European Journal of Operational Research 252 (2016) 888–899

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Multithreat multisite protection: A security case study

David Ríos Insua[a], Javier Cano[b,*], Michael Pellot[c], Ricardo Ortega[c]

[a] Instituto de Ciencias Matemáticas ICMAT-CSIC, Spain
[b] Department of Computer Science and Statistics, Rey Juan Carlos University, Spain
[c] TMB, Spain

ARTICLE INFO

Article history:
Received 28 April 2015
Accepted 20 January 2016
Available online 30 January 2016

Keywords:
Risk analysis
Multiple intelligent attackers
Adversarial risk analysis
Multisite protection
Transportation

ABSTRACT

We provide a novel adversarial risk analysis approach to security resource allocation decision processes for an organization which faces multiple threats over multiple sites. We deploy a Sequential Defend-Attack model for each type of threat and site, under the assumption that different attackers are uncoordinated, although cascading effects are contemplated. The models are related by resource constraints and results are aggregated over the sites for each participant and, for the Defender, by value aggregation across threats. We illustrate the model with a case study in which we support a railway operator in allocating resources to protect from two threats: fare evasion and pickpocketing. Results suggest considerable expected savings due to the proposed investments.

1. Introduction

Crime and terrorism constitute major global issues. As an example, among the threats considered in the World Economic Forum (2015) Global Risks report, there are several related with security, including large-scale terrorist attacks or a major escalation in organized crime. Similarly, we may find security among the seven thematic H2020 priorities for European research (ec.europa.eu/programmes/horizon2020). Governments and organizations worldwide are indeed increasingly committed to protecting themselves against various security threats. Recent large-scale terrorist events like 9/11 or the Madrid train bombings have led to significant national investments in protective responses, see (Haberfeld & von Hassell, 2009). However, public opinion has not always seen such expenditures as prudent or effective, see (Parnell et al., 2008) or (Sunstein, 2007).

In turn, this has motivated great interest in modeling issues in relation with security, with varied tools from areas such as reliability analysis, data mining, game theory or complex dynamic systems. Recent accounts of various techniques and applications in the field of counterterrorism may be seen in e.g. Ezell, Bennett, von Winterfeldt, Sokolowski, and Collins (2010) or Wein (2009). Parnell et al. (2008) and Enders and Sandler (2011) provide overviews on strategies, models, and research issues in security risk analysis. Other relevant work include e.g. Zhuang and Bier (2007), who discuss resource allocation for countering terrorism and natural disasters; (Brown, Carlyle, Salmerón, & Wood, 2006), where the protection of critical infrastructures is addressed; or (Yang, Kiekintveld, Ordóñez, Tambe, & John, 2013), who present mathematical models of adversarial behavior to support security forces in their fight against different adversaries.

We consider problems in which an organization needs to protect multiple sites from multiple threats. Our case study refers to deciding the security resource allocation for a railway system whose operator faces threats from fare evaders and pickpockets. The figures presented in the paper have been modified from actual figures to protect the confidentiality of the case study provider. Therefore, *the data is realistic data but not actual data*. We assume that the relevant multiple threats are uncoordinated, in that different attackers do not make common cause, although the outcome of different types of attacks might affect each other. In our case study, fare evaders and pickpockets will not be coordinated, although pickpockets alone and a group of fare evaders will be organized. Hausken and Levitin (2012) provide a classification of systems defense and attack models. Within such classification, we shall be facing a case of protection from attacks over multiple elements with incomplete information. For earlier work on protecting from multiple attackers, see (Hausken & Bier, 2011) and references therein. Haphuriwat and Bier (2011); Hausken (2014b) and Levitin, Hausken, and Dai (2014) provide ideas in relation with multiple site protection. Bier, Oliveros, and Samuelson (2007) and Hausken (2014a) refer to uncertainty in attacker resources and asset valuation. All of them perform game theoretical analyses under convenient common knowledge assumptions.

[*] Corresponding author. Tel.: +34 914 888 411.
E-mail addresses: david.rios@icmat.es (D. Ríos Insua), javier.cano@urjc.es (J. Cano), mpellot@tmb.cat (M. Pellot), rortegap@tmb.cat (R. Ortega).

http://dx.doi.org/10.1016/j.ejor.2016.01.041
0377-2217/© 2016 Elsevier B.V. All rights reserved.

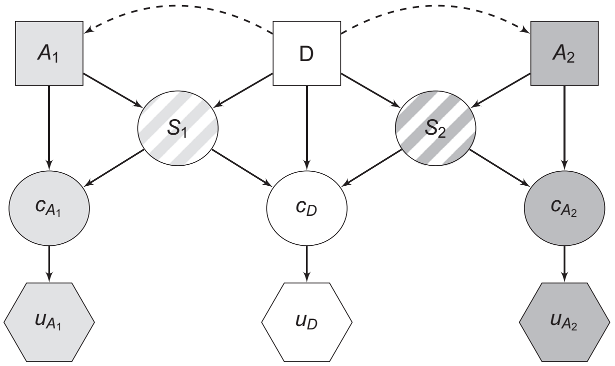

Fig. 1. Multiagent influence diagram for a bi-threat problem.

In contrast, we provide an adversarial risk analysis (ARA) approach, see (Ríos Insua, Ríos, & Banks, 2009), for such problems, combining multiple sites, multiple attackers and taking into account all relevant uncertainty sources. ARA builds a decision analysis model for one of the agents (she, the Defender), who forecasts the actions of her adversaries. Then, she will be able to decide her optimal defensive actions. Our approach will be based on the Sequential Defend-Attack model, see e.g. Brown et al. (2006) or Ríos and Ríos Insua (2012). In it, the Defender first chooses a defense and, then, after observing it, the Attacker decides his attack. We deploy one of such models for each type of threat and site, which we relate through resource constraints and aggregation of results over various sites for each participant and, for the case of the Defender, also by value aggregation over the various threats. We assume no particular spatial structure relating the sites, e.g. through proximity or a neighboring structure, see (Gil, Ríos, & Ríos Insua, 2016).

In Section 2, we provide a general framework for the basic problem of protecting a single site from multiple threats, illustrating it with our case in Section 3. Section 4 extends the previous model to the protection of multiple sites, applying it to an expanded version of the case study in Section 5. As described below, all the involved parameters have been assessed with the aid of transportation experts, using expert judgment elicitation techniques, see (O'Hagan et al., 2006) or Farquhar (1984), then validated at a security transportation workshop, and finally checked for robustness through sensitivity analysis.

2. Single site multithreat protection

We start with the basic multithreat protection problem over a single site. We consider a Defender, D, who needs to deploy defensive resources $d \in \mathcal{D}$ to protect the site from m uncoordinated attackers $A_1, \ldots, A_m$. These observe her decision d and, then, respectively, launch attacks $a_i \in \mathcal{A}_i$, $i = 1, \ldots, m$. The interaction between D and A_i through their corresponding decisions d and a_i, leads to a random result $S_i \in \mathcal{S}_i$. The Defender faces multiattribute consequences c_D which depend on her defense effort d and the results $s_1, \ldots, s_m$. She then gets her utility u_D. Each attacker will get his multiattribute consequences c_{A_i}, which depend on his attack effort a_i and his result s_i, and then gets his utility u_{A_i}.

The problem is illustrated in the multiagent influence diagram in Fig. 1, see (Koller & Milch, 2003). For simplicity, we only display two attackers, that is $m = 2$. White nodes correspond to the Defender, solid (light and dark) gray nodes to attackers A_1 and A_2, respectively. Striped nodes refer to interactions between the Defender and the attackers. Dashed arrows between node D and nodes A_1 and A_2 indicate that the attackers decide their alternatives after having observed the decision by D.

As an example, a port authority (D) is trying to protect a port against actions from drug smugglers (A_1) and terrorists (A_2) ready to introduce nuclear weapons. The Defender decisions d are portfolios which could include sniffer dogs, metal detectors, inspectors and others. Drug smuggler decisions a_1 typically would refer to drug smuggling (timing, placing, quantities) strategies. Terrorist decisions a_2 could refer to weapon smuggling strategies, like whether or not to infiltrate a nuclear weapon. S_1 could refer to the amount of drugs actually smuggled and S_2 could refer to the number of weapons smuggled.

The Defender aims at finding her optimal defense strategy d^*. She evaluates her consequences through her utility $u_D(d, s_1, \ldots, s_m)$. Assuming conditional independence between the outcomes S_i of different attacks, given the defensive resources d and attacks a_i, she needs to assess the probability models $p_D(s_i | d, a_i)$, $i = 1, \ldots, m$, reflecting which outcomes she finds more likely when attacker A_i launches attack a_i and she has deployed defensive resources d. She gets her expected utility, given the attacks, integrating out the uncertainty over the outcomes of the attacks:

$$\psi_D(d | a_1, \ldots, a_m) = \int \cdots \int u_D(d, s_1, \ldots, s_m) \times p_D(s_1 | d, a_1) \cdots p_D(s_m | d, a_m) \, \mathrm{d}s_1 \ldots \mathrm{d}s_m. \quad (1)$$

Suppose that the Defender is able to build the models $p_D(a_i | d)$, $i = 1, \ldots, m$, expressing her beliefs about which attack a_i will be chosen by the ith attacker after having observed d. Since attacks are uncoordinated, we assume conditional independence of $a_1, \ldots, a_m$ given d. Then, D may compute

$$\psi_D(d) = \int \cdots \int \psi_D(d | a_1, \ldots, a_m) \, p_D(a_1 | d) \cdots p_D(a_m | d) \times \mathrm{d}a_1 \ldots \mathrm{d}a_m,$$

and solve $\max_d \psi_D(d)$ to find her optimal defense resource allocation d^*.

The only nonstandard assessments in this formulation are those of $p_D(a_i | d)$. To obtain them, the Defender may put herself into the shoes of each attacker, and solve their corresponding problem separately, since they are uncoordinated. For instance, for the problem faced by attacker A_1, assuming that he is an expected utility maximizer, see (French & Ríos Insua, 2000), the Defender would need his utility $u_{A_1}(a_1, s_1)$ and probabilities $p_{A_1}(s_1 | d, a_1)$. Then, she would solve

$$a_1^*(d) = \operatorname*{argmax}_{a_1 \in \mathcal{A}_1} \int u_{A_1}(a_1, s_1) \, p_{A_1}(s_1 | d, a_1) \, \mathrm{d}s_1, \quad (2)$$

to find his optimal attack given that she has implemented d. However, she lacks knowledge about u_{A_1} and p_{A_1}. Suppose we may model her uncertainty about them through random utilities and probabilities $\left(U_{A_1}, P_{A_1}\right)$, and propagate that uncertainty to obtain the random optimal attack, given her defense d,

$$A_1^*(d) = \operatorname*{argmax}_{a_1 \in \mathcal{A}_1} \int U_{A_1}(a_1, s_1) \, P_{A_1}(s_1 | d, a_1) \, \mathrm{d}s_1. \quad (3)$$

Then, we would get $p_D(a_1 | d) = \Pr(A_1^*(d) \leq a_1)$, which may be approximated by Monte Carlo through Algorithm 1.

A similar scheme could be implemented in parallel for the other attackers, $A_2, \ldots, A_m$, leading to estimates $\widehat{p_D}(a_i | d)$ of the required probabilities $p_D(a_i | d), i = 2, \ldots, m$.

The approach may be generalized in several ways. For example, the simultaneous, but uncoordinated, implementation of attacks $a_1, \ldots, a_m$ could be jointly detrimental in face of defensive resources d, which could be shared against various types of attacks, see Fig. 2a. Then, we could rewrite the probability model in (1) as

$$p_D(s_1 | d, a_1, \ldots, a_m) \cdots p_D(s_m | d, a_1, \ldots, a_m),$$

and proceed in a similar fashion.

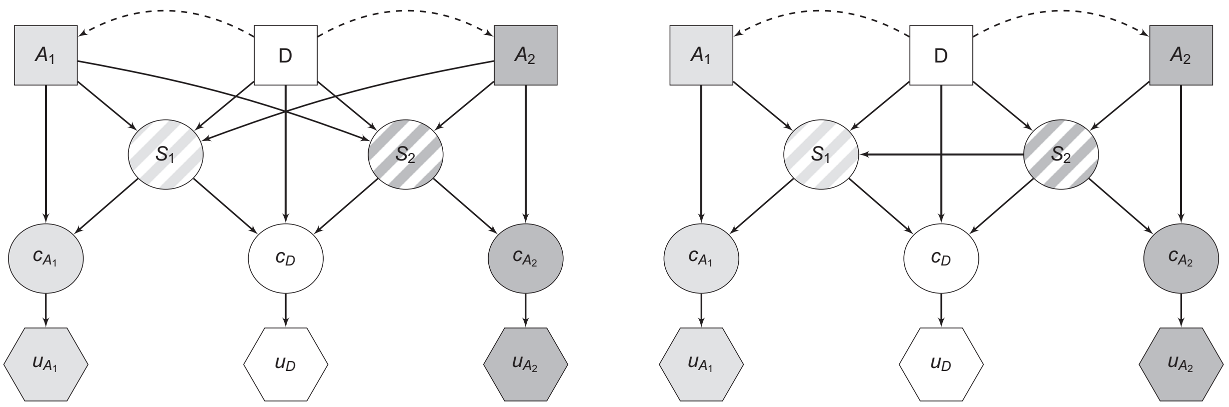

(a) Detrimental effect.

(b) Cascading effect.

Fig. 2. Some generalizations for the bi-threat problem.

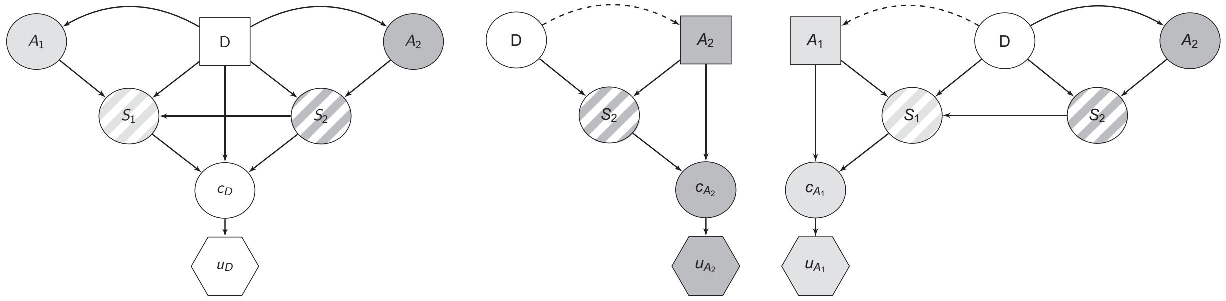

(a) Defender's problem.

(b) Attacker A_2.

(c) Attacker A_1.

Fig. 3. Solving the bi-threat problem with cascading effect.

Algorithm 1: Simulating the problem for attacker $\mathtt{A}_1$.

```
for $d \in \mathcal{D}$
    for $k = 1$ to $K$
        for $a_1 \in \mathcal{A}_1$
            Draw $\left(u^k_{A_1}, p^k_{A_1}\right) \sim \left(U_{A_1}, P_{A_1}\right)$
            Compute $\psi^k_{A_1}(d, a_1) = \int u^k_{A_1}(a_1, s_1)\, p^k_{A_1}(s_1|d, a_1)\, ds_1$
        Compute $a^k_1(d) = \arg\max_{a_1 \in \mathcal{A}_1} \psi^k_{A_1}(d, a_1)$
    Approximate $\widehat{p_D}(a_1|d) \approx \#\{1 \le k \le K : a^k_1(d) \le a_1\}/K$
```

Alternatively, there could be a cascading effect between the attack results. For example, see Fig. 2b, assuming that $m = 2$, it could be that s_2 affects s_1, so that $p_D(s_1|d,\ a_1)p_D(s_2|d,\ a_2)$ in (1) becomes $p_D(s_1|d,\ a_1,\ s_2)p_D(s_2|d,\ a_2)$. Under this assumption, the general scheme required to estimate $\widehat{p_D}(a_i|d),\ i = 1, \ldots, m$ cannot be implemented in parallel, but requires some sequentiality, as shown below and illustrated in our case study.

The influence diagram for the Defender's problem in this latter case is shown in Fig. 3a, where the attacks appear now as chance nodes.

Now, the expected utility for the Defender is

$$\psi_D(d|a_1, a_2) = \iint u_D(d, s_1, s_2)\, p_D(s_1|d, a_1, s_2) \\ \times p_D(s_2|d, a_2)\, ds_1\, ds_2.$$

After integrating out the uncertainty over the attacks, we obtain her expected utility

$$\psi_D(d) = \iint \psi_D(d|a_1, a_2)\, p_D(a_1|d)\, p_D(a_2|d)\, da_1\, da_2,$$

and she obtains the optimal portfolio d^* through $\max_d \psi_D(d)$.

We need to assess $p_D(a_1|d)$ and $p_D(a_2|d)$. We start with $p_D(a_2|d)$. The influence diagram for A_2 is sketched in Fig. 3b. Since the result of the attack performed by A_2 influences—but is not influenced by—that of A_1, the estimation of $p_D(a_2|d)$ is similar to that of $p_D(a_1|d)$ in the initial basic formulation, as outlined in (2), (3) and Algorithm 1.

Once we have dealt with the problem referring to A_2, we may address the problem for A_1, whose influence diagram is shown in Fig. 3c. We need to consider the influence of chance node $\textcircled{A_2}$ (for him). Indeed, he needs to maximize his expected utility incorporating his uncertainty about the attacking decision a_2 (given the defense d)

$$a^*_1(d) = \operatorname*{argmax}_{a_1 \in \mathcal{A}_1} \iiint u_{A_1}(a_1, s_1)\, p_{A_1}(s_1|d, a_1, s_2)\, p_{A_1}(s_2|d, a_2) \\ \times p_{A_1}(a_2|d)\, ds_1\, ds_2\, da_2.$$

In general, the Defender will not know u_{A_1}, $p_{A_1}(s_1|d, a_1, s_2)$, $p_{A_1}(s_2|d, a_2)$ and $p_{A_1}(a_2|d)$, but she may acknowledge her uncertainty about them through random utilities and probabilities $\left(U_{A_1}, P_{A_1}(s_1|\cdot), P_{A_1}(s_2|\cdot), P_{A_1}(a_2|\cdot)\right)$, and propagate that uncertainty to obtain the random optimal attack for each d

$$A^*_1(d) = \operatorname*{argmax}_{a_1 \in \mathcal{A}_1} \iiint U_{A_1}(a_1, s_1)\, P_{A_1}(s_1|d, a_1, s_2)\, P_{A_1}(s_2|d, a_2) \\ \times P_{A_1}(a_2|d)\, ds_1\, ds_2\, da_2.$$

Once with $A^*_1(d)$, she can obtain $p_D(a_1|d) = \Pr(A^*_1(d) \le a_1)$, which can be estimated following a sampling scheme similar to that in Algorithm 1.

For this, we need to assess $\left(U_{A_1}, P_{A_1}(s_1|d, a_1, s_2), P_{A_1}(s_2|d, a_2), P_{A_1}(a_2|d)\right)$, and $\left(U_{A_2}, P_{A_2}(s_2|d, a_2)\right)$. With respect to the random probabilities, we could base them on the corresponding assess-

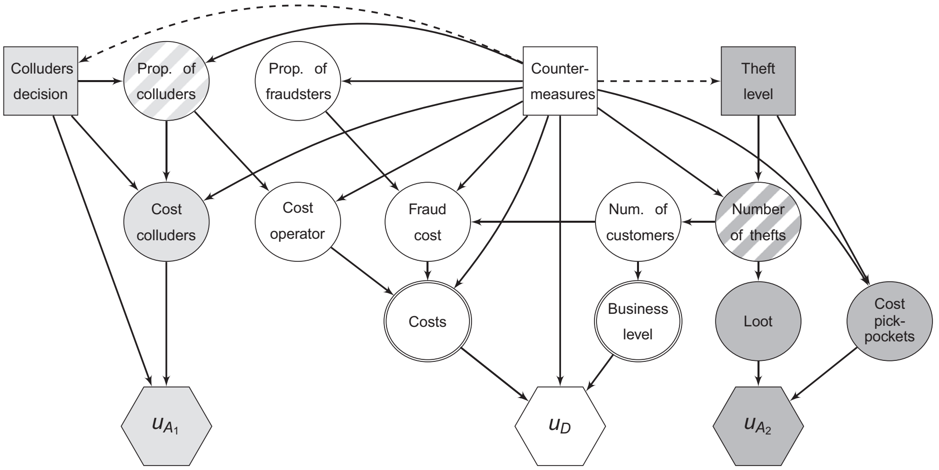

Fig. 4. Influence diagram when evaders and pickpockets are present.

ments for the Defender, $p_D(s_1|d,\, a_1,\, s_2)$, $p_D(s_2|d,\, a_2)$ and $p_D(a_2|d)$, possibly as we illustrate for $P_{A_2}(s_2|d, a_2)$:

- Should S_2 be discrete, $P_{A_2}(s_2|d, a_2)$ could be modeled as a Dirichlet distribution with mean $p_D(s_2|d,\, a_2)$ and variance accounting for the incumbent uncertainty. In particular, when S_2 is binary, $P_{A_2}(s_2|d, a_2)$ would be modeled as a beta distribution, see (French & Ríos Insua, 2000).
- Should S_2 be continuous, then $P_{A_2}(s_2|d, a_2)$ could be a Dirichlet process with base distribution $p_D(s_2|d,\, a_2)$ and concentration parameter δ expressing our uncertainty about such base, see (Ferguson, 1973).

For both cases, when lacking information, we could use a sufficiently large value for the variance or concentration parameter, respectively. Note also that, for some probability distributions $p_D(\cdot)$, we might have available a sample through Monte Carlo simulation, as e.g. for $p_D(a_2|d)$, which would have been obtained after applying Algorithm 1 to the problem for attacker A_2. Should this happen, we could use the sample variance associated with $p_D\,(\cdot)$ to describe our uncertainty about the corresponding $P_{A_2}(\cdot)$, adjusting the variance (or concentration parameter) accordingly.

As for the random utility, typically we shall have information about the interests and objectives of the attackers, see (Keeney, 2007) or (Keeney & von Winterfeldt, 2010; 2011), which we would aggregate with a weighted measurable value function, as by Dyer and Sarin (1979). Based on the relative risk aversion concept, see (Dyer & Sarin, 1982), we could assume risk proneness on the attackers. Finally, the uncertainty would be reflected by distributions over the weights and risk proneness coefficients. The proposed models for random probabilities and utilities are illustrated within our case study.

Wang and Bier (2013) provide another approach for assessing adversary preferences using ordinal judgments and the probabilistic inversion method, see (Kraan & Bedford, 2005).

3. Protecting a railway station from fare evasion and pickpocketing

We consider the case of a railway operator (D) which needs to protect a single station from two threats: fare evasion and pickpocketing. The operator has to deal with two types of fare evaders: (1) Standard, regarded as unintentional, who do not pay for the service in a casual manner; and (2) Colluders (A_1), who are intentional fare evaders preparing their evasion actions in an organized way. Reddy, Kuhls, and Lu (2011) provide related work addressing fare evasion, studied in detail by Ríos Insua, Cano, Pellot, and Ortega (2015b). We model pickpockets (A_2) as a single organized group, see (Troelsen & Barr, 2012) on combating pickpocketing in public transportation. For this threat, we focus on both security and image costs, as pickpocketing may decrease the feeling of security, affecting the operator business level.

When the operator faces both threats simultaneously, this can be viewed as a bi-threat Sequential Defend-Attack model with cascading effect, whose influence diagram, adapted from Fig. 2b, is shown in Fig. 4. Light gray nodes refer to fare evasion, dark gray nodes to pickpocketing and white nodes correspond to the operator's problem.

The decision node “Countermeasures” refers to the portfolio deployed by the operator to reduce: (1) The theft level; and (2) The proportion of standard fraudsters and colluders. With respect to pickpockets, we have uncertainty about the number of thefts and, consequently, on its impact over business level. Pickpockets face costs when preparing their actions, as well as the possibility of being fined if caught red-handed. When successful, they obtain a loot. For fare evasion, we have uncertainty about the proportion of (standard) fraudsters and the number of customers (influenced, in turn, by the theft level), from which we obtain the fraud cost. On the other hand, colluders would decide the proportion of fare evasion they would undertake, although the actual proportion, as reflected in node “Prop. of colluders”, would depend also on the means implemented by the operator. For clarity, we keep fraud costs due to standard evaders and colluders separately, but aggregate them in the deterministic node “Costs”. We assume that colluders and pickpockets do not coordinate their criminal activities.

The operator can deploy eight different types of countermeasures, displayed in Table 1. The first five fight fare evasion, whereas the last four are addressed towards pickpockets (the fifth one could affect both threats). We aim at supporting the operator in devising an optimal security portfolio. We provide a detailed model of the problem and, then, fully illustrate the required assessments and implementation.

3.1. Case modeling

Let $(d_1,\, d_2,\, d_3,\, d_5,\, d_6,\, d_7)$ be, respectively, the inspectors, door guards, secured automatic access doors, guards, patrols and cameras to be deployed. Their associated unit costs over the planning period will be, respectively, q_1, q_2, q_3, q_5, q_6 and q_7. We also use a binary decision variable $d_4 \in \{0,\, 1\}$, with $d_4 = 1$ indicating the involvement of ticket clerks in observation tasks, and $d_4 = 0$, otherwise. As clerks are already hired by the operator, there are no additional direct costs associated with the reassignment of their

Table 1
Relevant features of countermeasures.

	Role		Features
	Fare evasion	Pickpocketing	
Inspectors	Preventive/recovery	–	Inspect customers. Collect fines
Door guards	Preventive	–	Control access points
Doors	Preventive	–	New secured automatic access doors
Ticket clerks	Preventive	–	Currently, no implication in security
Guards	Preventive	Preventive/recovery	Patrol along the facility
Patrols	–	Preventive/recovery	Trained guard+security dog
Cameras	–	Preventive	Complicate pickpocket actions
Awareness campaign	–	Preventive	Alert users about pickpockets

duties. However, making them switch from a passive to a proactive attitude towards the fare evasion problem could entail negative implications in terms of troubles with unions. We monetize this assuming a cost q_4 for that. We also assume that a public awareness plan costs q_8. A binary decision variable $d_8 \in \{0, 1\}$ models its implementation, with $d_8 = 1$ meaning that the operator pays for the awareness plan, and $d_8 = 0$, otherwise. The budget available for the relevant planning period will be B. Then, the feasible security portfolios $d = (d_1, d_2, d_3, d_4, d_5, d_6, d_7, d_8)$ will satisfy

$$\begin{aligned} & q_1 d_1 + q_2 d_2 + q_3 d_3 + q_5 d_5 + q_6 d_6 + q_7 d_7 + q_8 d_8 \le B, \\ & d_1, d_2, d_3, d_5, d_6, d_7 \ge 0, \\ & d_1, d_2, d_3, d_5, d_6, d_7 \text{ integer}, \\ & d_3 \le \bar{d}_3, \\ & d_4, d_8 \in \{0, 1\}, \end{aligned} \tag{4}$$

where $\bar{d}_3$ is the maximum number of access doors that may be replaced in the station.

3.1.1. Pickpocketing

Pickpocket gangs are organized groups, taking advantage of crowded situations. Typically, a pickpocket perpetrates the theft, while his colleagues cover him and try to run away with the loot. The event flow for a pickpocket attempt is: (a) Some pickpockets will succeed in committing their theft; (b) Out of them, some will not be caught, getting the loot; (c) Otherwise, they will be caught red-handed, losing the loot and being fined. Should they try to evade the fine, they could be imprisoned, but this rarely happens, since it is better for pickpockets to pay off the fine immediately and return to “business as usual”. The gang will attempt to commit t thefts over the relevant planning period. t will be their decision variable.

Operator's dynamics. In the pickpocketing problem, the operator: (a) Invests $d_p = (d_5, d_6, d_7, d_8)$; (b) Faces the corresponding delinquency level; (c) Possibly observes a reduction in business; (d) Obtains her utility, which will depend on the change in business level and its operational costs.

The security investment costs for the operator against pickpocketing are

$$c_{\text{inv}}^{(2)}(d_5, d_6, d_7, d_8) = q_5 d_5 + q_6 d_6 + q_7 d_7 + q_8 d_8. \tag{5}$$

As for the business level b, the operator considers that it will not change much, unless there is a large number of thefts. We use an average logistic response to model this, see Fig. 5:

$$E[b|t] = \frac{b_0 - b_r}{1 + \exp[\gamma_b (t - t_{0.5})]} + b_r, \; t > 0. \tag{6}$$

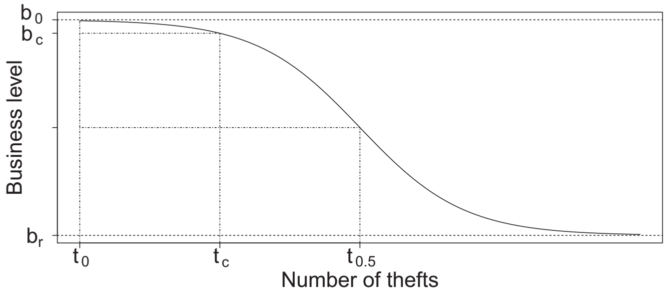

Fig. 5. Reduction in business level due to pickpockets.

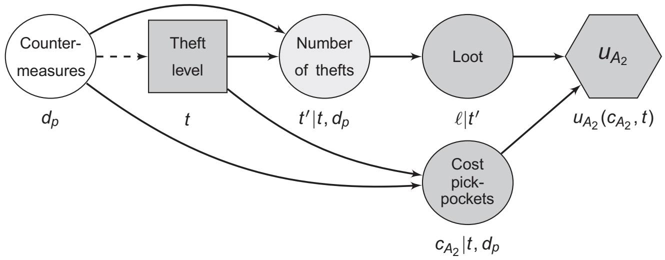

Fig. 6. Influence diagram for the Attacker's problem.

Here, b_0 is the ideal business level, free of the pickpocketing threat (for $t_0 \ll t_{0.5}$, $E[b|t_0]$ does not differ much from b_0); b_c is the business level given the current theft level t_c; and b_r is the business level for a large number of thefts: as pickpocketing surpasses a threshold, the business level will tend to attenuate its reduction stabilizing around b_r. $t_{0.5}$ designates the number of thefts for which the total expected business level reduction $b_0 - b_r$ would be halved. γ_b reflects how business level would drop as the number of thefts increases. We assume a normal distribution for $p_D(b|t)$, centered at $E[b|t]$, with standard deviation σ_b accounting for the uncertainty over $b|t$. Depending on the scale, we might need to truncate $p_D(b|t)$ at 0. The total balance for the operator will be

$$c_D^{(2)}(d_5, d_6, d_7, d_8, b) = -c_{\text{inv}}^{(2)}(d_5, d_6, d_7, d_8) - (b_0 - b).$$

Pickpockets' dynamics. The problem faced by pickpockets is depicted in Fig. 6.

Their event flow involves the following steps, where we provide, when pertinent, the operator's assessments $p_D(\cdot)$ about the involved quantities together with her assessments $P_{A_2}(\cdot)$ about the corresponding pickpocket's random distributions, following the suggestion at the end of Section 2, concerning Dirichlet distributions and processes.

(a) They see the operator's security investments $d_p = (d_5, d_6, d_7, d_8)$.
(b) They decide the number $t \in \mathcal{A}_2$ of theft attempts they will undertake, being $\mathcal{A}_2$ the set of plausible values.
(c) They implement the actual number of theft operations, $t' = (1 - \tau)t$: due to the measures deployed by the operator, some operations may need to be aborted. We

use a model $t' = (1 - \tau(d_5, d_6, d_7))t$, where $\tau(d_5, d_6, d_7)$ is the proportion of aborted thefts with respect to the original plan, which depends on d_5, d_6, d_7. A typical assumption for τ would be a beta distribution $p_D(\tau(d_5, d_6, d_7)) \sim \mathcal{B}e(\alpha_\tau(d_5, d_6, d_7), \beta_\tau(d_5, d_6, d_7))$, with τ close to zero if we feel that pickpockets are very committed to their plan, thus with $\alpha_\tau \ll \beta_\tau$. Then, we model $P_{A_2}(\tau(d_5, d_6, d_7))$ as a Dirichlet process with base distribution $p_D(\tau(d_5, d_6, d_7))$ and concentration parameter δ_τ. The random distribution $P_{A_2}(\tau(d_5, d_6, d_7)) = P_{A_2}(\tau | d_5, d_6, d_7)$ induces $P_{A_2}(t' | t, d_5, d_6, d_7)$.

(d) The implementation costs of their actions are $q_p t$, where q_p is the preparation cost per operation.

(e) They face their operational costs per each effectively attempted operation:

1. With probability $(1 - \xi)$, they will not succeed in the theft attempt. The only consequences are the preparation costs. The proportion ξ of successful operations depends on the pickpockets' ability. However, the presence of patrols and/or guards, and the influence of awareness campaigns will reduce such value. We model the operator's beliefs about ξ through a beta distribution $p_D(\xi(d_5, d_6, d_8)) \sim \mathcal{B}e(\alpha_\xi(d_5, d_6, d_8), \beta_\xi(d_5, d_6, d_8))$ with standard deviation σ_ξ and mean
$$E[\xi | d_5, d_6, d_8] = \xi_0 \cdot \exp(-\mu_5 d_5 - \mu_6 d_6 - \mu_8 d_8) + \xi_r.$$
Here, μ_5, μ_6 and μ_8 account for the fact that each additional unit of (d_5, d_6, d_8) is expected to reduce the proportion of successful operations. $(\xi_0 + \xi_r)$ represents the current success proportion, if no additional countermeasures are deployed, and ξ_r the residual proportion that would persist, even if infinite resources (d_5, d_6, d_8) were deployed. Again, we use a Dirichlet process for $P_{A_2}(\xi(d_5, d_6, d_8)) \sim \mathcal{DP}\big(p_D(\xi(d_5, d_6, d_8)), \delta_\xi\big)$.
2. With probability $\xi\theta$, they succeed in their theft attempts but are detained afterwards, facing the possibility of being fined with average fine f_p. The detention proportion θ depends on the number of patrols and guards. However, such dependence will not be linear: the operator believes that the additional contribution of each new patrol or guard will be smaller. To acknowledge the operator's uncertainty about θ, we use a beta distribution $p_D(\theta(d_5, d_6)) \sim \mathcal{B}e(\alpha_\theta(d_5, d_6), \beta_\theta(d_5, d_6))$ with standard deviation σ_θ and mean
$$E[\theta | d_5, d_6] = 1 - \exp(-\rho_5 d_5 - \rho_6 d_6),$$
where (ρ_5, ρ_6) account for the fact that each additional unit (d_5, d_6) is expected to increase the detention proportion. Then, $P_{A_2}(\theta(d_5, d_6)) \sim \mathcal{DP}\big(p_D(\theta(d_5, d_6)), \delta_\theta\big)$.
3. With probability $\xi(1 - \theta)$, they succeed and avoid getting caught. They get the theft benefit ℓ, which is uniformly distributed $\ell \sim \mathcal{U}(\ell_a, \ell_b)$.

(f) The pickpockets face an expected cost/benefit balance
$$c_{A_2}(t_1, t_2, t_3) = -q_p t - f_p t_2 + \ell t_3,$$
where (t_1, t_2, t_3) come from a multinomial distribution
$$\mathcal{M}\big(t'; 1 - (1 - \tau)\xi, (1 - \tau)\xi\theta, (1 - \tau)\xi(1 - \theta)\big).$$
When necessary, we shall use
$$p_{t_1 t_2 t_3 d_p} = \Pr(t_i \text{ theft attempts with outcome } i,$$
$$i = 1, 2, 3 | d_p \text{ is invested}),$$
where $\text{outcome} = \{1, 2, 3\}$ corresponds to the pickpocket cases described in (e). The random distributions $P_{A_2}(\xi(d_5, d_6, d_8)) = P_{A_2}(\xi | d_5, d_6, d_8)$ and $P_{A_2}(\theta(d_5, d_6)) = P_{A_2}(\theta | d_5, d_6)$ induce the random distribution $P_{A_2}(c_{A_2} | t, d_p)$.

(g) The pickpockets get their utility, which depends on the loot and costs entailed in implementing their decision. We assume that pickpockets are constant risk prone in benefits, see (Dyer & Sarin, 1982). Therefore, their utility function is strategically equivalent to
$$u_{A_2}(c_{A_2}) = \exp(k_{A_2} \cdot c_{A_2}),\ k_{A_2} > 0.$$
A random utility model for the pickpockets could be
$$U_{A_2}(c_{A_2}) = \exp(k_{A_2} \cdot c_{A_2}),\ k_{A_2} \sim \mathcal{U}(0, K_{A_2}).$$

As a consequence, the pickpockets get their random expected utility
$$\Psi_{A_2}(t', t, d_p) = \iint \Bigg[\sum_{t_1, t_2, t_3} p_{t_1 t_2 t_3 d_p} U_{A_2}\Big(-q_p t - f_p t_2 + \ell t_3\Big) \Bigg]$$
$$\times P_{A_2}(\xi | d_5, d_6, d_8) P_{A_2}(\theta | d_5, d_6)\, \mathrm{d}\xi\, \mathrm{d}\theta.$$
We integrate out the uncertainty over t' to get the random expected utility
$$\Psi_{A_2}(t, d_p) = \int \Psi_{A_2}(t', t, d_p) P_{A_2}(t' | t, d_p)\, \mathrm{d}t'.$$
Then, we find the pickpocket's random optimal theft level, given the defense d_p, through
$$T^*(d_p) = \operatorname*{argmax}_{t \in \mathcal{A}_2} \Psi_{A_2}(t, d_p).$$
We would simulate as in Algorithm 1 to obtain an estimate of the pickpocket's random optimal decision T given the security investment d_p, so that $p_D(T \le t | d_p) = \Pr(T^*(d_p) \le t)$.

3.1.2. Fare evasion

The fare evasion problem is described in detail by Ríos Insua et al. (2015b). We provide a brief summary of the uncertainty models involved, emphasizing the issue that the number of customers depends on the theft level, reflecting the cascading effect mentioned above.

In relation with standard fare evasion, we distinguish three types of customers: (1) Civic, who pay the ticket; (2) Standard evaders who do not pay the ticket but are not caught, therefore producing a loss of ν_c, the cost of the ticket; and (3) Standard evaders who are caught without a ticket, producing an expected income f_c due to fines. Denote by N_1, N_2 and N_3 the number of customers of each type, with $N = N_1 + N_2 + N_3$ the total number of customers. Let $p_{N_i d_c}$ be the probability that there are N_i customers of type i, $i = 1, 2, 3$ when the security plan $d_c = (d_1, d_2, d_3, d_4, d_5)$ is implemented.

With respect to colluders, we view them as a "club" which entails M operations over the incumbent planning period. Their event flow is: (1) Some colluders eventually change their mind deciding to pay when using the facility; (2) The rest decide not to pay; (3) Some of these will be inspected and fined. This will partly mitigate the operator losses associated with fare evasion. Colluders benefit from evading the ticket fare, but they face the possibility of being fined, in addition to having some preparation costs. Denote by (M_1, M_2, M_3) the number of aborted, successful, and failed operations, respectively, which we assume come from a multinomial distribution. Let $p_{M_1 M_2 M_3 d_c}$ be the probabilities that there are M_i colluders of type i, $i = 1, 2, 3$ when d_c is invested.

The benefit/cost balance for the operator, due to fare evasion, is
$$c_D^{(1)}(N_1, N_2, N_3, M_1, M_2, M_3, d_c) = -\nu_c(N_2 + M_2) + f_c(N_3 + M_3)$$
$$- \sum_{k=1}^{5} q_k d_k. \qquad (7)$$

Table 2
Planned investments (thousands of dollars).

Measure	Inspectors	Door guards	Doors	Guards	Patrols	Cameras	Campaign
Max.	4	4	1	4	4	3	1
Unit costs	50	25	15	30	35	4.5	40

3.1.3. Solving the operator's bi-threat problem

In Section 3.1.1, we sketched how to solve the pickpocketing problem from the point of view of the attacker. We obtained an estimate of $p_D(t|d_p)$. By Ríos Insua et al. (2015b), we show how to obtain $p_D(r|d_c)$, which models the operator beliefs over the proportion r of fare evasion attempted by colluders, when the security investment was d_c. Note that we need to adapt the methodology described there to incorporate the inherent cascading effect, as the number of customers will be affected by the presence of pickpockets through the business level. We explain this in Section 3.1.4.

We aggregate now all benefits and costs for the operator, using (5) and (7). This includes the investment in countermeasures, the fraud/fine balance associated with fare evasion, and the reduction in business level due to pickpocketing:

$$c_D(N_1, N_2, N_3, M_1, M_2, M_3, d, b) = -v_c(N_2 + M_2) + f_c(N_3 + M_3)$$
$$-\sum_{k=1}^{8} q_k d_k - (b_0 - b).$$

The operator is constant risk averse to an increase in income, see (Dyer & Sarin, 1982). Therefore, her utility function u_D will be strategically equivalent to $u_D(c_D) = -\exp(-k_D \cdot c_D)$, with $k_D > 0$.

We may now evaluate a security portfolio d by computing its expected utility:

$$\psi_D(d) = \int \left\{ \iint \left[\sum_{\substack{N_1,N_2,N_3\\ M_1,M_2,M_3}} p_{M_1M_2M_3d_c} \cdot p_{N_1d_c}\, p_{N_2d_c}\, p_{N_3d_c} \cdot u_D(c_D) \right] \times p_D(t|d_p)\, p_D(b|t)\, \mathrm{d}t\, \mathrm{d}b \right\} \times p_D(r|d_c)\, \mathrm{d}r. \quad (8)$$

We would then maximize $\psi_D(d)$, subject to (4), to find the optimal security plan.

3.1.4. Model assessments

We illustrate the model considering a specific railway station, with features representative of many others in the operator network, with a single street level entrance, and a moderate daily flow of customers. We choose one year as the relevant planning period, since the security budget is planned annually. The available annual security budget for new countermeasures at this station is $100,000, to be shared against both threats.

Table 2 displays the maximum additional investments over the planning period that the operator contemplates for each countermeasure, as well as their associated unit costs.

For human resources, we have provided unit annual gross salaries. The operator would have to hire four full-time (35 hours/week) workers of each category to fully cover service, since the station operates approximately 140 hours weekly. We also include the overall cost of installing a secured automatic access door over a whole year, including maintenance and repair, taking into account the typical door lifetime and, similarly, for costs associated with cameras. There are 324 feasible portfolios. As for the redefinition of clerk duties, the operator estimates that the negotiation with unions over a whole year could cost, in terms of labor troubles, approximately, $15,000 (per station).

We discuss now the assessment of the relevant parameters for the pickpocketing threat. With respect to theft level, the operator estimates that, approximately, only one out of every nine thefts is actually reported, and acknowledges, on average, three reported thefts per day across the network. This is roughly equivalent to 10,000 thefts throughout a year. The network has over 100 stations, of which the incumbent station cannot be regarded as a hotspot. Then, the operator estimates that the current annual number of thefts at such station would be around one half of the average value per station. Thus, she chooses a binomial distribution for the current theft level $t_c \sim \mathcal{B}in(100, 0.5)$ (expected value 50, standard deviation 5).

We estimate the reduction in business level due to the presence of pickpockets through the number of annual sold tickets, which has been around 1,000,000 over the last five years. As there are different transportation passes, the effective fare ticket is estimated at v_p =$0.75. Then, we set b_c =$750,000. The operator believes that the ideal business level, without pickpockets, would be around b_0 =$760,000. In addition, she thinks that business level would never drop below 80% of its current value, i.e. b_r =$600,000, even for an excessively large number of thefts. She would expect one half of such reduction if the number of thefts doubles, i.e. $t_{0.5} = 100$, see Fig. 5. She does not think that deterioration in business will happen drastically and she assesses $\gamma_b = 0.08$. Therefore, it seems reasonable to use $t \in \{0, 1, \dots, 150\} = \mathcal{A}_2$ as possible values for the pickpockets' decision variable. Finally, the operator does not have great uncertainty about the expected value of $E[b|t]$, choosing $\sigma_b = \$8,000$.

We have estimated the costs and consequences for pickpockets based on expert judgment:

- Preparation costs are estimated at $7 per attempted operation. This accounts for the ticket fare plus some expenses for daily food, drink and clothes.
- The fine in case of being caught red-handed depends, to some extent, on the amount robbed. For simplicity, we assume an average fine f_p =$600 for the whole gang.
- According to the data collected by the operator from theft complaints, the loot obtained by the whole gang varies uniformly between $\ell_a = \$150$ and $\ell_b = \$375$. For this, we have used a flat improper prior distribution ($\pi(\ell'_b) = 1$ for $\ell'_b = \ell_b - \ell_a > 0$, see (Rossman, Short, & Parks, 1998).
- The estimation of the proportion ξ of successful operations is an involved issue, given that little data is available. We have assessed the current proportion $\xi_0 + \xi_r$ through a beta $\mathcal{B}e(3, 1)$ distribution. The operator aims at reducing this proportion to a target value of $\xi_r = 0.05$. We assessed (μ_5, μ_6, μ_8) through expert elicitation. As an illustration, assume that the d_5 guards were the only countermeasure available. Our experts considered that having one guard would reduce the success proportion from 0.75 to approximately 0.55. Consequently, we assessed $\mu_5 = 0.35$. We checked for robustness of the assessment, asking the experts about the expected reduction in success proportion if more than one guard were hired, obtaining consistent results. We repeated the same reasoning when varying the number of patrols, obtaining $\mu_6 = 0.55$. The estimation of μ_8 was accomplished using the only possible values, $d_8 = \{0, 1\}$, leading to $\mu_8 = 0.25$. The operator expressed little uncertainty about her assessment of the Attacker's distribution. Thus, we assessed $\delta_\xi = 0.05$.

Table 3
Expected utilities for representative portfolios for different fare evasion proportions.

$\phi_0 + \phi_r = 0.03$				$\phi_0 + \phi_r = 0.06$			$\phi_0 + \phi_r = 0.12$		
d	Invest.	$\psi(d)$	Income	d	Invest.	$\psi(d)$	d	Invest.	$\psi(d)$
(1, 0, 0, 0, 0, 1, 0, 0)	85000	−2.03	−171585	(2, 0, 0, 0, 0, 0, 0, 0)	100000	−2.59	(2, 0, 0, 0, 0, 0, 0, 0)	100000	−2.99
(0, 4, 0, 0, 0, 0, 0, 0)	100000	−4.72	−310277	(0, 4, 0, 0, 0, 0, 0, 0)	100000	−6.02	(0, 4, 0, 0, 0, 0, 0, 0)	100000	−9.82
(0, 0, 1, 0, 0, 0, 0, 0)	15000	−3.32	−239770	(0, 0, 1, 0, 0, 0, 0, 0)	15000	−4.90	(0, 0, 1, 0, 0, 0, 0, 0)	15000	−10.45
(0, 0, 0, 0, 0, 0, 3, 0)	13500	−3.54	−252791	(0, 0, 0, 0, 0, 0, 3, 0)	13500	−5.77	(0, 0, 0, 0, 0, 0, 3, 0)	13500	−15.31
(0, 0, 0, 0, 0, 0, 0, 1)	40000	−4.07	−280640	(0, 0, 0, 0, 0, 0, 0, 1)	40000	−6.52	(0, 0, 0, 0, 0, 0, 0, 1)	40000	−17.72
(0, 4, 0, 1, 0, 0, 0, 0)	100000	−5.09	−325518	(0, 4, 0, 1, 0, 0, 0, 0)	100000	−6.49	(0, 4, 0, 1, 0, 0, 0, 0)	100000	−10.53
(0, 1, 1, 1, 2, 0, 0, 0)	100000	−6.82	−383989	(0, 1, 1, 1, 2, 0, 0, 0)	100000	−8.69	(0, 1, 1, 1, 2, 0, 0, 0)	100000	−14.12
(0, 0, 0, 1, 2, 0, 0, 1)	100000	−6.86	−385086	(0, 0, 0, 1, 2, 0, 0, 1)	100000	−8.88	(0, 0, 0, 1, 2, 0, 0, 1)	100000	−15.15

- Mimicking the argument, we estimated $\rho_5 = 0.15$ and $\rho_6 = 0.4$ for the detention proportion θ, with concentration parameter $\delta_\theta = 0.08$.

Finally, we assessed the risk coefficient k_D in the operator's utility function. We used the probability equivalent (PE) method, see (Farquhar, 1984), to assess a few values for the utility function and, then, fit an appropriate curve through least squares, obtaining a good fit for $k_D = 5 \cdot 10^{-6}$. With respect to pickpockets, based also on the PE method, and taking into account the Defender's uncertainty about the pickpockets' behavior, we assessed the maximum risk coefficient in the pickpockets' utility function, $K_{A_2} = 10^{-5}$.

The assessment of the parameters in relation with the fare evasion threat is based on the discussion by Ríos Insua et al. (2015b), although, as mentioned, here we have to incorporate the cascading effect arising from pickpocketing. We briefly outline the results for those parameters needed to solve the bi-threat problem. The average number N of customers will depend on the number t of thefts through the average business level b as expressed in (6). We modeled N as a Poisson distribution, $N \sim \mathcal{P}ois(\lambda)$, with $\lambda = b/0.75$, inheriting the uncertainty in b. In addition, (N_1, N_2, N_3) will follow (conditionally independent given d_c) Poisson distributions with parameters $\lambda_1 = \lambda(1 - \phi(d_c))$, $\lambda_2 = \lambda\phi(d_c)(1 - q(d_1))$ and $\lambda_3 = \lambda\phi(d_c)q(d_1)$, respectively, being $\phi(d_c)$ the proportion of fraudsters and $q(d_1)$ that of customers inspected, see (Ríos Insua et al., 2015b) for details. For the number of colluder operations M, we use a Poisson-gamma model, with diffuse, but proper, prior for the Poisson parameter μ, see (French & Ríos Insua, 2000). We obtained that, a posteriori, $\mu|data \sim \mathcal{G}(150000.1, 5.1)$, which shall be estimated, when necessary, through its posterior expectation $E(\mu|data) \approx 30000$.

For the proportion of fraudsters we assume a model $\phi(d_c) = \phi_0 \cdot \exp\left(-\sum_{k=1}^{5} \gamma_k d_k\right) + \phi_r$. We use a beta-binomial model for the current fraud proportion, $(\phi_0 + \phi_r)$, with a noninformative prior, see (French & Ríos Insua, 2000). Based on information provided by the operator, we get the posterior $\mathcal{B}e(3 \cdot 10^4 + 1, 10^6 + 1)$, with expected value 0.03 and negligible variance. The reduced target fare evasion proportion is 0.01. The γ_k's were assessed through expert elicitation, much the same as we did for the μ_j's, obtaining $\gamma_1 = 0.13$, $\gamma_2 = 0.72$, $\gamma_3 = 0.45$, $\gamma_4 = 0.23$ and $\gamma_5 = 0.84$. Other relevant parameters are the fare ticket (for the fare evasion threat, we consider the single-ride fare, ν_c =$2) and the average fine in case someone is caught without a valid ticket ($100). However, according to the metro operator, approximately only one sixth of the imposed fines are actually paid off, giving an effective average fine per caught evader of, roughly, $17. We shall use this value for f_c in our computations.

3.1.5. Solution

We discuss the case solution. We have simulated 10,000 years of operations for each portfolio d, to identify the optimal one. Fig. 7 shows the estimated expected utility of the 324 feasible portfolios. From left to right, the portfolio numbering begins with portfolio #1, $d = (0, 0, 0, 0, 0, 0, 0, 0)$; portfolio #2, $d = (0, 0, 0, 0, 0, 0, 0, 1)$; and so on, increasing sequentially the values in $d_8, d_7, \ldots, d_2$ and d_1, being the last feasible portfolio #324, $d = (2, 0, 0, 1, 0, 0, 0, 0)$. The optimal portfolio is $d^* = (1, 0, 0, 0, 0, 1, 0, 0)$, #276, corresponding to hiring one inspector and one patrol. Its estimated expected utility is -2.03, associated investment $85,000, and global expected decrease in income for the operator $171,585 (due to the investment, plus the expected balance between the fraud and the collected fines, which is −$42,980, and the expected reduction in business level, which amounts to $43,605). The next two portfolios with highest expected utilities are $d^{**} = (1, 0, 0, 0, 0, 1, 1, 0)$, corresponding to one inspector, one patrol and one camera, with associated investment of $89,500 and expected losses of $177,492; and $d^{***} = (1, 0, 0, 1, 0, 1, 0, 0)$, corresponding to one inspector and one patrol, and the involvement of clerks in observation tasks, with associated investment of $85,000, and expected losses of $185,656. As we can observe, when the operator faces multiple threats with similar impact, she has to distribute her available resources to fight against all threats. The optimal portfolio includes countermeasures specific to each threat: inspectors and patrols. Note that although patrols are more expensive than guards, they are preferable because of their higher deterrent effect.

The left column of Table 3 shows similar results for other relevant portfolios. We have included (when feasible) those portfolios for which the investment is maximum in one of the countermeasures, with no investment in the other ones. We have also considered those portfolios with highest investments. As we can observe, these are not necessarily the most effective ones in terms of operator's expected utility. For instance, the portfolio in the last row exhausts the available budget, and incurs in additional costs associated with the change in clerk duties. However, its expected utility is worse (-6.86) than that of the optimal portfolio (-2.03).

Due to the complexity of the model and that many parameters are assessed judgmentally, we need to perform sensitivity analysis. We assess here possible sensitivity to variations in the fare evasion proportion. Note first that the proportion estimated above ($\phi_0 + \phi_r = 0.03$) is not constant but, rather, depends on the specific day and time considered, varying between 0.005 and 0.12, according to the operator. We are interested in evaluating the impact of higher proportions on the operator's costs. For this, we repeat the previous calculations for proportions 0.06 and 0.12, shown in the central and right columns of Table 3, respectively. As we can observe, under higher fare evasion proportions, the operator needs to make bigger investments. The optimal portfolio in both cases is (2, 0, 0, 0, 0, 0, 0, 0), corresponding to hiring two inspectors, with associated investment of $100,000 and expected losses of $190,207 and $137,218, respectively. The second best portfolio is (1, 0, 0, 0, 0, 1, 0, 0) when $\phi_0 + \phi_r = 0.06$; and (2, 0, 0, 1, 0, 0, 0, 0) when the proportion is 0.12: just the optimal portfolio plus the expenses associated with the change in clerk duties. Under these settings, hir-

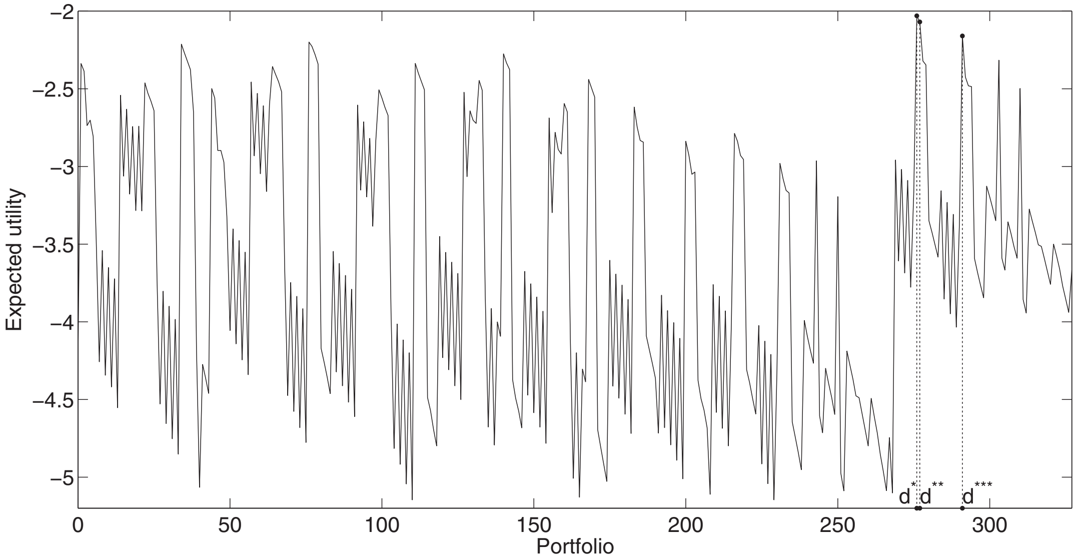

Fig. 7. Operator's estimated expected utility when both threats are present.

ing as many inspectors as possible becomes crucial, as they have authority for imposing fines.

These results suggest that when the relative impact of one of the threats dominates, the operator might need to reallocate resources to better fight against it, possibly unprotecting herself from the other threats. However, we found model performance sensitive to several other factors, especially to variations on the proportion of tickets inspected by each new inspector. Thus, it is essential that inspectors indeed carry out their task, so as to ensure an effective fight against fare evasion.

4. Multisite multithreat protection

We consider now the multisite case. An organization needs to protect from m threats over n sites. The threats are uncoordinated and the sites are not spatially related. The strategy we apply is to deploy one of the models in Section 2 over each site. Resource constraints for the Defender and for each of the attackers coordinate the models. The Defender and each of the attackers aggregate values attained at nodes, applying their utility function.

The Defender deploys defensive resources d_j over site j, $j = 1, \dots, n$. These must fulfill certain constraints which we represent through $g(\boldsymbol{d}) \in \mathcal{D}$, where $\boldsymbol{d} = (d_1, \dots, d_n)$. This might include, among others, financial constraints concerning a maximum budget; logistic constraints, like the impossibility of deploying certain resources separately (e.g. a sniffer dog has to be always accompanied by trained personnel), or the requirement of having some critical infrastructure protected 24/7; or political constraints, like the need of having each site minimally protected. The ith attacker will perform attack a_{ij} over the jth site. In turn, each attacker's strategy should satisfy certain constraints $h_i(\boldsymbol{a}_i) \in \mathcal{A}_i$, where $\boldsymbol{a}_i = (a_{i1}, \dots, a_{in})$, $i = 1, \dots, m$. This might also account for financial constraints, like a limited budget to buy sophisticated weapons or instruct hackers; or human resource constraints, like the need of having, ideally, a minimum number of attackers over each site, among others. The interaction between the Defender and the ith attacker over site j will yield a random result S_{ij}.

The Defender aggregates her results through $u_D(\boldsymbol{d}, \boldsymbol{s}_1, \dots, \boldsymbol{s}_m)$, where $\boldsymbol{s}_1 = (s_{11}, \dots, s_{1n}), \dots, \boldsymbol{s}_m = (s_{m1}, \dots, s_{mn})$. She needs to find her optimal defense strategy, $\boldsymbol{d}^*$, subject to the corresponding constraints. Under appropriate conditional independence assumptions over attack results, she needs to build the conditional models $p_D(s_{ij}|d_j,\ a_{ij})$, $i = 1, \dots, m$, $j = 1, \dots, n$, expressing her uncertainty about the outcome s_{ij} of the attack a_{ij} launched by attacker A_i over site j when she has deployed defensive resources d_j. By integrating out such uncertainty, she will get her expected utility:

$$\psi_D(\boldsymbol{d}|\boldsymbol{a}_1, \dots, \boldsymbol{a}_m) = \int \cdots \int u_D(\boldsymbol{d}, \boldsymbol{s}_1, \dots, \boldsymbol{s}_m)\, p_D(s_{11}|d_1, a_{11}) \cdots \\ p_D(s_{mn}|d_n, a_{mn})\, \mathrm{d}\boldsymbol{s}_1 \dots \mathrm{d}\boldsymbol{s}_m.$$

Suppose the Defender is able to build the models $p_D(a_{ij}|d_j)$, $i = 1, \dots, m$, $j = 1, \dots, n$, reflecting her beliefs about which attack will be chosen by attacker A_i against site j, when protected by defensive resources d_j. Then, she will be able to compute

$$\psi_D(\boldsymbol{d}) = \int \cdots \int \psi_D(\boldsymbol{d}|a_{11}, \dots, a_{mn})\, p_D(a_{11}|d_1) \cdots p_D(a_{mn}|d_n) \\ \times \mathrm{d}a_{11} \dots \mathrm{d}a_{mn},$$

and solve max $\psi_D(\boldsymbol{d})$ subject to $g(\boldsymbol{d}) \in \mathcal{D}$.

The assessments $p_D(a_{ij}|d_j)$, $i = 1, \dots, m$, $j = 1, \dots, n$ are non-standard. Following the strategy in Section 2, we can solve separately the problem for attacker A_i, who attacks sites 1 to n, subject to constraints $h_i(\boldsymbol{a}_i) \in \mathcal{A}_i$, $i = 1, \dots, m$. As an example, in order to solve the problem faced by attacker A_1, the Defender would need his utility $u_{A_1}(\boldsymbol{a}_1, \boldsymbol{s}_1)$ and probabilities $p_{A_1}(s_{1j}|d_j, a_{1j})$, $j = 1, \dots, n$. Then, she would solve the optimization problem

$$\boldsymbol{a}_1^*(\boldsymbol{d}) = \underset{h_1(\boldsymbol{a}_1) \in \mathcal{A}_1}{\operatorname{argmax}} \int \cdots \int u_{A_1}(\boldsymbol{a}_1, \boldsymbol{s}_1)\, p_{A_1}(s_{11}|d_1, a_{11}) \cdots \\ \times p_{A_1}(s_{1n}|d_n, a_{1n})\, \mathrm{d}s_{11} \dots \mathrm{d}s_{1n}.$$

However, the Defender does not know u_{A_1} and the p_{A_1}'s. To model her uncertainty about them, she uses random utilities and probabilities $\big(U_{A_1}, P_{A_1}(s_{11}|\cdot), \dots, P_{A_1}(s_{1n}|\cdot)\big)$ and, then, propagates the uncertainty to obtain the m-dimensional random optimal action

$$\boldsymbol{A}_1^*(\boldsymbol{d}) = \underset{h_1(\boldsymbol{a}_1) \in \mathcal{A}_1}{\operatorname{argmax}} \int \cdots \int U_{A_1}(\boldsymbol{a}_1, \boldsymbol{s}_1)\, P_{A_1}(s_{11}|d_1, a_{11}) \cdots \\ \times P_{A_1}(s_{1n}|d_n, a_{1n})\, \mathrm{d}s_{11} \dots \mathrm{d}s_{1n}.$$

Then, she would get $p_D(\boldsymbol{A}_1 \le \boldsymbol{a}_1|\boldsymbol{d}) = \Pr(\boldsymbol{A}_1^*(\boldsymbol{d}) \le \boldsymbol{a}_1)$. For an estimate of $p_D(\boldsymbol{a}_1|\boldsymbol{d})$, we could proceed through a sampling scheme similar to that in Algorithm 1. The problems faced by the other attackers $A_2 \dots, A_m$ will be solved in the same way, providing estimates $\widehat{p_D}(\boldsymbol{a}_i|\boldsymbol{d})$, $i = 2, \dots, m$ of the required probabilities. Extensions similar to those provided in Section 2 could be given here.

Table 4
Optimal portfolio for the bi-threat problem in four stations.

	d_1	d_2	d_3	d_4	d_5	d_6	d_7	d_8	Invest. (−)	Fines (+)	Loss fare (−)	Loss pick. (−)
S_1	0	0	0	–	0	1	0	–	35,000	–	101,938	42,595
S_2	0	0	0	–	0	1	0	–	35,000	–	114,280	33,757
S_3	1	0	1	–	0	0	0	–	65,000	162,688	234,401	127,994
S_4	0	0	2	–	0	1	0	–	65,000	–	394,731	78,290
Total	1	0	3	1	0	3	0	0	200,000	162,688	845,170	282,636

5. Protecting from fare evasion and pickpocketing at several stations

We extend the case in Section 3 to several stations. The railway system analyzed in our case study comprises more than 100 stations. To better illustrate ideas, we consider a small representative group of $n = 4$ stations. As described before, we build a model like that in Section 3 for each station. Thus, for station $j \in \{1, 2, 3, 4\}$, we deploy resources $d_j \equiv (d_{j1}, d_{j2}, d_{j3}, d_{j5}, d_{j6}, d_{j7})$, with meanings as in Section 3.1. The decision d_4 on whether or not to change the clerk duties is made for the whole network; the associated costs will be proportional to the number of stations. The investment in an awareness plan is also common for the entire network, and will be denoted $d_8 \in \{0, 1\}$, with the same meaning as in Section 3.1. If the operator has a global budget B for investing in new countermeasures, the resources will have to fulfill the constraints

$$\sum_{j=1}^{4}\left(\sum_{\substack{k=1\\k\neq 4}}^{7} q_k d_{jk}\right) + q_8 d_8 \leq B,$$

$$0 \leq \sum_{j=1}^{4} d_{jk} \leq \bar{d}_k,\ \ k = 1, \ldots, 7,\ \ k \neq 4,$$

$$d_{jk}\ \text{integer},\ \ j = 1, \ldots, 4,\ \ k = 1, \ldots, 7,\ \ k \neq 4,$$

$$d_{j3} \leq \bar{d}_{j3},\ \ j = 1, \ldots, 4,$$

$$d_4, d_8 \in \{0, 1\}.$$

Here, $\bar{d}_{j3}$ is the maximum number of secured automatic access doors that may be replaced at site j, and the $\bar{d}_k$'s, the overall maximum allowable number for each resource. Finally, note that some additional constraints could possibly apply for certain sites.

We provide specific constraints and parameters for the example. Stations 1–3 have an average annual flow of customers of 1,000,000, whereas Station 4 has 5,000,000 under current operational conditions. The security additional budget is $200,000 for the network, on top of the current one. There is an additional requirement that the investment at each station has to lie between $30,000 and $70,000, except for Station 4, in which the minimum investment has to be $50,000. Besides, for image reasons, the investment in the whole network has to be, at least, $120,000. Resource upper bounds are $\bar{d}_k = 4,\ k = 1, 2, 3, 5, 6,$ and $\bar{d}_7 = 8$. Moreover, the operator indicates that, at most, two units of each countermeasure should be deployed at a single station.

As the impact of both threats over the four incumbent stations regards:

- Stations 1 and 2. There are moderate levels of fare evasion, $\phi_0 + \phi_r = 0.03,\ M = 30000$; and pickpocketing: we assume a range $\{0, 1, \ldots, 100\}$ for t.
- Station 3. There is a high level of fare evasion, $\phi_0 + \phi_r = 0.12$, $M = 120000$ and a moderate level of pickpocketing. This is representative of stations not so well protected against fare evasion. It is required to hire, at least, one inspector.
- Station 4. There is a moderate level of fare evasion and a high pickpocketing level: we assume a range of values $t \in \{0, 1, \ldots, 150\}$. This is representative of pickpocketing hotspots, typically busy stations close to main transport hubs. In this station, the presence of, at least, a patrol is required.

For simplicity, we consider just one group of pickpockets operating at each station, although they frequently belong to the same gang and move between stations. We assume that countermeasures and attackers are static, in that they are not allowed to move between stations. This may sound unrealistic, but we have to keep in mind that we are planning security in annual terms. Operational decisions, like patrolling routes, may be decided at a later stage, see our final discussion.

There are 26 decision variables and 16 constraints. The number of security portfolios is too large to implement the enumeration strategy in Section 3. Alternatively, we use a genetic algorithm, see (Goldberg, 1989), with fitness function given by the operator's expected utility, which generalizes (8) by including contributions from the four stations. In our computations, we have used the built-in `ga` function with default options, implemented in `MATLAB R2013b` (The MathWorks, 2013). Specifically, the algorithm stopped when the weighted average relative change in the best fitness function value over 50 generations was less than or equal to 10^{-6}. After 10,000 replications, we obtained the optimal portfolio, shown in Table 4, together with relevant information about investments at each station, expected money collected through fines, and expected losses due to fare evasion and pickpocketing.

As we can observe, investing in door guards, cameras and the awareness plan seems not worthwhile for the operator, given the constraints. She should involve ticket clerks in observation tasks, with an impact on costs. Investments at Stations 1 and 2 coincide, as expected, since both have similar features: the operator invests $35,000 in one patrol. At Station 3, the main issue was fare evasion. Thus, in addition to the required inspector, the optimal portfolio suggests installing an automatic access door, with associated overall investment of $65,000. Finally, Station 4 was the busiest one, implying a potentially greater impact of fare evasion and, especially, pickpocketing. The presence of at least one patrol was mandatory. Additionally, two automatic access doors should be installed at this station, with a global investment of $65,000. Under this policy, the annual expected losses for the operator are $1,225,118, corresponding to $200,000 of investments (plus $60,000 of negotiation costs with the unions concerning clerk duties), $682,482 in the fraud/fine balance, and $282,636 of business lost due to pickpocketing. This might seem a large amount, but keep in mind that, should the operator not invest in new countermeasures, the expected loss would be around $2.5 m: thanks to the deployed portfolio, the operator is able to halve expected losses. Other portfolios entail minor changes with respect to the optimal portfolio and have similar expected utility. For instance, the second best portfolio, with annual expected losses of $1,229,250 for the operator, simply changes the decision on where to install an automatic access door: from Station 3 to either Station 1 or 2.

Computations took around seven hours on a standard laptop running `Windows`. Table 5 summarizes computing times for several stations and replications

As we can observe, times are linear in the number of replications but nonlinear in the number of stations. Indeed, for the

Table 5
Computing times (in hours).

# rep	$n = 4$	$n = 8$	$n = 16$	$n = 32$
1000	0.7	2.7	6.6	15.6
2000	1.4	5.5	13.2	31.2
3000	2.1	8.2	19.9	46.9

whole network, with more than 100 stations, the computation takes over one week with 10,000 replications. Even though computations could be improved with faster parallel machines, recall that this is a tactical decision made yearly, making such computing times actually acceptable.

Beyond the computational burden induced by a large number of sites, note that there might be considerable elicitation burden. In our case study, we have assumed the same responses at various stations, thus with the same parameters. But it could be the case that different stations have different responses. One way to mitigate such elicitation burden would be to somehow cluster sites with similar features and perform the elicitation for just one of the stations in the cluster.

6. Discussion

We have provided a methodology for protecting multiple sites from multiple uncoordinated threats, based on ARA. First, we have dealt with the multithreat problem over a single site, deploying a Sequential Defend-Attack model for each attacker. Then, we have extended the formulation to multiple sites, using one of those models over each site, with models coordinated by resource constraints for each participant, and value aggregation over various sites and, for the case of the Defender, also across various threats. We have illustrated the approach with a case study in railway security, in which an operator is concerned with the impact of fare evasion and pickpocketing over the quality and efficiency of the service. Although colluders and pickpockets do not coordinate their attacks, there is a cascading effect between them: the actions of pickpockets will affect the number of customers through the business level, and this might have an influence over the fare evasion result.

Our example referred to crime but similar ideas may be applied to fighting terrorism. The key difference would be in the adversarial objectives. Crime cases would refer more to business-like objectives, whereas terrorism cases would rather refer to political objectives. see (Keeney & von Winterfeldt, 2010) for references. Note that both types of adversaries might appear in a same problem. For example, our railway operator, besides fare evasion and pickpocketing, might be interested in fighting terrorism and graffiti.

Several issues remain to be addressed. We have assumed that attackers responsible of different types of threats are uncoordinated. However, it would be conceivable that they are coordinated. For instance, we could envisage a scenario in which a terrorist group shares its zone of influence with other criminal organizations as, e.g., drug dealers or a local mafia. By coordinating their attacks over different sites, the attackers could take advantage of their own and others' resources, allocating them so as to inflict as much damage as possible to the Defender, obtaining higher revenue than if attacking separately.

The chosen model is static, since we have not allowed for mobility of resources. This is sufficient for our purposes, as we refer to annual planning. One way to tackle this issue would be to consider a model allowing for further interactions among the Defender and the attackers: we could assume that the attackers have some degree of mobility between different sites, trying e.g. to move away from better protected sites or, alternatively, concentrating their attacks on the most valuable targets. This may be dealt with more dynamic models, like the Sequential Defend-Attack-Defend model, see (Brown et al., 2006) or (Ríos & Ríos Insua, 2012). Alternatively, the approach here may be seen at the tactical level, deciding what resources to deploy. Once this has been resolved, we would decide the patrolling schedule at an operational level, with models as in e.g. Alpern, Morton, and Papadaki (2011); Brown, Saisubramanian, Varakantham, and Tambe (2014); or Zoroa, Fernández-Sáez, and Zoroa (2012).

Beyond the standard parametric sensitivity analysis performed, a key assumption would be whether the adversary actually performs an expected utility analysis. Thus, we would face a concept uncertainty issue. By Ríos Insua, Banks, and Ríos (2015a), we have described how to deal with other rationality types of adversaries and how to mix such types.

A final important issue refers to validation. The model can be validated in, at least, four ways: (1) Assumptions and elicitations may be validated by third parties in dedicated seminars, as we performed in the workshops mentioned. (2) Some model components were data based and we used standard (Bayesian) goodness-of-fit approaches. Others were based on expert judgment and we could use consistency checks, as explained for the μ parameters in Section 3.1.4. (3) Other aspects could be validated through sensitivity analysis, as discussed in Section 3.1.5. (4) In addition, the actual validation could be done after applying the model under the same circumstances. Note, however, that these are dynamic. For example, an increasingly severe crisis might make population more prone to evade fares or turn more people into pickpocketing as a means of life. Finally, another way to ascertain the validity of the solution would be to test several competing solutions simultaneously over different homogeneous clusters of sites, and check whether the proposed solution is providing the best response.

Acknowledgments

The work of DRI has been supported by the Spanish Ministry of Economy and Innovation program MTM2014-56949-C3-1-R, the AXA-ICMAT Chair on Adversarial Risk Analysis and the ESF Cost Action IS1304 on Expert Judgement. This project has received funding from the European Union's Seventh Framework Programme for Research, Technological Development and Demonstration under grant agreement 285223. We are grateful to the transportation experts and stakeholders participating at validation workshops for fruitful discussions.

References

Alpern, S., Morton, A., & Papadaki, K. (2011). Patrolling games. *Operations Research, 59*(5), 1246–1257.

Bier, V. M., Oliveros, S., & Samuelson, L. (2007). Choosing what to protect: strategic defensive allocation against an unknown attacker. *Journal of Public Economic Theory, 9*(4), 563–587.

Brown, G., Carlyle, M., Salmerón, J., & Wood, K. (2006). Defending critical infrastructure. *Interfaces, 36*(6), 530–544.

Brown, M., Saisubramanian, S., Varakantham, P. R., & Tambe, M. (2014). STREETS: game-theoretic traffic patrolling with exploration and exploitation. In *Innovative applications in artificial intelligence (IAAI). Twenty Eighth AAAI Conference on Artificial Intelligence, AAAI-14*. Research Collection School of Information Systems.

Dyer, J. S., & Sarin, R. K. (1979). Measurable multiattribute value functions. *Operations Research, 27*(4), 810–822.

Dyer, J. S., & Sarin, R. K. (1982). Relative risk aversion. *Management Science, 28*(8), 875–886.

Enders, W., & Sandler, T. (2011). *The political economy of terrorism* (2nd edition). New York: Cambridge University Press.

Ezell, B. C., Bennett, S. P., von Winterfeldt, D., Sokolowski, J., & Collins, A. J. (2010). Probabilistic risk analysis and terrorism risk. *Risk Analysis, 30*(4), 575–589.

Farquhar, P. H. (1984). State of the art—utility assessment methods. *Management Science, 30*(11), 1283–1300.

Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics, 1*(2), 209–230.

French, S., & Ríos Insua, D. (2000). *Statistical decision theory*. London: Arnold.

Gil, C., Ríos, J., & Ríos Insua, D. (2016). Adversarial risk analysis for urban security resource allocation. *Risk Analysis, To appear*.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.

Haberfeld, M. R., & von Hassell, A. (2009). A new understanding of terrorism: case studies, trajectories and lessons learned. *Humanities, Social Sciences and Law*. New York: Springer.

Haphuriwat, N., & Bier, V. M. (2011). Trade-offs between target hardening and overarching protection. *European Journal of Operational Research, 213*(1), 320–328.

Hausken, K. (2014a). Choosing what to protect when attacker resources and asset valuations are uncertain. *Operations Research and Decisions, 24*(3), 23–44.

Hausken, K. (2014b). Individual versus overarching protection and attack of assets. *Central European Journal of Operations Research, 22*(1), 89–112.

Hausken, K., & Bier, V. M. (2011). Defending against multiple different attackers. *European Journal of Operational Research, 211*(2), 370–384.

Hausken, K., & Levitin, G. (2012). Review of systems defense and attack models. *International Journal of Performability Engineering, 8*(4), 355–366.

Keeney, G. L., & von Winterfeldt, D. (2010). Identifying and structuring the objectives of terrorists. *Risk Analysis, 30*(12), 1803–1816.

Keeney, R. L. (2007). Modeling values for anti-terrorism analysis. *Risk Analysis, 27*(3), 585–596.

Keeney, R. L., & von Winterfeldt, D. (2011). A value model for evaluating homeland security decisions. *Risk Analysis, 31*(9), 1470–1487.

Koller, D., & Milch, B. (2003). Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior, 45*(1), 181–221.

Kraan, B., & Bedford, T. (2005). Probabilistic inversion of expert judgments in the quantification of model uncertainty. *Management Science, 51*(6), 995–1006.

Levitin, G., Hausken, K., & Dai, Y. (2014). Optimal defense with variable number of overarching and individual protections. *Reliability Engineering & System Safety, 123*, 81–90.

O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., … Rakow, T. (2006). *Uncertain judgements: eliciting experts' probabilities*. Chichester, West Sussex: John Wiley & Sons.

Parnell, G. S., Banks, D., Borio, L., Brown, G., Cox Jr, L. A. T., Gannon, J., … Wilson, A. (2008). Report on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis. National Academies Press.

Reddy, A. V., Kuhls, J., & Lu, A. (2011). Measuring and controlling subway fare evasion. *Transportation Research Record: Journal of the Transportation Research Board, 2216*(1), 85–99.

Ríos, J., & Ríos Insua, D. (2012). Adversarial risk analysis for counterterrorism modeling. *Risk Analysis, 32*(5), 894–915.

Ríos Insua, D., Banks, D., & Ríos, J. (2015a). Modeling opponents in adversarial risk analysis. *Risk Analysis*, n/a–n/a. doi:10.1111/risa.12439.

Ríos Insua, D., Cano, J., Pellot, M., & Ortega, R. (2015b). *Current trends in bayesian methodology with applications* (pp. 311–334). Boca Raton, Florida: Chapman and Hall/CRC.

Ríos Insua, D., Ríos, J., & Banks, D. (2009). Adversarial risk analysis. *Journal of the American Statistical Association, 104*(486), 841–854.

Rossman, A. J., Short, T. H., & Parks, M. T. (1998). Bayes estimators for the continuous uniform distribution. *Journal of Statistics Education, 6*(3), 1–7.

Sunstein, C. R. (2007). *Worst-case scenarios*. Boston: Harvard University Press.

The MathWorks (2013). Matlab and global optimization toolbox release 2013b. The MathWorks, Inc.Natick, Massachusetts, United States.

Troelsen, L. H., & Barr, L. (2012). Combating pickpocketing in public transportation. *Public Transport International, 61*(1), 32–33.

Wang, C., & Bier, V. M. (2013). Expert elicitation of adversary preferences using ordinal judgments. *Operations Research, 61*(2), 372–385.

Wein, L. M. (2009). OR forum—Homeland security: from mathematical models to policy implementation. *Operations Research, 57*(4), 801–811.

World Economic Forum (2015). *World economic forum. global risks*. URL http://www3.weforum.org/docs/WEF_Global_Risks_2015_Report15.pdf.

Yang, R., Kiekintveld, C., Ordóñez, F., Tambe, M., & John, R. (2013). Improving resource allocation strategies against human adversaries in security games: an extended study. *Artificial Intelligence, 195*, 440–469.

Zhuang, J., & Bier, V. M. (2007). Balancing terrorism and natural disasters—defensive strategy with endogenous attacker effort. *Operations Research, 55*(5), 976–991.

Zoroa, N., Fernández-Sáez, M. J., & Zoroa, P. (2012). Patrolling a perimeter. *European Journal of Operational Research, 222*(3), 571–582.