-instance. Suppose the vertex cover instance is a yes-instance. Then there exists $V' \subseteq V$ with $|V'| \leq n$ such that for every $(v_1, v_2) \in E$, either $v_1 \in V'$ or $v_2 \in V'$. Given V' , we construct a solution to (4.27) as follows. For every $v \in V$, let $x_{\bar{i}(v)} = x_{\bar{i}'(v)} = x_{\bar{j}(v)} = x_{\bar{j}'(v)} = 0$. For every $v \in V'$, let $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 1$ and $x_{\bar{i}'(v)} = x_{\bar{j}'(v)} = 0$, and let $\theta^{\bar{\omega}(v)} = 0$ and $\theta^{\omega(v)} = 1$. For

every $v \in V \setminus V'$, let $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 0$ and $x_{\underline{i}(v)} = x_{\underline{j}(v)} = 1$, and let $\theta^{\bar{\omega}(v)} = 1$ and $\theta^{\underline{\omega}(v)} = 0$. Finally, since for every $(v_1, v_2) \in E$ either $x_{\bar{i}(v_2)} = 1$ or $x_{\bar{j}(v_2)} = 1$, we can let $\theta^{\omega'} = 0$. Since $|V'| \leq n$, we have $z \leq 2|V|\epsilon_2 + |V|\epsilon_3 + k\epsilon_4 = \alpha$, and thus the instance of the decision problem of (4.24) is a yes-instance.

To show the reverse direction, suppose the transformed instance of the decision problem of (4.24) is a yes-instance. Then there exists (x, θ) satisfying (4.27). We know by Lemma 32(b) that either $x_{\bar{i}(v)} = x_{\bar{j}(v)} = 1$ and $x_{\underline{i}(v)} = x_{\underline{j}(v)} = 0$ or vice-versa. Let $V' = \{v \in V : x_{\bar{i}(v)} = x_{\bar{j}(v)} = 1\}$. By Lemma 32(c) we know that $|V'| \leq k$, and by Lemma 32(d) we know that for every $(v_1, v_2) \in E$, either $v_1 \in V'$ or $v_2 \in V'$. So V' is a valid vertex cover and thus the vertex cover instance is a yes-instance. \square

4.4.2 Three-Country Deterministic Network Interdiction Problem

For some stochastic network interdiction problems, the bounds provided by the solution to the wait-and-see problem can be used to improve computational performance. Since these bounds are computed by solving a single-scenario problem for each smuggler scenario, the effectiveness of such a strategy is linked to our ability to efficiently solve the deterministic version of the interdiction problem. In the deterministic setting, we can solve the one-country problem via a greedy algorithm and the two-country problem via a minimum cut problem. In this section, we prove that it is NP-complete to solve the three-country deterministic interdiction problem, even with unit interdiction costs and perfectly reliable detectors.

We describe the three-country deterministic interdiction problem as follows. Let $h \in H$, $i \in I$, and $j \in J$ index sets of checkpoint arcs in a transportation network $G(N, A)$. First, an interdicator chooses b checkpoint arcs to remove from the network. Then, a smuggler with a given origin $o \in N$ and destination $d \in N$ chooses a o - d path which maximizes the probability that he avoids detection. Without being detected, the smuggler can traverse a path from the origin o to the tail of a checkpoint arc $h \in H$ with probability p_o^h , from the head of checkpoint arc $h \in H$ to the tail of checkpoint arc $i \in I$ with probability p_h^i , from the head of checkpoint arc $i \in I$ to the tail of checkpoint arc $j \in J$ with probability p_i^j , and from the head of checkpoint arc $j \in J$ to the destination d with probability p_j^d . The smuggler may traverse a checkpoint arc without being detected with probability 1 unless the interdicator removes the arc. A checkpoint arc becomes impassible when removed. Let $K = H \cup I \cup J$, and let decision variable $x_k = 1$ if checkpoint arc $k \in K$ is interdicted and $x_k = 0$ otherwise. We define a subset of checkpoints $K_0 \subset K$, which indexes all checkpoints which cannot receive detectors, i.e., $x_k = 0$, $k \in K_0$. Note that we can form an equivalent model in which every checkpoint can receive a detector by creating a total of $b + 1$ copies of every checkpoint in K_0 . We modify a dynamic programming based LP formulation of the maximum reliability path problem in which π_k , $k \in K$, is the probability that the smuggler can traverse path from the tail of checkpoint arc $k \in K$ to the destination d without being detected, and π_o is the probability that the smuggler can traverse an o - d path

undetected. The objective of the interdicator is to minimize π_o as follows:

$$\min_{x, \pi} \quad \pi_o \quad (4.33a)$$

$$\text{s.t.} \quad x \in X \quad (4.33b)$$

$$\pi_j = (p_j^t - x_j)^+, \quad j \in J \quad (4.33c)$$

$$\pi_i = (\max_{j \in J} p_i^j \pi_j - x_i)^+, \quad i \in I \quad (4.33d)$$

$$\pi_h = (\max_{i \in I} p_h^i \pi_i - x_h)^+, \quad h \in H \quad (4.33e)$$

$$\pi_o = \max_{h \in H} p_o^h \pi_h. \quad (4.33f)$$

where $X = \{x \in \mathbb{B}^{|K|} : \sum_{k \in K} x_k \leq b, x_k = 0, k \in K_0\}$.

The decision problem of (4.33) is to determine if there exists (x, π) satisfying (4.33b)-(4.33f) and $\pi_s \leq \alpha$ for some target α . We show that the decision problem is NP-complete via reduction from (unweighted) vertex cover. In the undirected graph $G(V, E)$ associated with the vertex cover instance, we represent every edge $e \in E$ as $e = (v_1(e), v_2(e))$ where $v_1(e)$ and $v_2(e)$ are ordered arbitrarily. We also define $E_l(v) = \{e \in E : v_l(e) = v\}$ for $l = 1, 2$ and $v \in V$. We transform an instance of vertex cover into an instance of the decision problem of (4.33) as follows. For every $v \in V$, create checkpoints $h(v) \in H$ and $i_0(v) \in I \cap K_0$. For every $e \in E$, create checkpoints $i(e) \in I$ and $j(e) \in J$. Finally, create checkpoints $h_0 \in H \cap K_0$ and $j_0 \in J \cap K_0$, and choose $\alpha = \epsilon^2$, where $0 < \epsilon < 1$, and $b = n + |E|$. The transformed instance of the decision problem of (4.33) is a yes-instance if and only if there exists

(x, π) satisfying:

$$x \in X \quad (4.34a)$$

$$\pi_o \leq \epsilon^2 \quad (4.34b)$$

$$\pi_{j(e)} = 1 - x_{j(e)}, \quad e \in E \quad (4.34c)$$

$$\pi_{j_0} = \epsilon \quad (4.34d)$$

$$\pi_{i(e)} = \left(\max(\pi_{j(e)}, \pi_{j_0}) - x_{i(e)} \right)^+, \quad e \in E \quad (4.34e)$$

$$\pi_{i_0(v)} = \max_{e \in E_2(v)} \epsilon \pi_{j(e)}, \quad v \in V \quad (4.34f)$$

$$\pi_{h(v)} = \left(\max \left(\max_{e \in E_1(v)} \pi_{i(e)}, \pi_{i_0(v)} \right) - x_{h(v)} \right)^+, \quad v \in V \quad (4.34g)$$

$$\pi_{h_0} = \max_{e \in E} \epsilon \pi_{i(e)} \quad (4.34h)$$

$$\pi_o = \max \left(\max_{v \in V} \pi_{h(v)}, \pi_{h_0} \right). \quad (4.34i)$$

We connect a solution to the vertex cover instance and a solution to (4.34) in the following way. If (4.34) is feasible, then there exists a solution (x, π) with $\sum_{v \in V} x_{h(v)} \leq n$ in which the set $V' = \{v \in V : x_{h(v)} = 1\}$ is a valid vertex cover. And, given a vertex cover V' , there exists a solution (x, π) to (4.34) with $x_{h(v)} = 1$ if and only if $v \in V'$, $x_{i(e)} = 1$ if $v_1(e) \notin V'$, and $x_{j(e)} = 1$ if $v_2(e) \notin V'$. Since either $v_1(e) \in V'$ or $v_2(e) \in V'$ for any feasible vertex cover V' , we have that exactly $n + |E|$ checkpoints are interdicted. If $v_1(e)$ and $v_2(e)$ are both in V' , we arbitrarily choose one of $x_{i(e)}$ and $x_{j(e)}$ to equal 1. Figure 4.7 shows the network corresponding to (4.34) for $V = \{1, 2, 3\}$ and $E = \{e_1, e_2\}$ where $e_1 = (1, 2)$ and $e_2 = (2, 3)$. For $n = 2$, there exists a vertex cover $V' = \{1, 2\}$. The corresponding solution to the interdiction problem is

to interdict checkpoints $h(1)$, $h(2)$, $j(e_2)$ and either $i(e_1)$ or $j(e_1)$. The only remaining paths in the residual network go through checkpoints h_0 and j_0 and have reliability ϵ^2 .




Figure 4.7: Network corresponding to (4.34) for $V = \{1, 2, 3\}$ and $E = \{e_1, e_2\}$ where $e_1 = (1, 2)$ and $e_2 = (2, 3)$. Solid arrows indicate arcs with reliability 1, while dotted arrows indicate arcs with reliability ϵ .

We formalize this connection in the following results.

Lemma 34. *If system (4.34) is feasible, then there exists (x, π) satisfying both (4.34) and:*

$$x_{i(e)} + x_{j(e)} = 1, \quad e \in E \quad (4.35a)$$

$$\sum_{v \in V} x_{h(v)} \leq n. \quad (4.35b)$$

Proof. Suppose (x, π) satisfies (4.34). By (4.34b) and (4.34i) we must have $\pi_{h_0} \leq \epsilon^2$, which is true only if $\pi_{i(e)} \leq \epsilon$ for all $e \in E$, which in turn is true only if $x_{i(e)} + x_{j(e)} \geq 1$ for all $e \in E$ by (4.34c) and (4.34e). Let $E' =$

$\{e \in E : x_{i(e)} + x_{j(e)} = 2\}$. Then we can perturb (x, π) in the following way and maintain feasibility of (4.34). For every $e \in E'$, let $x_{j(e)} = 0$ and $x_{h(v_2(e))} = 1$. This can only decrease $\sum_{k \in K} x_k$ and so we maintain $x \in X$. For every $e \in E'$ set $\pi_{j(e)} = 1$ to maintain feasibility of (4.34c), and set $\pi_{i_0(v_2(e))} = \epsilon$ to maintain feasibility of (4.34f). Note that we still have $x_{i(e)} = 1$ for all $e \in E'$, so constraint (4.34e) is unaffected by the perturbation. Similarly (4.34g) is unaffected even when $v = v_2(e)$ since we set $x_{h(v_2(e))} = 1$. The perturbed solution is feasible to (4.34) and satisfies (4.35a). Condition (4.35a) implies $\sum_{e \in E} x_{i(e)} + x_{j(e)} = |E|$, and therefore condition (4.35b) also holds since $\sum_{k \in K} x_k \leq n + |E|$. \square

Theorem 35. *The decision version of (4.33) is NP-complete.*

Proof. We first establish that the decision version of (4.33) belongs to the class NP. A polynomial-length guess consists of $S \subseteq H \cup I \cup J$ with $|S| = b$. Given that all checkpoints in subset S are interdicted, we can compute the smuggler's evasion probability in polynomial time by a complete enumeration of all paths. We can then verify whether or not S verifies an instance of the decision version of (4.33) as a yes-instance by comparing the evasion probability to α .

We show that (4.34) is feasible if and only if there exists a vertex cover of size n or less. Let $V' \subseteq V$ be a vertex cover with $|V'| \leq n$. Then we construct a solution to (4.34) as follows. Let $x_{h(v)} = 1$ if $v \in V'$ and let $x_{h(v)} = 0$ otherwise. For every $e \in E$, if $v_1(e) \notin V'$ let $x_{i(e)} = 1$ and $x_{j(e)} = 0$, and let $x_{i(e)} = 0$ and $x_{j(e)} = 1$ otherwise. Note that since V' is a vertex cover,

for every $e \in E$ either $v_1(e)$ or $v_2(e)$ is in V' , and so $x_{j(e)} = 1$ if $v_2(e) \notin V'$. We have $\sum_{k \in K} x_k = n + |E|$ and so $x \in X$. Since $x_{i(e)} + x_{j(e)} = 1$ for all $e \in E$, we can set $\pi_{i(e)} = \epsilon(1 - x_{i(e)})$ and therefore $\pi_{h_0} = \epsilon^2 \max_{e \in E} (1 - x_{i(e)}) \leq \epsilon^2$. So $\pi_o \leq \epsilon^2$ if $\pi_{h(v)} \leq \epsilon^2$. But $\pi_{h(v)} > \epsilon^2$ implies that $x_{h(v)} = 0$ and that either $x_{i(e)} = 0$ for some $e \in E_1(v)$ or $x_{j(e)} = 0$ for some $e \in E_2(v)$. This leads to a contradiction, since $x_{h(v)} = 0$ implies that $v \notin V'$, and thus $x_{i(e)} = 1$ if $v_1(e) \notin V'$ and $x_{j(e)} = 1$ if $v_2(e) \notin V'$.

Next, we show that if there exists a solution to (4.34), then the vertex cover instance is a yes-instance. By Lemma 34, if (4.34) is feasible, then there exists a solution satisfying (4.35a) and (4.35b). Let (x, π) be such a solution. Then for every $v \in V$, either $x_{h(v)} = 1$ or $\pi_{i(e)} \leq \epsilon^2$ for all $e \in E_1(v)$ by (4.34b) and (4.34g). But $\pi_{i(e)} \leq \epsilon^2$ only if $x_{i(e)} = 1$, so

$$x_{h(v_1(e))} + x_{i(e)} \geq 1, \quad e \in E. \quad (4.36)$$

Similarly, for every $v \in V$, either $x_{h(v)} = 1$ or $\pi_{i_0(v)} \leq \epsilon^2$. But $\pi_{i_0(v)} \leq \epsilon^2$ only if $x_{j(e)} = 1$ for all $e \in E_2(v)$ and so,

$$x_{h(v_2(e))} + x_{j(e)} \geq 1, \quad e \in E. \quad (4.37)$$

Since for every $e \in E$, either $x_{i(e)}$ or $x_{j(e)}$ equals 0 by (4.35a), we have that either $x_{h(v_1(e))}$ or $x_{h(v_2(e))}$ equals 1. So $V' = \{v \in V : x_{h(v)} = 1\}$ is a valid vertex cover, and $|V'| \leq n$ by (4.35b). \square

Table 4.2 gives a summary of the complexity results of the Stackelberg games considered in this dissertation. All problems in P remain polynomially-

solvable even with arbitrary interdiction costs, and all NP-complete problems remain hard even with unit interdiction costs.

	Number of Countries		
	1	2	3
Deterministic	P	P	NP-complete
Stochastic with Soft Budget	P	NP-complete	NP-complete
Stochastic with Hard Budget	NP-complete	NP-complete	NP-complete

Table 4.2: Complexity landscape of the maximum-reliability network interdiction problem.

Chapter 5

Conclusion

This dissertation has developed models and solution techniques for a class of stochastic network interdiction problems. In these models, an interdicator installs detectors on arcs in a network subject to a budget constraint, while a smuggler selects a path in the residual network. The interdicator's goal is to minimize the reliability of the smuggler's chosen path. Relevant smuggler characteristics such as the origin-destination pair, mass and type of material being smuggled, and the thickness of shielding are known only through a probability distribution at the time the detectors are installed. The models considered vary in the number of countries in which the interdicator can install detectors and whether the two parties act simultaneously or sequentially. The work in this dissertation was motivated by the Second Line of Defense (SLD) Program, which is a cooperative program between the US DOE and the Russian Federation State Customs Committee. The SLD Program aims to minimize the risk of illicit trafficking of nuclear material, equipment and technology.

Chapter 2 considers a Stackelberg game in which the interdicator can only install detectors at border checkpoints of a single country. The single

country in question will typically be either a country from which we expect material may be stolen, such as Russia, or a country that we wish to keep smugglers from entering, such as the United States. We present conditions under which smuggler scenarios with similar attributes may be aggregated. While the problem is NP-complete with a hard budget constraint, the problem becomes solvable in polynomial time when the budget constraint is dualized. This implies that solutions on the concave envelope of the efficient frontier can be found in polynomial time. We use the fact that the decrease in the smuggler's evasion probability as a function of the set of interdicted checkpoints is supermodular to show that solutions on the concave envelope are nested. A naive mixed-integer programming (MIP) formulation of the problem can lead to loose linear-programming (LP) relaxations. We present an extended formulation based on a polyhedral analysis which tightens the LP relaxation and develop an associated branch-and-bound algorithm which utilizes easily computed wait-and-see bounds and performs well on particularly challenging instances.

Chapter 3 considers a two-person zero-sum Cournot game in which the interdictor and the smuggler act simultaneously. The challenge here lies in the fact that the interdictor must place a probability distribution over an exponentially-sized set of pure strategies. We show that in the single-country case if the detectors have unit installation costs, we may determine the value of the game by solving a polynomially-sized LP in which the decision variables are the marginal probabilities that a checkpoint receives a detector. We present an

easily-implementable version of the weighted majority algorithm to find a joint distribution which approximates the marginals. Finally, we present a model in which the interdicator can install two types of assets; the first type is visible to the smuggler and the second is not. This model may be appropriate if, for example, the interdicator can “upgrade” some subset of the installed detectors or if the interdicator can install “decoy” detectors.

Chapter 4 extends the results for the Stackelberg game considered in Chapter 2 to the case in which the interdicator can install detectors at border checkpoints of both the origin and destination country. If the smuggler characteristics are known before the detectors are installed, the problem can be solved in polynomial time by solving a sequence of vertex cover problems on bipartite graphs. Thus, the wait-and-see bounds for the stochastic problem are easily obtained. We use these bounds to tighten the LP relaxation of the associated MIP formulation. We conclude with complexity results for the two-country problem with a dualized budget constraint and for the three-country problem. These results fill out the complexity landscape for the Stackelberg version of the problem.

The main contributions of this dissertation lie in both the development of a Cournot model and associated solution techniques for the maximum-reliability stochastic network interdiction problem, and significant algorithmic advances for the Stackelberg model. In particular, the customized branch-and-bound scheme developed for the one- and two-country models allows us to solve significantly larger problem instances than was possible using previ-

ous methods. We also make several connections between network interdiction models and other research areas, i.e., the selection problem, supermodularity, and nestedness, which to our knowledge had not been made previously.

Appendices

Appendix A

Customized Branch-and-Bound for BiSNIP

The following is a pseudo-code representation of the customized branch-and-bound algorithm for solving BiSNIP. For a subproblem P in the branch-and-bound tree, we define $\bar{S}(P) = \{\bar{S}_1, \dots, \bar{S}_m\}$ as the set of subsets $\bar{S}_i \subset K$, $i = 1, \dots, m$, for which $x_k = 1, k \in \cup_{i=1}^m \bar{S}_i$ is enforced, and $\underline{S}(P) = \{\underline{S}_1, \dots, \underline{S}_n\}$ as the set of subsets $\underline{S}_i \subset K$, $i = 1, \dots, n$, for which $\sum_{k \in \underline{S}_i} x_k \leq |\underline{S}_i| - 1, i = 1, \dots, n$ is enforced. We use a standard last-in, first-out stack to store the subproblems in the tree.

Algorithm 3: *GetFixedVariables*(P)

Input: Problem P

Output: Set \underline{U} of all (k, ω) pairs corresponding to u_k^ω which should be fixed to 0 in problem P

Let \bar{S} be the union of the elements of $\bar{S}(P)$

Let $\underline{S}_1, \dots, \underline{S}_n$ be the elements of $\underline{S}(P)$

$\underline{U} \leftarrow \emptyset$

for all $\omega \in \Omega, k \in K$ **do**

for $i = 1, \dots, n$ **do**

if $\underline{S}_i \subseteq K_k^\omega$ **then**

 Add (k, ω) to \underline{U}

end if

end for

if $\sum_{k' \in K_k^\omega \cup \bar{S}} c_{k'} > b$ **then**

 Add (k, ω) to \underline{U}

end if

end for

return \underline{U}

Algorithm 4: *GetCriticalSubset*(\hat{u})

Input: Partial solution \hat{u} to the LP relaxation of a BiSNIP instance

Output: Subset of checkpoints to branch on S

$S_0 \leftarrow \emptyset$

for $t = 1, \dots, b$ **do**

 Compute $S_t \in \operatorname{argmax}_{k \in K \setminus S_{t-1}} \operatorname{Loss}(S_{t-1} \cup \{k\})$

$\operatorname{Loss}_t \leftarrow \operatorname{Loss}(S_t)$

end for

Let $t^* \in \operatorname{argmax}_{1 \leq t \leq b} \operatorname{Loss}_t$

$S \leftarrow S_{t^*}$

return S

Algorithm 5: *BranchAndBound*(p, r, c, b)

Input: Scenario probabilities $p^\omega > 0$, evasion probabilities $r_k^\omega \geq 0$, detector installation costs $c_k \geq 1$, installation budget b

Output: Optimal installation plan x^* , minimum evasion probability z^*

for all $\omega \in \Omega$ **do**

Sort the components of r_k^ω in decreasing order

for $i = 1, \dots, |K|$ **do**

Let $k(i, \omega) \in K$ denote the i th checkpoint in the sorted list

Let $K_{k(i, \omega)}^\omega = \{k(i', \omega) : 1 \leq i' \leq i\}$

Compute $s_{k(i, \omega)}^\omega = r_{k(i, \omega)}^\omega - r_{k(i+1, \omega)}^\omega$ where $r_{k(|K|+1, \omega)}^\omega \equiv 0$

end for

end for

Create problem P with $\bar{S}(P) = \underline{S}(P) = \emptyset$

Create an empty stack of problems and push P onto the stack

$LB \leftarrow -\infty$

while stack not empty **do**

Pop problem P off the stack

$\underline{U} \leftarrow \text{GetFixedVariables}(P)$

Let (\hat{x}, \hat{u}) be the optimal solution to the LP relaxation of (2.29) with the added constraints $x_k = 1, k \in \bar{S}$, and $u_k^\omega = 0, (k, \omega) \in \underline{U}$

$\hat{z} \leftarrow \sum_{\omega \in \Omega} \sum_{k \in K} p^\omega s_k^\omega \hat{u}_k^\omega$

if $\hat{z} > LB$ **then**

if \hat{x} is integral **then**

$LB \leftarrow \hat{z}$

$x^* \leftarrow \hat{x}$

else

$S \leftarrow \text{GetCriticalSubset}(\hat{u})$

Create problems $\bar{P} \leftarrow P$ and $\underline{P} \leftarrow P$

Add S to $\bar{S}(\bar{P})$ and to $\underline{S}(\underline{P})$

Push \underline{P} onto stack

Push \bar{P} onto stack

end if

end if

end while

$z^* \leftarrow \sum_{\omega \in \Omega} p^\omega r_{k(1, \omega)}^\omega - LB$

return x^*, z^*

Appendix B

Customized Branch-and-Bound for the Two-Country Stochastic Network Interdiction Problem

The following is a pseudo-code representation of the customized branch-and-bound algorithm for solving the two-country stochastic network interdiction problem. For a subproblem P in the branch-and-bound tree, we define $\bar{S}(P) = \{\bar{S}_1, \dots, \bar{S}_m\}$ as the set of subsets $\bar{S}_i \subset \bar{K}$, $i = 1, \dots, m$, for which $x_{i(k)} + x_{j(k)} \geq 1, k \in \cup_{i=1}^m \bar{S}_i$ is enforced, and $\underline{S}(P) = \{\underline{S}_1, \dots, \underline{S}_n\}$ as the set of subsets $\underline{S}_i \subset \bar{K}$, $i = 1, \dots, n$, for which $\theta^\omega \geq \max_{1 \leq i \leq n} \min_{k \in \underline{S}_i} r_k^\omega, \omega \in \Omega$, is enforced. We use a pair of priority queues, $pq1$ and $pq2$, to store subproblems in the tree. The former stores those subproblems which are still eligible for customized branching, and the latter stores those which are to be solved by a commercial branch-and-bound solver. For each priority queue, a pop operation returns the subproblem in the queue with the smallest lower bound.

Algorithm 6: *GetCriticalSubset(P)*

Input: Problem P
Output: Subset of paths to branch on S
Let \bar{S} be the union of the elements of $\bar{S}(P)$
 $S_0 \leftarrow \emptyset$
for $t = 1, \dots, |\bar{K} \setminus \bar{S}|$ **do**
 Compute $S_t \in \operatorname{argmax}_{k \in \bar{K} \setminus (S_{t-1} \cup \bar{S})} \text{Value}(S_{t-1} \cup \{k\})$
 $\text{Value}_t \leftarrow \text{Value}(S_t)$
end for
Let $t^* \in \operatorname{argmax}_{1 \leq t \leq |\bar{K} \setminus \bar{S}|} t \cdot \text{Value}_t$
 $S \leftarrow S_{t^*}$
return S

Algorithm 7: *UpdateBoundsAndPush(pq, P, UB)*

Input: Priority queue pq , problem P , objective function upper bound UB
Output: Updated $\underline{\theta}(P, \omega)$, $\omega \in \Omega$, problem P pushed onto pq if lower bound for P less than UB
Let \bar{S} be the union of the elements of $\bar{S}(P)$
Let $\underline{S}_1, \dots, \underline{S}_N$ be the elements of $\underline{S}(P)$
 $LB \leftarrow 0$
for all $\omega \in \Omega$ **do**
 $\underline{\theta}(P, \omega) \leftarrow \max(\underline{\theta}^\omega(\bar{S}), \max_{1 \leq n \leq N} \min_{k \in \underline{S}_n} r_k^\omega)$
 $LB \leftarrow LB + \underline{\theta}(P, \omega)$
end for
if $LB < UB$ **then**
 Push P onto pq with priority LB
end if

Algorithm 8: *BranchAndBound*($p, r, c, b, UB, MAXD, MAXP$)

Input: Scenario probabilities $p^\omega > 0$, evasion probabilities $r_k^\omega \geq 0$, interdiction costs c_i, c_j , installation budget b , objective function upper bound UB , maximum depth of branch-and-bound tree $MAXD$, maximum size of branch-and-bound tree $MAXP$

Output: Optimal interdiction plan x^* , minimum evasion probability z^*

Create two empty priority queues of problems $pq1$ and $pq2$

Create problem P with $\bar{S}(P) = \underline{S}(P) = \emptyset$ and $\underline{\theta}(P, \omega) = \underline{\theta}^\omega(\emptyset)$, $\omega \in \Omega$

Push P onto $pq1$ with priority $\sum_{\omega \in \Omega} p^\omega \underline{\theta}(P, \omega)$

while $pq1$ not empty and $pq1.size + pq2.size < MAXP$ **do**

 Pop problem P off of $pq1$

if $|\bar{S}(P)| + |\underline{S}(P)| \geq MAXD$ **then**

 Push P onto $pq2$ with priority $\sum_{\omega \in \Omega} p^\omega \underline{\theta}(P, \omega)$

 continue

end if

$S \leftarrow GetCriticalSubset(P)$

 Create $\bar{P} \leftarrow P$ and $\underline{P} \leftarrow P$

 Add S_{t^*} to $\bar{S}(\bar{P})$ and to $\underline{S}(\underline{P})$

 UpdateBoundsAndPush($pq1, \bar{P}, UB$)

 UpdateBoundsAndPush($pq1, \underline{P}, UB$)

end while

while $pq1$ not empty **do**

 Pop problem P off of $pq1$

 Push problem P onto $pq2$ with priority $\sum_{\omega \in \Omega} p^\omega \underline{\theta}(P, \omega)$

end while

while $pq2$ not empty **do**

 Pop problem P off of $pq2$

 Let \bar{S} be the union of the elements of $\bar{S}(P)$

 Let $\hat{x}, \hat{\theta}$ be the optimal solution to (4.10) with $\underline{\theta}^\omega = \underline{\theta}(P, \omega)$ and with the added constraints $x_{i(k)} + x_{j(k)} \geq 1, k \in \bar{S}$

$\hat{z} \leftarrow \sum_{\omega \in \Omega} \sum_{k \in K} p^\omega \hat{\theta}^\omega$

if $\hat{z} < UB$ **then**

$UB \leftarrow \hat{z}$

$x^* \leftarrow \hat{x}$

end if

end while

return x^*, z^*

Bibliography

- [1] M. J. Atallah, editor. *Algorithms and theory of computation handbook*. CRC Press LLC, 1999.
- [2] M. L. Balinski. On a selection problem. *Management Science*, 17:230–231, 1970.
- [3] H. Bayrak and M. Bailey. Shortest path network interdiction with asymmetric information. *Networks*, 52:133–140, 2008.
- [4] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252, 1962.
- [5] V. Bier, S. Oliveros, and L. Samuelson. Choosing what to protect: strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory*, 9:563–587, 2007.
- [6] D. W. Blackett. Pure strategy solutions to Blotto games. *Naval Research Logistics Quarterly*, 5:107–109, 1958.
- [7] G. G. Brown, M. Carlyle, J. Salmerón, and R.K. Wood. Defending critical infrastructure. *Interfaces*, 36:530–544, 2006.
- [8] H. W. Corley and H. Chang. Finding the n most vital nodes in a flow network. *Management Science*, 21:362–64, 1974.

- [9] K. Cormican, D.P. Morton, and R.K. Wood. Stochastic network interdiction. *Operations Research*, 46:184–197, 1998.
- [10] ILOG CPLEX 10.0 User’s Manual, 2006.
- [11] K. Fan. Minimax theorems. *Proceedings of the National Academy of Sciences of the United States of America*, 39:42–47, 1953.
- [12] Y. Freund and R. E. Schapire. Game theory, on-line prediction and boosting. In *Proceedings of the Ninth Annual Conference on Computational Learning Theory*, pages 325–332, Desenzano del Garda, Italy, 1996.
- [13] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55:119–139, 1997.
- [14] D. R. Fulkerson and G. C. Harding. Maximizing the minimum source-sink path subject to a budget constraint. *Mathematical Programming*, 13:116–118, 1977.
- [15] G. Gallo, M. D. Grigoriadis, and R. E. Tarjan. A fast parametric maximum flow algorithm and applications. *SIAM Journal on Computing*, 18:30–55, 1989.
- [16] A. V. Goldberg and R. E. Tarjan. A new approach to the maximum flow problem. *Journal of the ACM*, 35:921–940, 1988.

- [17] D. S. Hochbaum. Dynamic evolution of economically preferred facilities. *European Journal of Operational Research*, 193:649–659, 2009.
- [18] E. Israeli and R. K. Wood. Shortest-path network interdiction. *Networks*, 40:97–111, 2002.
- [19] J. E. Kelley. The cutting-plane method for solving convex programs. *Journal of the Society for Industrial and Applied Mathematics*, 8:703–712, 1960.
- [20] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, 1951.
- [21] C. Lim and J. C. Smith. Algorithms for discrete and continuous multi-commodity flow network interdiction problems. *IIE Transactions*, 39:15–26, 2007.
- [22] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.
- [23] D. P. Michalopoulos. *Prioritization and optimization in stochastic network interdiction problems*. PhD thesis, The University of Texas at Austin, 2008.
- [24] A. J. Miller and L. A. Wolsey. Tight formulations for some simple mixed integer programs and convex objective integer programs. *Mathematical Programming*, 98:73–88, 2003.

- [25] D. P. Morton and F. Pan. Minimizing a stochastic maximum-reliability path. *Networks*, 52:111–119, 2008.
- [26] D. P. Morton, F. Pan, and K.J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions*, 39:3–14, 2007.
- [27] D. P. Morton and R. K. Wood. Restricted-recourse bounds for stochastic linear programming. *Operations Research*, 47:943–956, 1999.
- [28] J. Nash. Non-cooperative games. *The Annals of Mathematics*, 54:286–295, 1951.
- [29] F. Pan. *Models and Methods for Stochastic Network Interdiction*. PhD thesis, The University of Texas at Austin, 2005.
- [30] R. Powell. Defending against terrorist attacks with limited resources. *American Political Science Review*, 101:527–541, 2007.
- [31] T. K. Ralphs, M. J. Saltzman, and M. M. Wiecek. An improved algorithm for solving biobjective integer programs. *Annals of Operations Research*, 147:43–70, 2006.
- [32] J. M. W. Rhys. A selection problem of shared fixed costs and network flows. *Management Science*, 17:200–207, 1970.
- [33] J. C. Smith, C. Lim, and F. Sudargho. Survivable network design under optimal and heuristic interdiction scenarios. *Journal of Global Optimization*, 38:181–199, 2007.

- [34] D. Topkis. *Supermodularity and complementarity*. Princeton University Press, Princeton, N.J., 1998.
- [35] J. W. Tukey. Systems defense games. *Naval Research Logistics Quarterly*, 28:281–87, 1981.
- [36] R. M. Van Slyke and R. J.-B. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17:638–663, 1969.
- [37] J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, third edition, 1953.
- [38] A. Washburn. *Two-Person Zero-Sum Games*. Institute for Operations Research and the Management Sciences, third edition, 2003.
- [39] A. R. Washburn and R. K. Wood. Two-person zero-sum games for network interdiction. *Operations Research*, 43:243–251, 1994.
- [40] K. M. Witt. *Development of a probabilistic network model to simulate the smuggling of nuclear materials*. Nuclear and Radiation Engineering, The University of Texas at Austin, 2003. M.S. Thesis.
- [41] R. D. Wollmer. Removing arcs from a network. *Journal of the Operations Research Society of America*, 12:934–940, 1964.

- [42] R. D. Wollmer. Algorithms for targeting strikes in a lines-of-communication network. *Operations Research*, 18:497–515, 1970.
- [43] R. K. Wood. Deterministic network interdiction. *Mathematical and Computer Modeling*, 17:1–18, 1993.
- [44] J. Zhuang and V. Bier. Balancing terrorism and natural disasters - defensive strategy with endogenous attacker effort. *Operations Research*, 55:976–991, 2007.

Vita

Michael Victor Nehme, the son of David A. Nehme and Linda K. Nehme, was born in Stuart, Fl. on July 27, 1980. In 1998, he graduated from Martin County High School and began his studies at Duke University in Durham, NC. There, he obtained a Bachelor of Science in Electrical Engineering and Computer Science with a minor in Mathematics in May 2002 and a Masters degree in Engineering Management in May 2003. In August 2004, he began to pursue a Ph.D. in the Operations Research and Industrial Engineering graduate program at the University of Texas at Austin. During his time at UT, he worked as a Teaching Assistant and a Graduate Research Assistant and studied parallel computing at the San Diego Supercomputer Center during the summer of 2008.

Permanent address: 2412 W 12th Street
Austin, Texas 78703

This dissertation was typeset with \LaTeX^\dagger by the author.

[†] \LaTeX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's \TeX Program.