European Journal of Operational Research 203 (2010) 409–418

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Modeling secrecy and deception in a multiple-period attacker–defender signaling game

Jun Zhuang [a,*], Vicki M. Bier [b], Oguzhan Alagoz [b]

[a] Department of Industrial and Systems Engineering, University at Buffalo, The State University of New York, United States
[b] Department of Industrial and Systems Engineering, University of Wisconsin-Madison, United States

ARTICLE INFO

Article history:
Received 23 February 2009
Accepted 29 July 2009
Available online 3 August 2009

Keywords:
Game theory
Signaling game
Secrecy and deception
Truthful disclosure
Dynamic programming
Multi-period game

ABSTRACT

In this paper, we apply game theory to model strategies of secrecy and deception in a multiple-period attacker–defender resource-allocation and signaling game with incomplete information. At each period, we allow one of the three possible types of defender signals—truthful disclosure, secrecy, and deception. We also allow two types of information updating—the attacker updates his knowledge about the defender type after observing the defender's signals, and also after observing the result of a contest (if one occurs in any given time period). Our multiple-period model provides insights into the balance between capital and expense for defensive investments (and the effects of defender private information, such as defense effectiveness, target valuations, and costs), and also shows that defenders can achieve more cost-effective security through secrecy and deception (possibly lasting more than one period), in a multiple-period game.

This paper helps to fill a significant gap in the literature. In particular, to our knowledge, no past work has studied defender secrecy and deception in a multiple-period game. Moreover, we believe that the solution approach developed and applied in this paper would prove useful in other types of multiple-period games.

1. Introduction

Most applications of game theory to homeland-security resource allocation so far have involved only one-period games (in which the payoffs are realized all at once, even if the players move sequentially). Dresher (1961) was among the first researchers to apply game theory to military strategic interactions. However, he did not explicitly model deception and secrecy. More recently, Hausken (2008) has applied game theory to study the strategic interaction between a defender and an attacker for series and parallel reliability systems, and Levitin and Hausken (2009) studied the benefits of using "false" targets that the attacker cannot distinguish from the desired targets in an attacker–defender game.

Crawford (2003) modeled an attacker–defender sequential game allowing for bounded rationality. In particular, bounded rationality can yield an equilibrium with deceptive signals even when signaling is costless and noiseless. Powell (2007) studied an attacker–defender, multiple-target game where the defender has some private information about the vulnerability of the various targets. He found an equilibrium in which the defender always "pools"; i.e., the defender allocates her defensive investments without regard to the vulnerability of various targets, so that the attacker cannot infer their vulnerability. We interpret this as a form of secrecy.

Secrecy has also been modeled as simultaneous play in game theory (see for example Zhuang and Bier, 2007), since in a simultaneous game, each player moves without knowing the moves of other players. Note that this does not actually require both players to make their decisions at the same time—the players can be viewed as being engaged in a simultaneous game as long as neither party knows the other's decision at the time he makes his own decision.

Brown et al. (2005) studied secrecy in a zero-sum attacker–defender game in the context of ballistic missile deployment, but failed to include the potential for the attacker to endogenously update his beliefs. In particular, in Brown et al. (2005), the attacker is assumed to be unaware of even the defender *options* when the defender chooses secrecy, while a typical endogenous model usually assumes only that the attacker is unaware of the specific choice made by the defender.

While the definition of secrecy is relatively straightforward, many kinds of deception have been discussed in the literature. Some researchers have modeled deception as sending noisy or imperfect signals to mislead one's opponents. For instance, Hendricks and McAfee (2006) and Oliveros (2005) used the

* Corresponding author. Tel.: +1 716 645 4707; fax: +1 716 645 3302.
E-mail address: jzhuang@buffalo.edu (J. Zhuang).

0377-2217/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2009.07.028

Normandy invasion as an example to argue that the first mover (the Allies) optimally allocated resources to targets that they did not intend to attack, in order to mislead the Germans about their true landing place by sending noisy signals. Similarly, Hespanha et al. (2000) and Brown et al. (2005) defined deception in a zero-sum attacker–defender game as occurring when the defender discloses only a subset of the defenses, in an attempt to route attacks to heavily-defended locations.

There is a substantial economics literature on *principal-agent* (or, more generally, *mechanism-design*) problems, which addresses how the first mover (principal) can provide incentives (usually by contract) to the second mover (agent), to ensure that the second mover chooses the preferred action. This literature usually allows hidden actions and/or private information on the part of the second mover (Zhang and Zenios, 2008; Doepke and Townsend, 2006), and may further address the issue of information disclosure by the second mover (Prat, 2005), but not the first mover. By contrast, Zhuang and Bier (in press) argued that hidden actions and private information on the part of the second mover did not change the defender's preference for truthful disclosure in their model.

Another body of economics literature focuses on *revelation of private information*, addressing how individual players might either truthfully or deceptively disclose their *private information* (but not their actions); See for example Neeman, 2004. Likewise, several researchers have addressed the question of information disclosure about player attributes (rather than actions) in supply-chain management (see for example Li, 2002; Raghunathan, 2003; Chu et al., 2006; Yao et al., 2008; Leng and Parlar, 2009).

Zhuang and Bier (2007) studied the balance between protecting from natural disasters and protecting from terrorism in a one-period game; their results indicate that truthful disclosure should always be (weakly) preferred to secrecy, which is not surprising, since their model is a game of complete information. Much other work to date (e.g., Bier et al., 2007) also recommends disclosure. By contrast, Zhuang and Bier (in press) found that defender secrecy and/or deception could be strictly preferred in a one-period game in which the defender has private information (i.e., the attacker is uncertain about the defender type).

In the real world, however, attackers and defenders frequently interact repeatedly over time either through successive attacks (as in Israel, for example), or through successive attacker attempts to "probe" a system before a successful attack (as in the case of computer security). Moreover, many critical defender decisions (e.g., whether to invest in capital defenses versus short-term expenses such as police patrol) also involve time as a critical dimension.

To our knowledge, non-zero-sum, multiple-period, attacker–defender resource-allocation and signaling games with incomplete information and hidden actions have not yet been extensively studied in the literature. The model developed by Coleb and Kocherlakotad (2001) might be the closest to ours. They assumed that at the beginning of each period, the players learn the complete history of the previous period, including both private information and hidden actions. Like Coleb and Kocherlakotad (2001), in Section 3, we assume that the players learn the hidden actions in the history, an assumption that is then relaxed in Section 4. Unlike Coleb and Kocherlakotad (2001), our model also allows private information (such as the defender discount rate, the target valuation, or the relative effectiveness of defensive capital versus expense) to remain secret throughout the game.

This paper begins to fill the gap in analysis of multi-period security games by analyzing an N-period game between a single attacker and a single defender. In particular, this paper extends the model in Zhuang and Bier (in press) to an N-period game. Since Zhuang and Bier (in press) suggested that the defender should always prefer truthful disclosure to secrecy and deception when she does not have private information, in this paper we focus on the case where the defender does have private information (i.e., the attacker does not know some defender attributes, such as asset values, costs, or the relative effectiveness of capital versus expenses), while the attacker does not. In this case, we allow two types of updates about the defender type—the attacker updates his knowledge about the defender type after observing the defender's signals, and also after observing the result of a contest (if one occurs in any given time period). Our analysis shows that there exist equilibria in which secrecy and/or deception are strictly preferred by some types of defenders, in order to mimic defender types that are of less interest to attackers (e.g., defender types that may be less valuable, less cost-effective to attack, or better defended), or to distinguish themselves from defender types that are of greater interest to attackers. To take advantage of the multiple-period nature of our model, we specifically explore secrecy and deception regarding the tradeoff between capital investments versus short-term defenses, since this tradeoff can be studied only in a multi-period model. For simplicity, we consider only a binary choice between capital and expense, but our model could be adapted to non-binary decision variables (e.g., differing levels of investment).

The next section puts forth an N-period model of secrecy and deception in the case where the defender has private information, while the attacker does not. Due to its complexity, our model is not in general analytically tractable. However, some special cases of our model are readily solvable using numerical methods. In particular, Sections 3 and 4 present special cases of our model, and use numerical simulation to show that defender secrecy and deception can sometimes be strictly preferred to truthful disclosure. In e we assume that the attacker learns the defender's defensive investment in any given period at the beginning of the next period; in Section 4, the defensive investment can remain unknown to the attacker. Our numerical simulations provide insights into the most appropriate balance between short-term expenses and capital investments in homeland security, as well as the possible benefits of secrecy and deception. Section 5 summarizes the results of this paper.

2. Model formulation for repeated game

Our game has two players: an attacker (he, signal receiver, A); and a defender (she, signal sender, D). Our model involves a N-period game with private defender information. Fig. 1 provides the sequence of actions for this game. At the beginning of the first period ($t = 1$), nature chooses the defender type θ, which could include numerous defender attributes (e.g., the target valuation, relative effectiveness of short-term expenses versus long-term capital investment, the fraction of capital investment that remains effective in subsequent periods, etc.). For simplicity, we consider only a two-type model; i.e., the defender type θ equals θ_1 with probability p_1 and θ_2 with probability $1 - p_1$, respectively, where p_t is the attacker's probability that the defender is of type θ_1 at the beginning of period t. We assume that p_1, the attacker's prior probability at the beginning of the period 1, is common knowledge to both the attacker and the defender.

For each period $t = 1, \ldots, N$, the decision process is as follows: First, a defender of type θ chooses a strategy $d_t(\theta)$ and a signal $s_t(\theta)$ for $\theta = \theta_1, \theta_2$. We let $d_t(\theta) \in \{0, 1\}$ be a binary decision variable; $d_t(\theta) = 0$ if a defender of type $\theta \in \{\theta_1, \theta_2\}$ invests in short-term expenses (such as police patrol) in period t, and $d_t(\theta) = 1$ if the defender invests in capital defenses in period t. We also let $s_t(\theta) \in \{0, 1, S\}$ be the signal sent by a defender of type θ about its defensive choice (where the "signal" is allowed to include secrecy; or the absence of a signal). We occasionally use quotation

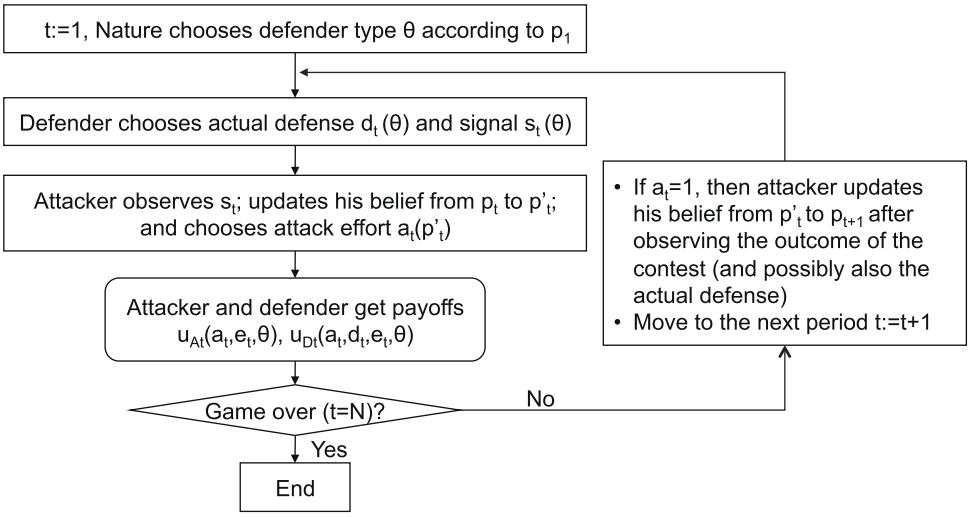

Fig. 1. Sequence of actions for the N-period game with private defender information.

marks, '·', to distinguish a signal from the actual defense, to avoid confusion.

The attacker observes the signal s_t, updates his belief from the prior $p_t \in [0,1]$ (the attacker's prior probability that the defender type is θ_1 at the beginning of period t) to the posterior p'_t (the attacker's posterior probability that the defender type is θ_1), and chooses an attacker response $a_t(s_t) \in \{0,1\}$, where $a_t(s_t) = 0$ is the decision to do nothing during period t, and $a_t(s_t) = 1$ represents the decision to launch an attack. Note that after observing the defender signal, the attacker's belief is updated in a discrete manner (provided that the defenders do not choose which signal to send at random in a mixed-strategy equilibrium). In particular, if both defender types send the same signal at equilibrium, $s_t(\theta_1) = s_t(\theta_2)$, then the attacker is not able to update his belief about the defender type, and we have $p'_t = p_t$; by contrast, if different defender types send different signals at equilibrium, $s_t(\theta_1) \neq s_t(\theta_2)$, then the attacker is able to recognize the defender type with certainty, in which case p'_t equals 1 with probability p_t (which happens when the observed signal s_t is the equilibrium signal for a defender of type θ_1), and 0 with probability $1 - p_t$ (when the observed s_t is the equilibrium signal for a defender of type θ_2).

Then, the attacker and defender payoffs in period t are realized. We assume for simplicity that the actual level of damage to the target is either 100% or zero; therefore, the expected payoffs to the attacker and defender, respectively, are given by

$$u_t^A[a_t(s_t), e_t(\theta), \theta] = \begin{cases} 0 & \text{if } a_t(s_t) = 0 \\ -g_A + c[a_t(s_t), e_t(\theta)]\, v_A & \text{if } a_t(s_t) = 1 \end{cases} \quad (1)$$

and

$$u_t^D[a_t(s_t), d_t(\theta), e_t(\theta), s_t(\theta), \theta] = -g_D[d_t(\theta), s_t(\theta), \theta] - c[a_t(s_t), e_t(\theta)]\, v_D. \quad (2)$$

Here, v_A and v_D are the attacker's and defender's target valuations, respectively (assuming that these valuations are stationary over time); $e_t(\theta)$ is the effective total defense for a defender of type θ in period t, as given by

$$e_t(\theta) = \alpha[1 - d_t(\theta)] + \sum_{k=1}^{t} \rho_{t-k} d_k(\theta), \quad (3)$$

where $\alpha > 1$ is the effectiveness of defender short-term expenses (e.g., police patrol) relative to defender capital investment in security; ρ_{t-k} is the fraction of defensive capital from period k that is still effective in period t and we assume $\rho_0 = 1$; g_A is the cost to the attacker of choosing to attack; $g_D[d_t(\theta), s_t(\theta), \theta]$ is the cost to a defender of type θ of choosing the signal $s_t(\theta)$ in period t, given the actual defense $d_t(\theta)$ (i.e., the cost of implementing truthful disclosure, secrecy, or deceptive disclosure); and $c[a_t(s_t), e_t(\theta)]$ is the conditional probability that an attack would succeed, given the attacker effort a_t and the effective defense e_t. This contest success function is assumed to be of the form

$$c[a_t(s_t), e_t(\theta)] = \Pr\{w_t = 1 | a_t(s_t), e_t(\theta)\} = \frac{a_t(s_t)}{a_t(s_t) + e_t(\theta)}. \quad (4)$$

Hence, the success probability is zero if the attacker chooses not to attack, and is decreasing in the effective defense $e_t(\theta)$ if the attacker does choose to attack. The above functional form is a special case of a contest success function given in Skaperdas (1996).

Finally, if $t = N$, then the game ends, since we are assuming a finite number of periods. Otherwise, if the attacker has chosen to attack, he updates his belief from p'_t to p_{t+1} based on observing the result of the contest (as discussed in detail in the next section), and the game moves to the next period.

2.1. Definition of equilibrium

We let $a(s) \equiv \{a_1(s_1), \ldots, a_N(s_N)\}, d(\theta) \equiv \{d_1(\theta), \ldots, d_N(\theta)\}, s \equiv \{s_1, \ldots, s_N\}, p \equiv \{p_1, \ldots, p_N\}$, and $p' \equiv \{p'_1, \ldots, p'_N\}$ be vectors of state variables and decision variables. Recall that the prior probability p_1 at the beginning of the game is assumed to be common knowledge. Let β_A and β_D be the attacker and defender discount factors, respectively. We also let $U_A[a(s), d(\theta_1), d(\theta_2), p']$ and $U_D[a(s), d(\theta), s(\theta)]$ be the objective functions maximized by the attacker and a defender of type θ, respectively, which are assumed to be sums of the (expected) discounted total payoffs:

$$U_A[a(s), d(\theta_1), d(\theta_2), p'] = \sum_{t=1}^{N} \beta_A^{t-1} \sum_{\theta=\theta_1,\theta_2} u_t^A[a_t(s_t), e_t(\theta), \theta] p'_t(\theta), \quad (5)$$

$$U_D[a(s), d(\theta), s(\theta)] = \sum_{t=1}^{N} \beta_D^{t-1} u_t^D[a_t(s_t), d_t(\theta), e_t(\theta), s_t(\theta), \theta], \quad (6)$$

where the effective defense $e_t(\theta)$ depends on $d_t(\theta)$ and $d_{t-1}(\theta)$ as specified by Eq. (3).

By analogy to perfect Bayesian equilibrium for signaling games without first-mover hidden actions, we define the equilibrium of

our game with defender hidden actions as follows. As in Zhuang and Bier (in press), for computational simplicity, we do not consider mixed-strategy equilibria here, and focus solely on pure-strategy equilibria.

Definition 1. We call the collection $\{a^*(s), d^*(\theta), s^*(\theta), p^*, p'^*\}$ an *equilibrium* if the following four conditions are satisfied:

1. A defender of type θ chooses both the defense strategy $d^*(\theta)$ and the signal strategy $s^*(\theta)$ to maximize her expected payoff, assuming that the attacker will choose his equilibrium response a^*. That is, for all $\theta = \theta_1, \theta_2$, we have

$$d^*(\theta), s^*(\theta) \in \arg\max_{d,s} U_D[a^*(s), d, s, \theta]. \tag{7}$$

2. The attacker chooses his response $a^*(s)$ to maximize his expected payoff, according to his equilibrium posterior distribution p'^* for the defender type, assuming that the two defender types choose their equilibrium defense strategies $d^*(\theta)$. That is, we have

$$a^*(s) \in \arg\max_{a} U_A[a, d^*(\theta_1), d^*(\theta_2), p'^*(s)]. \tag{8}$$

3. The attacker may update his beliefs $p_t'^*$ from p_t^* in period t after observing the signals sent by the defender in a discrete manner. If $s_t^*(\theta_1) = s_t^*(\theta_2)$, then the attacker does not update his beliefs and we have $p_t'^* = p_t^*$. Otherwise, if $s_t^*(\theta_1) \neq s_t^*(\theta_2)$, then $p_t'^*$ equals 1 with probability p_t^* (which happens when the signal observed comes from a defender of type θ_1), and 0 with probability $1 - p_t^*$ (which happens when the signal comes from a defender of type θ_2).
4. At the end of period t, if $a_t = 0$, then no contest happens and we have $p_{t+1}^* = p_t'^*$. Otherwise, if $a_t = 1$, a contest occurs. The attacker then observes the outcome $w_t \in \{0, 1\}$ (the random variable representing success or failure of an attack in period t; $w_t = 1$ if the attack succeeds, and zero otherwise), and updates his belief about the defender's type from $p_t'^*$ to p_{t+1}^* using Bayes' theorem, in which case for all t, we have:

$$p_{t+1}^*(w_t, p_t') = \begin{cases} \frac{p_t'^* c[a_t^*(s_t^*(\theta_1)), e_t^*(\theta_1)]}{p_t'^* c[a_t^*(s_t^*(\theta_1)), e_t^*(\theta_1)] + (1-p_t'^*) c[a_t^*(s_t^*(\theta_2)), e_t^*(\theta_2)]} & \text{if } w_t = 1, \\ \frac{p_t'^* [1 - c[a_t^*(s_t^*(\theta_1)), e_t^*(\theta_1)]]}{p_t'^* [1 - c[a_t^*(s_t^*(\theta_1)), e_t^*(\theta_1)]] + (1-p_t'^*)[1 - c(a_t^*(s_t^*(\theta_2)), e_t^*(\theta_2))]} & \text{if } w_t = 0. \end{cases} \tag{9}$$

2.2. Definition of secrecy and deception

Based on the equilibrium concept given in Definition 1, we define truthful disclosure, secrecy, and deception as follows.

Definition 2. In an equilibrium $\{a^*(s), d^*(\theta), s^*(\theta), p'^*, p^*\}$, we say that in period t, a defender of type θ chooses:

1. *truthful disclosure* if and only if $s_t^*(\theta) = {}'d_t^*(\theta)'$;
2. *secrecy* if and only if $s_t^*(\theta) = \{S\}$; and
3. *deceptive disclosure* if and only if $s_t^*(\theta) \neq \{S\}$ and $s_t^*(\theta) \neq {}'d_t^*(\theta)'$.

When the attacker knows the defender type, and the cost of implementing truthful disclosure is lower than the costs of secrecy and deception, then Zhuang and Bier (in press) suggests that truthful disclosure will always be preferred to secrecy or deception. An analogous condition for the case with private defender information is

$$g_D(d_t, {}'d_t', \theta) \leqslant g_D(d_t, s_t, \theta) \ \forall t, d_t, s_t. \tag{10}$$

When inequality (10) is satisfied for a defender of type θ, then the cost of implementing truthful disclosure is lower than the costs of implementing secrecy and deception, respectively. We focus on the case when (10) is satisfied, since otherwise it would be trivial to find an equilibrium in which a defender would prefer secrecy or deception, for the simple reason that truthful disclosure is costly.

The numerical examples in the rest of this paper show that in games with private defender information, defenders may strictly prefer secrecy or deception at equilibrium even if inequality (10) is satisfied (that is, even if truthful disclosure is less costly to implement than secrecy or deception, as seems plausible). For simplicity, Section 3 studies the case when the attacker observes the defensive investment from the previous period at the beginning of the next period. Section 4 relaxes that assumption.

3. Attacker observes defensive investment from the previous period

The model presented in Section 2 above does not appear to be analytically solvable in a straightforward or tractable manner. In this section, we therefore provide a backward induction algorithm to solve it numerically. In particular, we solve the model under the assumption that the attacker can observe the previous period's defensive choice, d_{t-1}, at the beginning of period t, as in Coleb and Kocherlakotad (2001). However, unlike Coleb and Kocherlakotad (2001), we still allow the defender's private information to remain secret throughout the entire game, if not revealed by the defender's choices. This assumption (that the attacker can observe the previous defenses, which is relaxed in Section 4) greatly simplifies the computation. However, with this assumption, the defender cannot choose deception or secrecy at optimality for more than one time period. Once deception or secrecy occurs in a single period, the attacker is subsequently able to recognize the defender type by comparing the signal and the observed defense, so further deception will not be beneficial after the defender type has been revealed.

For computational convenience, we assume that capital can be carried over only to the immediate next period; that is, we have $\rho_k = 0$ for $k \geqslant 2$, and $\rho_1 = \rho$. (Note that this assumption is not a major limitation of our model, and is made only for computational convenience; cases where capital investment is carried over for more than one period can be analyzed in a straightforward manner simply by expanding the state space of our model, although this significantly increases the computational burden.) Under this assumption, Eq. (3) becomes

$$e_t(\theta) = \alpha[1 - d_t(\theta)] + d_t(\theta) + \rho d_{t-1}(\theta). \tag{11}$$

Since the player payoffs in period t depend only on the previous defensive investment d_{t-1}, the attacker prior belief p_t at the beginning of period t, and the player strategies $a_t(s_t)$, $d_t(\theta)$, and $s_t(\theta)$ in period t, there must exist a Markovian optimal strategy (see Strauch, 1966; Puterman, 1994). That is, if we let the "state" of the system consist of the variables d_{t-1} and p_t, the current-period payoff depends only on the current state and the current strategies. If we model the case where the capital defense can be carried over to more than one period, this property still holds, provided that we have a suitable state space (assuming that the attacker knows all of the previous-period defenses). The backward algorithm (Puterman, 1994) is therefore guaranteed to find an optimal equilibrium strategy (if one exists).

One element of the state of the system in this model is the attacker's belief about the defender type, p_t, at the beginning of period t. In order to use the backward algorithm to solve for the optimal strategy in period $t-1$, we need to calculate the payoff for each possible value of p_t. Note that although the attacker updates his beliefs about the defender type, p_t', in a discrete manner after observing the defender's signal ($p_t' = 0, p_t$, or 1), the update

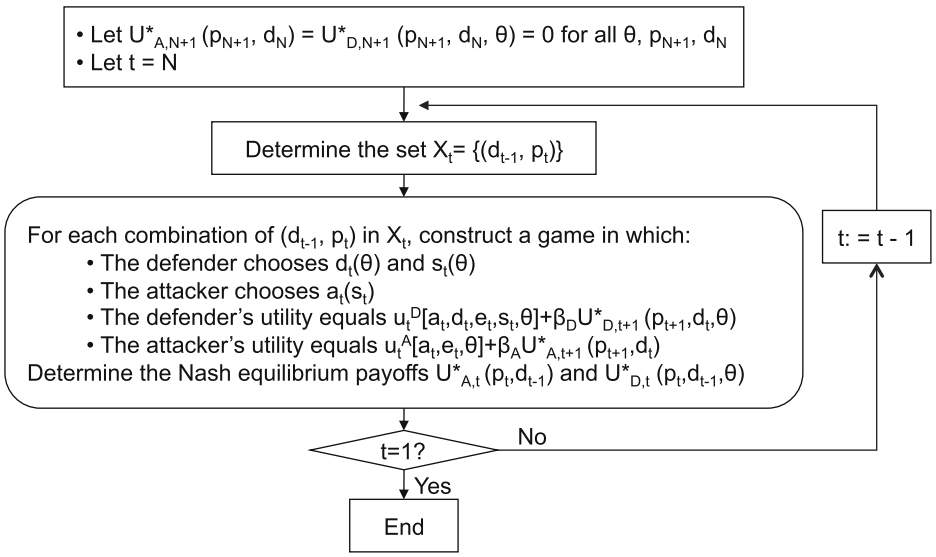

Fig. 2. Overview of the algorithm when attacker observes defensive investment.

after observing the contest outcome w_t is continuous. Therefore, for reasons of computational tractability, we discretize the resulting probability p_{t+1}. In particular, for a step size given by $0 < \delta << 1$ where $\frac{1}{\delta}$ is an integer, we approximate p_t by $\lfloor \frac{p_t}{\delta} \rfloor \delta$, where the floor function $\lfloor x \rfloor$ returns the largest integer less than or equal to x. We use $\delta = 0.01$ in the numerical examples of this section.

3.1. Overview of the algorithm

Fig. 2 provides an overview of our backward algorithm for the case when the attacker observes the previous-period defensive investment. This algorithm consists of four steps, as specified below:

1. Set the optimal attacker and defender payoffs earned at the end of the game (denoted as period $N+1$), $U^*_{A,N+1}(p_{N+1}, d_N)$ and $U^*_{D,N+1}(p_{N+1}, d_N, \theta)$, to zero for all possible states (i.e., all possible values of p_{N+1} and d_N) and for both defender types θ. (Recall that the states of the game in period t are the attacker's beliefs p_t and the previous defense d_{t-1}, and that p_1 and d_0 are assumed to be common knowledge.) Then set the current period to be $t = N$.
2. Create a set X_t whose elements consist of pairs of the form (d_{t-1}, p_t). In particular, when $t = 1$, define $X_t = \{(d_0, p_1)\}$. When $t > 1$, let the set X_t consist of all possible combinations of the defensive investment $d_{t-1} = 0, 1$ and the (discretized) probability p_t. Thus the set of possible probabilities is either $\{0, \delta, 2\delta, \ldots, 1-\delta, 1\}$ (if the contest result depends on the defender type θ) or $\{0, p_1, 1\}$ otherwise.
3. For each pair (d_{t-1}, p_t) in the set X_t, construct a one-period game in period t in which the player strategies include the attacker response $a_t(s_t)$, the defense $d_t(\theta)$, and the defender signal $s_t(\theta)$ for all $s_t = 0, 1, S$ and $\theta = \theta_1, \theta_2$. These player strategies will (stochastically) determine the states of the model (d_t and p_{t+1}) in the next period $t+1$. The (stochastic) attacker and defender payoffs of this one-period game include the current payoffs $u_{A,t}$ and $u_{D,t}$, as well as the optimal discounted future payoffs $\beta_A U^*_{A,t+1}(p_{t+1}, d_t)$ and $\beta_D U^*_{D,t+1}(p_{t+1}, d_t, \theta)$, respectively. Then calculate the Nash equilibrium of this one-period normal-form game, and let the equilibrium player payoffs equal $U^*_{A,t}(p_t, d_{t-1})$ and $U^*_{D,t}(p_t, d_{t-1}, \theta)$ for $\theta = \theta_1, \theta_2$. (Section 3.2 describes how to construct the strategies and payoffs of this game in more detail.)
4. If $t = 1$, end the algorithm and retrieve the optimal equilibrium path (if one exists). Otherwise, let t equal $t-1$, and go back to Step 2.

3.2. Solving the one-period games in period t

In this section, we explain how to implement Step 3 in the above algorithm; i.e., how to construct and solve the game in period t corresponding to any possible pair of the previous-period defense d_{t-1} and the attacker belief p_t. Depending on whether the attacker already knows the defender type at the beginning of period t, there are two possibilities:

Case A ($p_t = 0$ **or** $p_t = 1$): In this case, at the beginning of period t, the attacker already knows whether the defender is of type $\theta = \theta_2$ or $\theta = \theta_1$ (corresponding to $p_t = 0$ or $p_t = 1$, respectively). Therefore, we have an 8×6 game between the attacker and one known defender. The attacker has eight possible strategies; for each of the three possible signals $s_t = 0, 1, S$, he can respond with one of the two possible actions, $a_t(s_t) = 0, 1$. The known defender of type θ has six possible strategies; she can choose one of the two possible defense levels, $d_t(\theta) = 0, 1$, and one of three possible signals, $s_t(\theta) = 0, 1, 2$. See Table 1 for all possible attacker and defender strategies.

For all 48 cases, we calculate e_t using Eq. (11), and let $p'_t = p_{t+1} = p_t$. The attacker and defender total expected payoffs are calculated as the sum of the current payoff plus the discounted future equilibrium payoff:

$$u^A_t[a_t(s_t), e_t, \theta] + \beta_A U^*_{A,t+1}[p_{t+1}, d_t] \tag{12}$$

and

$$u^D_t[a_t(s_t), d_t, e_t, s_t, \theta] + \beta_D U^*_{D,t+1}[p_{t+1}, d_t, \theta]. \tag{13}$$

Table 1
Attacker and defender strategies at period t.

Eight attacker strategies			Six defender strategies	
$a_t(\text{'0'}) = 0$	$a_t(\text{'1'}) = 0$	$a_t(S) = 0$	$d_t(\theta) = 0$	$s_t(\theta) = \text{'0'}$
$a_t(\text{'0'}) = 0$	$a_t(\text{'1'}) = 0$	$a_t(S) = 1$	$d_t(\theta) = 0$	$s_t(\theta) = \text{'1'}$
$a_t(\text{'0'}) = 0$	$a_t(\text{'1'}) = 1$	$a_t(S) = 0$	$d_t(\theta) = 0$	$s_t(\theta) = S$
$a_t(\text{'0'}) = 0$	$a_t(\text{'1'}) = 1$	$a_t(S) = 1$	$d_t(\theta) = 1$	$s_t(\theta) = \text{'0'}$
$a_t(\text{'0'}) = 1$	$a_t(\text{'1'}) = 0$	$a_t(S) = 0$	$d_t(\theta) = 1$	$s_t(\theta) = \text{'1'}$
$a_t(\text{'0'}) = 1$	$a_t(\text{'1'}) = 0$	$a_t(S) = 1$	$d_t(\theta) = 1$	$s_t(\theta) = S$
$a_t(\text{'0'}) = 1$	$a_t(\text{'1'}) = 1$	$a_t(S) = 0$		
$a_t(\text{'0'}) = 1$	$a_t(\text{'1'}) = 1$	$a_t(S) = 1$		

We now determine the Nash equilibrium of the above two-player game. In this case, $U^*_{A,t}(p_t, d_{t-1})$ is the attacker's equilibrium payoff. If $p_t = 0$, then $U^*_{D,t}(p_t, d_{t-1}, \theta_2)$ is the defender's equilibrium payoff, and we set $U^*_{D,t}(p_t, d_{t-1}, \theta_1)$ to a large negative number. Otherwise, if $p_t = 1$, then $U^*_{D,t}(p_t, d_{t-1}, \theta_1)$ is the defender's equilibrium payoff, and we set $U^*_{D,t}(p_t, d_{t-1}, \theta_2)$ to a large negative number.

Note that multiple equilibria will typically occur in tandem with multiple possible off-equilibrium signals. For example, if $a^*_t(s_t)$ and $s^*_t(\theta)$ are part of an equilibrium in which $s^*_t(\theta) = 0$, then there may exist up to four equilibria in which $a^*_t(1) = 0$ or 1 and $a^*_t(S) = 0$ or 1, since the attacker's responses to off-equilibrium signals (1 or S in this example) will not directly change the player equilibrium payoffs. If there are multiple equilibria, we first find all equilibria, then discard those that have less than the maximum total social payoffs, and finally choose any of the remaining equilibria. Note, however, that equilibria with the maximum total social payoffs may not always be realistic, since this is an adversarial rather than a cooperative game. If no equilibrium exists for some values of p_t and d_{t-1}, then we set the corresponding player payoffs $U^*_{A,t}(p_t, d_{t-1})$ and $U^*_{D,t}(p_t, d_{t-1}, \theta)$ to be highly negative, to ensure that the corresponding strategies will not be accessed as part of an optimal equilibrium path from the previous period.

Case B ($0 < p_t < 1$): In this case, at the beginning of period t, the attacker is uncertain about the defender type, and we have a three-player, $8 \times 6 \times 6$ game. As in Case A, the attacker has eight possible strategies, and each defender type has six possible strategies (as specified in Table 1 above). For all 288 cases, we calculate $e_t(\theta)$ using Eq. (11), and then determine p'_t stochastically as a function of $s_t(\theta_1)$, $s_t(\theta_2)$, and p_t, using condition 3 of Definition 1.

We next calculate the attacker and defender total expected payoffs (the current payoff, plus the discounted expected future equilibrium payoff). In particular, the attacker payoff is given by:

$$\mathbb{E}_{p'_t}[p'_t u^A_t[a_t(\theta_1), s_t(\theta_1), e_t(\theta_1), \theta_1] + (1 - p'_t)u^A_t[a_t(s_t(\theta_2)), e_t(\theta_2), \theta_2] + \beta_A \mathbb{E}_{w_t} U^*_{A,t+1}(p_{t+1}(w_t, p'_t), d_t)] \quad (14)$$

and the payoff to a defender of type θ is given by

$$\mathbb{E}_{p'_t}[u^D_t[a_t(\theta), d_t(\theta), e_t(\theta), s_t(\theta), \theta] + \beta_D \mathbb{E}_{w_t} U^*_{D,t+1}(p_{t+1}(w_t, p'_t), d_t, \theta)], \quad (15)$$

where the first and second expectations are taken with respect to the distribution of p'_t and w_t, respectively. We calculate $\Pr\{w_t = 1\}$ at period t for given values of $a_t(s_t), e_t(\theta)$, and p'_t by conditioning on the possible defender types:

$$\begin{aligned}\Pr\{w_t = 1\} &= \Pr\{w_t = 1|\theta_1\}p'_t + \Pr\{w_t = 1|\theta_2\}(1 - p'_t) \\ &= c[a_t(s_t(\theta_1)), e_t(\theta_1)]p'_t + c[a_t(s_t(\theta_2)), e_t(\theta_2)](1 - p'_t).\end{aligned}$$

Finally, to determine the Nash equilibrium of the above three-player game, we let $U^*_{A,t}[p_t, d_{t-1}(\theta)]$ be the equilibrium attacker payoff, and $U^*_{D,t}[p_t, d_{t-1}(\theta), \theta]$ be the equilibrium defender payoff for a defender of type θ. Cases with multiple equilibria or no equilibrium are handled as in Case A above.

3.3. Baseline parameter values

In the examples in the following sections, we use the following baseline parameter values: $N = 2; p_1 = 0.9; \beta_A = 0.9; \beta_D(\theta_1) = \beta_D(\theta_2) = 0.9; \rho(\theta_1) = \rho(\theta_2) = 0.5; \alpha(\theta_1) = \alpha(\theta_2) = 2; v_A(\theta_1) = v_A(\theta_2) = 20; v_D(\theta_1) = v_D(\theta_2) = 20$. Moreover, we use the following baseline costs:

	$g_D(d, s, \theta_1)$			$g_D(d, s, \theta_2)$		
	$s =$ '0'	$s =$ '1'	$s = S$	$s =$ '0'	$s =$ '1'	$s = S$
$d = 0$	0	2	1	0	2	1
$d = 1$	2	0	1	2	0	1

Note that these values ensure that implementing truthful disclosure is less costly than implementing deception or secrecy. Also, in practice, defenders must usually spend more resources to implement deception than secrecy, since a successful deception requires both keeping the truth secret, and publicizing the deceptive information. We have therefore chosen the cost of secrecy to be less than the cost of deception.

Sections 3.4 and 3.5 show that secrecy can be an equilibrium strategy when the attacker is uncertain about the defender expense effectiveness and the asset valuation. In those cases, the stronger (or less valuable) defender type uses secrecy to differentiate herself from the weaker (or more valuable) defender type, in order to deter (or disinterest) the attacker. Section 3.6 shows that deception can also be an equilibrium strategy when the attacker is uncertain about the defender costs; however, in this case, we find deception used to mimic the other defender type.

Table 2
Equilibrium output when the expense effectiveness α is defender private information.

	p_t	d_{t-1}	$d_t(\theta_1)$	$s_t(\theta_1)$	$d_t(\theta_2)$	$s_t(\theta_2)$	a_t('0')	a_t('1')	$a_t(S)$	Remark
$g_A = 2$:										
$t = 1$	0.90	0	0	0	0	0	1	1	1	–
$w_1 = 0 : t = 2$	0.88	0	0	0	0	0	1	1	1	–
$w_1 = 1 : t = 2$	0.93	0	0	0	0	0	1	1	1	–
$g_A = 4$:										
$t = 1$	0.90	0	0	0	0	S	1	1	1	θ_2 Secrecy
$t = 2$	1.00	0	0	0	–	–	1	1	1	–
$t = 2$	0.00	0	–	–	0	0	0	0	0	–
$g_A = 6$:	*No pure-strategy equilibrium exists. The results below show a mixed strategy in period 1*									
$t = 1$	0.90	0	0	$P\{0\} = \frac{2}{3}$, $P\{S\} = \frac{1}{3}$	0	S	1	1	$P\{0\} = \frac{3}{20}$, $P\{1\} = \frac{17}{20}$	Secrecy
$s_1 = 0 : t = 2$	1.00	0	0	0	–	–	1	1	1	–
$s_1 = S, w_1 = 0 : t = 2$	0.72	0	0	0	0	0	0	0	0	–
$s_1 = S, w_1 = 1 : t = 2$	0.83	0	0	0	0	0	1	1	1	–
$g_A = 8$:										
$t = 1$	0.90	0	0	0	0	0	0	0	0	–
$t = 2$	0.90	0	0	0	0	0	0	0	0	–

3.4. Effectiveness of expenses as defender private information

When the defender private information θ involves parameters affecting the contest result (such as α or ρ_k), the attacker may be able to update his beliefs in a continuous manner. Here, we let $\alpha(\theta_1) = 2$ and $\alpha(\theta_2) = 4$ be the defender private information. Table 2 gives results for attack costs $g_A = 2$, 4, 6, and 8.

Case of $g_A = 2$: In this case, the attack cost is relatively low, so that at optimality, the attacker chooses to attack for any observed signal in all periods. Therefore, there is no benefit to deception, since the attacker response is independent of the signals, and a contest occurs in each period. In period 1, both defender types choose to invest in short-term expenses, and choose truthful disclosure. After observing the result of the contest, the attacker updates his knowledge about the defender type in a continuous manner. In particular, when the attack fails ($w = 0$), the attacker concludes that the defender is more likely to be of the strong type (in this case, type θ_2), resulting in $p_2 = 0.88$ at the beginning of period 2 (in other words, the probability that the defender is of type θ_1 has decreased from $p_1 = 0.9$ in period 1 to $p_2 = 0.88$ in period 2). By contrast, when the attack succeeds ($w = 1$), the attacker concludes that the defender is more likely to be of the weak type (in this case, type θ_1), resulting in $p_2 = 0.93$ at the beginning of period 2.

Case of $g_A = 4$: In this case, the attack cost is intermediate, so the attacker would be deterred by investment in expenses by the defender of type θ_2 (if the defender type were known), but not by the (weaker) defender of type θ_1. We observe secrecy by the defender of type θ_2 in the first period. The purpose of this secrecy is for the defender of type θ_2 to distinguish herself from the defender of type θ_1, and thereby benefit from attack deterrence in future periods. Note also that when the discount factor of the defender of type θ_2 is sufficiently small (so that she places little value on future attack deterrence), we no longer obtain secrecy in this case.

Case of $g_A = 6$: In this case, the attack cost is moderately high and there exists no pure-strategy equilibrium. This is because the weak defender type always wants to mimic the stronger defender type in the first period, in order to deter the attacker, while the stronger defender type wants to differentiate herself from the weaker type.

Although there exists no pure-strategy equilibrium in this case, we have a mixed-strategy equilibrium, as follows: In period 1, both defender types choose to invest in expenses, $d_1(\theta_1) = d_1(\theta_2) = 0$. The strong defender of type θ_2 chooses secrecy (as in the case where $g_A = 4$) with probability one. The weak defender of type θ_1 chooses a truthful signal of having invested in expenses with probability $\frac{2}{3}$, and a signal of secrecy with probability $\frac{1}{3}$, in order to decrease the chance of being attacked by (probabilistically) mimicking the stronger (secretive) defender. The attacker will respond to a signal of expense defense by attacking, since he will know that this signal can come only from a weak defender; he will respond to a signal of secrecy by attacking with probability $\frac{17}{20}$, and not attacking with probability $\frac{13}{20}$, respectively. Note that when observing a signal of secrecy, the attacker is indifferent about whether to attack, since he cannot determine whether the secrecy is coming from a weak or a strong defender. Given the attacker's mixing probabilities at equilibrium, the weaker defender is also indifferent between secrecy (which is costly, but achieves attack deterrence) or a truthful signal (which avoids any signaling cost, but results in an attack).

Case of $g_A = 8$: In this case, the attack cost is sufficiently high that the attacker will not attack regardless of the nature of the defender's investment (capital or expense). Therefore, we do not find deception or secrecy at equilibrium, and both defender types choose to invest in expenses.

Table 3
Equilibrium output when the target valuation is defender private information.

	p_t	d_{t-1}	$d_t(\theta_1)$	$s_t(\theta_1)$	$d_t(\theta_2)$	$s_t(\theta_2)$	a_t('0')	a_t('1')	$a_t(S)$	Remark
$g_A = 2$:										
$t = 1$	0.90	0	0	0	0	0	1	1	1	–
$w_1 = 0 : t = 2$	0.90	0	0	0	0	0	1	1	1	–
$w_1 = 1 : t = 2$	0.90	0	0	0	0	0	1	1	1	–
$g_A = 4$:										
$t = 1$	0.90	0	0	S	0	0	1	1	1	θ_1 Secrecy
$t = 2$	1.00	0	0	0	–	–	0	0	0	–
$t = 2$	0.00	0	–	–	0	0	1	1	1	–
$g_A = 6$:										
$t = 1$	0.90	0	0	0	0	0	0	0	0	–
$t = 2$	0.90	0	0	0	0	0	0	0	0	–

Table 4
Equilibrium output when the cost is defender private information.

	p_t	d_{t-1}	$d_t(\theta_1)$	$s_t(\theta_1)$	$d_t(\theta_2)$	$s_t(\theta_2)$	a_t('0')	a_t('1')	$a_t(S)$	Remark
$g_A = 4$:										
$t = 1$	0.90	0	0	0	1	1	1	1	1	–
$t = 2$	1.00	0	0	0	–	–	1	1	1	–
$t = 2$	0.00	1	–	–	1	1	1	1	1	–
$g_A = 6$:										
$t = 1$	0.90	0	1	1	1	1	1	1	1	–
$w_1 = 0 : t = 2$	0.90	1	0	0	1	0	0	1	1	θ_2 Deception
$w_1 = 1 : t = 2$	0.90	1	0	0	1	0	0	1	1	θ_2 Deception
$g_A = 8$:										
$t = 1$	0.90	0	0	0	1	0	0	1	1	θ_2 Deception
$t = 2$	1.00	0	0	0	–	–	0	0	0	–
$t = 2$	0.00	1	–	–	1	1	0	0	0	–
$g_A = 10$:										
$t = 1$	0.90	0	0	0	1	1	0	0	0	–
$t = 2$	1.00	0	0	0	–	–	0	0	0	–
$t = 2$	0.00	1	–	–	1	1	0	0	0	–

3.5. Target valuation as private information

In this subsection, we consider $\alpha(\theta_1) = \alpha(\theta_2) = 1.5$, $v_A(\theta_1) = v_D(\theta_1) = 10$ and $v_A(\theta_2) = v_D(\theta_2) = 20$, and inherit all other parameter values from the previous subsection. Therefore, defenders of type θ_1 differ from defenders of type θ_2 only with regard to their target valuations. Table 3 shows the results for attack costs of $g_A = 2$, 4, and 6.

Case of $g_A = 2$: In this case, the attack cost is relatively low, so that at optimality, the attacker chooses to attack for any observed signal in all periods. Therefore, there is no benefit to deception, since the attacker response is independent of the signals. A contest happens in each period. In both periods 1 and 2, both defender types choose to invest in short-term expenses, and choose truthful disclosure. However, unlike the result in Table 2, after observing the result of the contest, the attacker is still not able to update his knowledge about the defender type, and therefore we have $p_2 = p_1 = 0.9$. This is because the target valuation, which is the defender's private information in this case, has no direct impact on the success probability of an attack (and is assumed not to be observable by the attacker even after a successful attack).

Case of $g_A = 4$: In this case, the attack cost is intermediate. The attack cannot be deterred in period 1, and both defender types choose to invest in short-term expenses. However, we observe secrecy by the defender of type θ_1, to differentiate herself from the more valuable defender of type θ_2, so that the attacker will not be interested in attacking her in period 2 (since her asset valuation is relatively low).

Case of $g_A = 6$: In this case, the attack cost is high, and the attacker will not attack regardless of whether the defender invests in capital or expense. Therefore, we do not find deception or secrecy at equilibrium, and both defender types choose to invest in short-term expenses.

3.6. Defender costs as private information

In this subsection, we consider $\alpha(\theta_1) = \alpha(\theta_2) = 2$. We let the cost be the defender's private information. As shown below, the defender of type θ_2 has higher costs for all signals than the defender of type θ_1 when the defenses are given by $d = 0$.

	$g_D(d,s,\theta_1)$			$g_D(d,s,\theta_2)$		
	$s =$ '0'	$s =$ '1'	$s = S$	$s =$ '0'	$s =$ '1'	$s = S$
$d = 0$	0	2	1	4	6	5
$d = 1$	2	0	1	2	0	1

The differences between costs for defenders of type θ_1 and θ_2 could be justified by opportunity costs or political costs. For example, the constituents or stakeholders of the defender of type θ_2 may dislike short-term expenses. Table 4 gives the results for attack costs $g_A = 4, 6, 8,$ and 10.

Case of $g_A = 4$: In this case, the attack cost is relatively low, so that at equilibrium, the attacker chooses to attack for any observed signal in all periods, and a contest occurs in each period. Therefore, there is no benefit of deception, since the attacker response is independent of the signals. The defender of type θ_2 chooses long-term defense, because it is cheaper for her; the defender of type θ_1 chooses short-term expenses, for the same reason. Since we assume that the attacker observes the previous-period defense, at the beginning of period 2 the attacker recognizes the defender type with certainty (either $p_2 = 1$ or $p_2 = 0$), as reflected in the two rows for $t = 2$ in Table 4.

Case of $g_A = 6$: In this case, the attack cost is intermediate. The attack cannot be deterred in period 1, and both defender types choose capital defenses (which are partially carried over to period 2). However, unlike in Table 2, after observing the result of the contest, the attacker is still not able to update his knowledge about the defender type ($p_2 = 0.9$). This is because the cost, which is the defender's private information in this case, has no direct impact on the success of an attack. In period 2, the attacker is deterred by expense defenses, but not by capital defenses. We observe deception by the defender of type θ_2, to mimic the defender of type θ_1. This is because the cost of truthfully disclosing an expense defense is higher for the defender of type θ_2 than the cost of choosing a capital defense but signaling that expense was chosen.

Table 5
Equilibrium output when the attacker does not observe defensive investment and the cost is defender private information.

	$d_t(\theta_1)$	$s_t(\theta_1)$	$d_t(\theta_2)$	$s_t(\theta_2)$	a_t('0')	a_t('1')	$a_t(S)$	Remark
$g_A = 6$:								
$t = 1$	0	0	1	1	1	1	1	–
$t = 2$	0	0	1	1	1	1	1	–
$g_A = 8$:								
$t = 1$	0	0	1	0	0	1	1	θ_2 Deception
$t = 2$	0	0	1	1	1	1	1	–
$g_A = 9$:								
$t = 1$	0	0	1	0	0	1	1	θ_2 Deception
$t = 2$	0	0	1	0	0	1	1	θ_2 Deception
$g_A = 10$:								
$t = 1$	0	0	1	1	0	0	0	–
$t = 2$	0	0	1	1	0	0	0	–

Case of $g_A = 8$: In this case, the attack cost is relatively high. The attacker can be deterred by expense defenses in each period. Since the expense defense is costly to the defender of type θ_2, she chooses deception to mimic the defender of type θ_1 in period 1, which successfully deters the attacker. This reveals her type at the beginning of period 2 (and therefore we have either $p_2 = 1$ or $p_2 = 0$, as reflected by the two rows for $t = 2$ in corresponding section of Table 4). Benefiting from the carried over defense in period 1, the defender of type θ_2 is then able to deter the attacker in period 2.

Case of $g_A = 10$: In this case, the attack cost is high, and the attacker will not attack regardless of whether the defensive investment is capital or expense. Therefore, we do not find either deception or secrecy at equilibrium, and defenders of type θ_1 and θ_2 choose to invest in expenses and capital, respectively.

3.7. Other parameters as defender private information

In cases where the defender's private information is associated only with future payoffs (such as the carry-over coefficients ρ_k and the discount rate β_D), we have not found deception or secrecy in our numerical model, despite an extensive computer search. Further, we speculate that deception and secrecy will not be equilibrium strategies in such cases, because we have assumed that the attacker always observes the defenses from the previous period; deception or secrecy seem unlikely to be beneficial if the defender's private information affects only the defender's future payoffs.

4. Attacker does not observe defensive investment

In this section, we consider the case where the attacker does not learn the previous-period defense d_{t-1} at the beginning of period t. For simplicity, we also assume that the attacker does not observe the result of the contest from the previous period. In this case, the strategies of the game must consist of the defenses in all periods. Therefore, we need to solve a three-player $8^N \times 6^N \times 6^N$ game, where N is the number of periods. The attacker has 8^N possible strategies; for each period $t = 1, \ldots, N$, and for each of the three

possible signals $s_t = 0, 1, 2$, he can respond with one of the two possible actions $a_t(s_t) = 0, 1$. Each defender type has 6^N possible strategies; for each period $t = 1, \ldots, N$, she can choose one of the two possible defense levels $d_t(\theta) = 0, 1$, and one of three possible signals, $s_t(\theta) = 0, 1, 2$.

For each of the $8^N \times 6^N \times 6^N$ attacker–defender strategies $\{a_t(s_t), d_t(\theta), s_t(\theta)\}_{t=1}^N = \{a, d, s\}$, we calculate the beliefs $p_t'^*$ using condition 3 in Definition 1, and calculate p_{t+1}^* using Eq. (9), recursively for each period t. Then we calculate the player payoffs $U_A(a, d, p'^*)$ and $U_D(a, d, s, \theta)$ using Eqs. (5) and (6). Finally, we identify the Nash equilibrium of the above three-player $8^N \times 6^N \times 6^N$ game. Cases with multiple equilibria or no equilibrium are handled as in Section 3. Note that the computational demands of this model increase exponentially in N, so we consider only $N = 2, 3$.

Using the same parameter values as in Section 3.6, we get the equilibrium outcomes specified in Table 5, showing that deception can be an equilibrium strategy when the attacker is uncertain about the defender cost. As in Section 3.6, deception is used here only to mimic the other defender type. However, it is noteworthy that we are able to get deception over more than one period in this case.

Case of $g_A = 6$: In this case, the attack cost is relatively low, so that at equilibrium, the attacker chooses to attack for any observed signal in all periods, and a contest occurs in each period. Therefore, there is no benefit of deception, since the attacker response is independent of the signals.

Case of $g_A = 8$: In this case, the attack cost is intermediate. In period 1, the defender of type θ_1 chooses short-term expenses, which deters the attacker. The defender of type θ_2 chooses long-term defense, but deceptively discloses a short-term expenses in order to deter the attacker. The defender of type θ_1 chooses short-term expenses because it is more effective, and the defender of type θ_2 chooses long-term defense because short-term expenses is more costly for that defender type.

Case of $g_A = 9$: In this case, the attack cost is relatively high. The attacker can be deterred by an expense defense in each period. However, since the expense defense is costly to the defender of type θ_2, she chooses deception in both periods in order to mimic the defender of type θ_1, which successfully deters the attacker.

Case of $g_A = 10$: In this case, the attack cost is high, and the attacker will not attack regardless of whether the defensive investment is capital or expense. Therefore, we do not find deception or secrecy at equilibrium, and defenders of type θ_1 and θ_2 choose to invest in expenses and capital, respectively.

5. Conclusions and future research

This work uses game theory and dynamic programming to model a multiple-period, attacker–defender, resource-allocation and signaling game with incomplete information. Our numerical examples show that defenders can sometimes achieve more cost-effective security through secrecy and deception in a multiple-period game. In particular, Sections 3.4 and 3.5 show that secrecy can be an equilibrium strategy when the attacker is uncertain about the defender expense effectiveness α and the asset valuation v, respectively. In those cases, the stronger (or less valuable) defender type uses secrecy to differentiate herself from the weaker (or more valuable) defender type, in order to deter (or disinterest) the attacker.

Section 3.6 shows that deception can also be an equilibrium strategy when the attacker is uncertain about the defender costs; however, in that case, deception is used by the defender with higher costs for short-term expenses, to mimic the other defender type when that defender chooses to invest in expenses rather than capital. In other words, the deceiver gets both the deterrence benefit of appearing to invest in short-term expenses (without the high cost of actually doing so), and also the long-term benefit of having actually invested in capital defenses.

Section 4 shows results similar to those in Section 3.6 for the case where the attacker is again uncertain about the defender costs, but now assuming that the attacker does not learn the previous-period defense d_{t-1} at the beginning of period t. With that assumption, we are able to find equilibria that involve sustained deception over more than one period.

This paper helps to fill a significant gap in the literature. In particular, to our knowledge, no past work has studied defender secrecy and deception in a multiple-period game. Moreover, we believe that the solution approach developed in this paper will prove useful in other types of multiple-period games.

One limitation to this paper is that our algorithm does not automatically identify mixed strategies. (However, we provided an example of a mixed-strategy equilibrium in one case where a pure-strategy equilibrium does not exist.) This limitation should ideally be relaxed. In fact, we anticipate that secrecy and/or deception may be observed at equilibrium for a wider range of parameter values if we allow mixed strategies (i.e., if the defender is allowed to choose secrecy and/or deception with some probability less than one).

Our multiple-period model can be used to address several other important considerations in homeland security, such as: (1) defender reputation effects (in which, for example, deception might be desirable in the short-term, but lead to loss of credibility in the long-term); and (2) attacker learning over time (e.g., through repeated attacks) regarding defender private information. A slight modification of our model could also be used to address evolving attacker and defender "technologies" (e.g., new attack strategies, or changing cost functions).

Our multiple-period model of capital versus expenses could also be extended to cases in which the effectiveness of capital defenses decays as a result of damage due to attacks. Similarly, the level of damage could be modeled using a continuous variable rather than a binary variable. Finally, our multiple-period model could be extended to incorporate multiple targets and/or multiple defensive measures, in which (for example) the defender might disclose part of her resource allocation and keep the rest secret, or disclose the total investment (perhaps over a subset of the targets) but not the detailed allocation among targets.

In cases with multiple equilibria, it would also be interesting to explore results involving objectives other than the maximum total social payoffs. While the maximum total social payoffs provide a convenient focal equilibrium in cases with multiple equilibria, this may not be a realistic equilibrium in this context, since defenders and attackers will in general have no reason to cooperate to help achieve high total payoffs.

Although we found secrecy and deception as equilibrium strategies, which is somewhat unusual in the literature, such equilibria were relatively rare and difficult to obtain in our model, compared to the frequency with which secrecy and deception are observed in practice. We suspect that this may be at least in part because of some of the more unrealistic assumptions of game theory (e.g., common knowledge, full rationality). Therefore, it may be worthwhile to develop models of optimal strategies for rational defenders when facing non-strategic (irrational and/or behaviorally realistic) players, making it possible to explore the sensitivity of optimal defender strategies to assumptions about the behavior of other players.

Moreover, once it is known that secrecy or deception may be optimal strategies, this opens up the question of whether it may sometimes be optimal for the defender to allocate defensive resources to targets that are not among those most attractive to attackers (unlike the recommendations in Bier et al., 2007). For example, if some targets can be defended more cost-effectively than the most attractive targets, then resources devoted to those

targets may not be wasted if the attacker can be misled about their attractiveness.

Although this paper studies secrecy and deception specifically in the homeland-security context, we believe that our model can also provide useful insights in other contexts, such as business competition or sustainability. For example, in a business-entry game, an established company (first mover, analogous to the defender in our model) may have some private information (such as cost structure, market information, or proprietary technology), and choose secrecy or deception rather than truthfully disclosing her actions (e.g., production and marketing plans), in order to optimally deter a possible entrant (second mover, analogous to the attacker in our model), who wants to enter the industry and compete for the same scarce market. Finally, we also believe that the solution approach adopted in this paper will prove useful in analyzing a wide variety of multiple-period games.

Acknowledgements

This research was supported by the United States Department of Homeland Security through the National Center for Risk and Economic Analysis of Terrorism Events (CREATE) under Grant Number 2007-ST-061-000001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security. We thank the editor and three anonymous reviewers for helpful comments.

References

Bier, V.M., Oliveros, S., Samuelson, L., 2007. Choosing what to protect. Journal of Public Economic Theory 9 (4), 563–587.

Brown, G., Carlyle, M., Diehl, D., Kline, J., Wood, K., 2005. A two-sided optimization for theater ballistic missile defense. Operations Research 53 (5), 263–275.

Chu, Julius, Wai Hung, Lee, Ching Chyi, 2006. Strategic information sharing in a supply chain. European Journal of Operational Research 174 (3), 1567–1579.

Coleb, H.L., Kocherlakotad, N., 2001. Dynamic games with hidden actions and hidden states. Journal of Economic Theory 98 (1), 114–126.

Crawford, V.P., 2003. Lying for strategic advantage: Rational and boundedly rational misrepresentation of intentions. American Economic Review 93 (1), 133–149.

Doepke, M., Townsend, R.M., 2006. Dynamic mechanism design with hidden income and hidden actions. Journal of Economic Theory 126 (1), 235–285.

Dresher, M., 1961. Games of Strategy—Theory and Application. Prentice-Hall, Englewood Cliffs, NJ.

Hausken, Kjell., 2008. Strategic defense and attack for series and parallel reliability systems. European Journal of Operational Research 186 (2), 856–881.

Hendricks, K., McAfee, P., 2006. Feints. Journal of Economics and Management Strategy 15 (2), 431–456.

Hespanha, J., Ateskan, Y., Kizilocak, H. 2000. Deception in non-cooperative games with partial information. In; Proc. of the Second DARPA-JFACC Symposym on Advances in Enterprise Control.

Leng, Mingming, Parlar, Mahmut, 2009. Allocation of cost savings in a three-level supply chain with demand information sharing: A cooperative-game approach. Operations Research 57 (1), 200–213.

Levitin, Gregory, Hausken, Kjell, 2009. False targets efficiency in defense strategy. European Journal of Operational Research 194 (1), 155–162.

Li, L., 2002. Information sharing in a supply chain with horizontal competition. Management Science 48 (9), 1196–1212.

Neeman, Z., 2004. The relevance of private information in mechanism design. Journal of Economic Theory 117 (1), 55–77.

Oliveros, S., 2005. Equilibrium bluffs: A model of rational feints. Working Paper, University of Wisconsin-Madison, Department of Economics.

Powell, R., 2007. Allocating defensive resources with private information about vulnerability. The American Political Science Review 101 (4), 799–809.

Prat, A., 2005. The wrong kind of transparency. American Economic Review 95 (3), 862–877.

Puterman, M.L., 1994. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley and Sons, New York, NY.

Raghunathan, S., 2003. Impact of demand correlation on the value of and incentives for information sharing in a supply chain. European Journal of Operational Research 146 (3), 634–649.

Skaperdas, S., 1996. Contest success functions. Economic Theory 7 (2), 283–290.

Strauch, R.E., 1966. Negative dynamic programming. Annals of Mathematical Statistics 37 (4), 871–890.

Yao, Z., Leung, Stephen C.H., Lai, K.K., 2008. Manufacturer's revenue-sharing contract and retail competition. European Journal of Operational Research 186 (2), 637–651.

Zhang, H., Zenios, S., 2008. A Dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation. Operations Research 56 (3), 681–696.

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

Zhuang, J., Bier, V.M., 2009. Secrecy and deception at equilibrium, with applications to anti-terrorism resource allocation. Defence and Peace Economics.