Citation for published version:

Anderson, G, McCusker, G & Pym, D 2016, A Logic for the Compliance Budget. in Q Zhu, T Alpcan, E Panaousis, M Tambe & W Casey (eds), *Proceedings, GameSec 2016- Decision and Game Theory for Security: 7th International Conference, GameSec 2016, New York, NY, USA, November 2-4, 2016*. Lecture Notes in Computer Science, vol. 9996, Springer Verlag, pp. 370-381. https://doi.org/10.1007/978-3-319-47413-7_21

DOI:

[10.1007/978-3-319-47413-7_21](https://doi.org/10.1007/978-3-319-47413-7_21)

Publication date:

2016

Document Version

Peer reviewed version

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A logic for the compliance budget

Gabrielle Anderson¹, Guy McCusker², and David Pym³

¹ University College London, UK
`gabrielle.anderson@cantab.net`

² University of Bath, UK
`g.a.mccusker@bath.ac.uk`

³ University College London, UK
`d.pym@ucl.ac.uk`

Abstract. Security breaches often arise as a result of users’ failure to comply with security policies. Such failures to comply may simply be innocent mistakes. However, there is evidence that, in some circumstances, users choose not to comply because they perceive that the security benefit of compliance is outweighed by the cost that is the impact of compliance on their abilities to complete their operational tasks. That is, they perceive security compliance as hindering their work. The ‘compliance budget’ is a concept in information security that describes how the users of an organization’s systems determine the extent to which they comply with the specified security policy. The purpose of this paper is to initiate a qualitative logical analysis of, and so provide reasoning tools for, this important concept in security economics for which quantitative analysis is difficult to establish. We set up a simple temporal logic of preferences, with a semantics given in terms of histories and sets of preferences, and explain how to use it to model and reason about the compliance budget. The key ingredients are preference update, to account for behavioural change in response to policy change, and an ability to handle uncertainty, to account for the lack of quantitative measures.

1 Introduction

The security of systems is not simply a technical problem. While encryption, robust protocols, verified code, and network defences are critical aspects of system security, the behaviour of system managers and users, and the policies that are intended to manage their behaviour, are also of critical importance.

Many security breaches are the result of users’ failure to comply with security policies. Failure to comply may simply be the result of a mistake, because of a misunderstanding, or derive from users’ being required to form an effectively impossible task.

In recent years, many effective tools for analysing security behaviour and investments have been provided by economics, beginning with significant work by Anderson and Moore [2,3], explaining the relevance of economics to information security, and Gordon and Loeb [11,12], considering optimal investment in information security. Since then, there has been a vast development in security

economics, too extensive to survey in this short paper. Game theory and decision theory have been significant parts of this development; see, for example, [1,22], and much more besides. Some of us have contributed to the use of methods from economics to assess the role of public policy in the management of information security [18] and in system management policy [15,17].

A key aspect of the management of system security policies concerns the relationship between the human users of systems and the security policies with which they are expected to comply. This relationship has been explored, in the context of security economics, by Beautelement et al. [6,7] through the concept of the *compliance budget*. The idea here is that users have a limited appetite for engagement in the behaviour that is required in order to ensure compliance with policy if that behaviour detracts from their primary operational tasks.

In Section 2, we explain the concept of the compliance budget as introduced in [6,7], building on earlier work in [8]. In Section 3, we introduce a simple temporal logic with a semantics that is based on histories of events and agents' preferences. In Section 4, we consider an example of how agents' behaviour can be understood in terms of the compliance budget and reasoned about logically. In Section 5, we consider our model of the compliance behaviour in the context of incomplete information, and briefly set out a programme of further work.

This paper is intended to be conceptual rather than technical in nature and, to this end, we deliberately employ a slightly informal style. Its purpose is to initiate a qualitative logical analysis of an important concept in security economics for which quantitative analysis is difficult to establish. We are not aware of any related work on logical analyses of the compliance budget or similar concepts.

This work was supported by UK EPSRC EP/K033042/1 and EP/K033247/1.

2 The compliance budget

Organizations' security policies are enforced using tools of different kinds, ranging from simple instructions from managers through to complex combinations of hardware, software, and tokens. For example, access control via 'something you are, something you have, and something you know'. In situations in which non-compliance with the policy is possible, most of an organization's employees will nevertheless comply provided compliance does not require additional effort.

If extra effort is required, individuals will weigh this extra effort, and the opportunity cost that it implies in terms of their production task, against the benefits they obtain from compliance. If there is good alignment (i.e., of incentives) between the individual's goals as an employee and the organization's goals, then there will be little or no conflict as the behaviour required from the individual for compliance causes no friction.

However, most individuals will tend not to choose to comply with the security behaviour required by an organization if that behaviour conflicts with the behaviour that they perceive to be required in order to achieve their own goals. In such a situation, goals are less likely to be met and effort is likely to be wasted. This benefit-cost analysis is illustrated in Figure 1.

Fig. 1. The relationship between the perceived cost to an individual of compliance and the effectiveness of a security policy [6,7]. More effective policies achieve greater effectiveness at a given cost to an individual.

Alternative rates of compliance expenditure are also shown for comparison. Once the compliance threshold is crossed security effectiveness drops sharply as employees elect to complete tasks that benefit them more directly as individuals rather than security tasks that more benefit the organization as a whole. A well-managed budget will spend perceived effort at a slower rate, so that more security policies can be accommodated before the compliance threshold is reached, resulting in a higher level of achieved security. If the limit is exceeded, security policies are less likely to be followed; achieved levels of security will decline.

Following [6,7], we remark that, in the absence of quantitative data, the precise shape of the graph in Figure 1 cannot be plotted precisely. Moreover, there will be variations from individual to individual, although the same core features will occur. These behaviours have been investigated in extensive empirical studies [8,6,7], supporting the formulation of the concept of *the compliance budget*.

3 A logic for the compliance budget

In this section, we introduce a (multi-modal) temporal logic of preferences with which we can reason about the compliance budget. For convenience, we name the logic CBL, for ‘compliance budget logic’.

The logic includes temporal modalities, modalities for agents’ preferences, and a modality for preference update [5]. Each modality has a specific role in our modelling of reasoning about the compliance budget.

- The temporal modalities, $\bigcirc$ (next time step) and $\mathcal{U}$ (until) are familiar from temporal logic [21] (see [14] for a tutorial exposition) and are used to reason about how properties of the system change over time.

- The modality $\Diamond_i$, together with its dual $\Box_i$, is used to reason about the preferences of the agents (or principals, or players) i in a system. It denotes decision-making capability between outcomes with different ‘worth’ to the agents in the system.
- The modality $[\Phi]$, for a finite set of formulae Φ , is reminiscent of the key modality in public announcement logic (see, e.g., [10] for a convenient summary); it is used to reason about how the preferences of agents in the system are changed by the imposition of policies by those who manage the system.

The semantics of the logic is specified in terms of history-based structures [20], and is explained below. Histories (sequences of events) can be used to represent the trace semantics of complex systems models, and can be seen as a simple version of the process-theoretic models explored in, for example, [9].

Definition 1 (Syntax of CBL). *Given a set of propositional variables P , with elements p, q , etc., the syntax of the logic CBL is defined as follows:*

$$\begin{array}{ll}
 \phi ::= p \mid \perp \mid \top \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \phi \rightarrow \phi & \text{classical propositionals} \\
 \mid \bigcirc\phi \mid \phi \mathcal{U} \phi & \text{temporal modalities} \\
 \mid \Diamond_i\phi \mid \Box_i\phi \mid [\Phi]\phi & \text{preference modalities.}
 \end{array}$$

We write formulae as ϕ, ψ , etc., and finite sets of formulae as Φ, Ψ , etc.. The existential preference modal operator for agent i is $\Diamond_i\phi$. The dual universal preference modal operator for agent i is $\Box_i\phi$. When there is only a single agent in the system, we sometimes drop the agent annotation. The temporal ‘next-time’ operator is $\bigcirc\phi$. The temporal ‘until’ operator is $\phi \mathcal{U} \psi$.

The preference update modality — which updates agents’ preferences — is written $[\Phi]\phi$ and denotes that ϕ holds when the model is updated to disregard preferences between pairs of histories that (respectively) do and do not satisfy some formula in Φ . We refer to a formula as *update-free* if it contains no uses of the preference update modality.

The compliance budget [6,7] is qualitative rather than quantitative concept, and accepts that accurate measures of the effort taken to follow a given policy, and the effort that an employee has to expend, can generally not be practically obtained. As a result, a preference update consists of a set of formulae according to which the preferences are updated without any formal notion of likelihood or probability between the different facts; that is, it is a qualitative update to preferences rather than a quantitative update. The preference-update modality $[\Phi]$ will be used to give a logical account of the behavioural changes brought about by the implementation of a policy. The set of formulae Φ represents the impact on agents’ decision-making under a new policy; we allow Φ to be a finite set rather than a single formula in order to incorporate uncertainty in this decision-making impact. This set-up will be essential to our logical description of the compliance budget.

First, we need some notation. Let $\mathcal{E}$ be the set of events (the ‘global event set of a model’) and $\mathcal{A}$ be the set of agents of a history-based model.

A history H over a set of events $\mathcal{E}$ is a possibly infinite sequence of events drawn from the set $\mathcal{E}$. ϵ denotes the empty history.

$$\begin{aligned}
& H, t \models_{\mathcal{M}} p \text{ iff } H_t \text{ is defined and } H_t \in \mathcal{V}(p) \\
& H, t \models_{\mathcal{M}} \perp \text{ never} \quad H, t \models_{\mathcal{M}} \top \text{ always} \quad H, t \models_{\mathcal{M}} \neg\phi \text{ iff } H, t \not\models_{\mathcal{M}} \phi \\
& H, t \models_{\mathcal{M}} \phi \vee \psi \text{ iff } H, t \models_{\mathcal{M}} \phi \text{ or } H, t \models_{\mathcal{M}} \psi \quad H, t \models_{\mathcal{M}} \phi \wedge \psi \text{ iff } H, t \models_{\mathcal{M}} \phi \text{ and } H, t \models_{\mathcal{M}} \psi \\
& H, t \models_{\mathcal{M}} \bigcirc \phi \text{ iff } H, t+1 \models_{\mathcal{M}} \phi \\
& H, t \models_{\mathcal{M}} \phi \mathcal{U} \psi \text{ iff there exists } k \in \mathbb{N} \text{ such that } t \leq k, H, k \models_{\mathcal{M}} \psi \\
& \quad \text{and, for all } l \in \mathbb{N}, t \leq l < k \text{ implies } H, l \models_{\mathcal{M}} \phi \\
& H, t \models_{\mathcal{M}} \Diamond_i \phi \text{ iff there exist } H' \in \mathcal{H} \text{ and } \pi \in \Pi \\
& \quad H\pi_i H', \text{ and } H', t \models_{\mathcal{M}} \phi \\
& H, t \models_{\mathcal{M}} \Box_i \phi \text{ iff for all } H' \in \mathcal{H} \text{ and all } \pi \in \Pi, \\
& \quad H\pi_i H' \text{ implies } H', t \models_{\mathcal{M}} \phi \\
& H, t \models_{\mathcal{M}} [\Phi] \phi \text{ iff } H, t \models_{\mathcal{M}[\Phi, t]} \phi
\end{aligned}$$

Fig. 2. Satisfaction relation

If a history H is of at least length $m \in \mathbb{N}$, then let $H(m)$ be the m th element of the sequence, H_m be the m -length prefix of the history. We emphasize that a history is finite using lower case. Let $h; H$ denote the concatenation of a finite history h with a (possibly infinite) history H . In this case, we say that h is a prefix of $h; H$. (Note that $\epsilon; H = H$.)

A *protocol* $\mathcal{H}$ is a set of histories closed under finite prefix.

Definition 2. A *preference relation* $\prec$ is a strict partial order on a protocol.

In a system with multiple agents, we use a different preference relation for each agent, to describe their separate interests. Such a collection of preferences is specified as a preference structure.

Definition 3 (Preference structure). A *preference structure for agents* $\mathcal{A}$ over histories $\mathcal{H}$ is given by a tuple $(\prec_1, \dots, \prec_n)$, where $\mathcal{A} = \{1, \dots, n\}$, and, for all $i \in \mathcal{A}$, $\prec_i$ is a preference relation on the protocol $\mathcal{H}$.

We write preference structures π, π' , etc., and sets of preference structures Π, Π' , etc..

We can now define models of CBL. Satisfaction and model update are defined mutually inductively.

Definition 4 (History-based preference models). A tuple $(\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \Pi)$ is a *history-based preference model (HBPM)*, or a *history-based model for short*, if $\mathcal{E}$ is a set of events, $\mathcal{A} = \{1, \dots, n\}$ is a set of agents, $\mathcal{H}$ is a protocol, $\mathcal{V}$ is a valuation function from propositions to subsets of $\mathcal{H}$ containing only finite histories, and Π is a set of preference structures for agents $\mathcal{A}$ over $\mathcal{H}$.

Models are denoted $\mathcal{M}, \mathcal{M}'$, etc.. The interpretation of the connectives and modalities is given in Figure 2, where satisfaction of a formula ϕ in history H at time t in a model $\mathcal{M}$ is written $H, t \models_{\mathcal{M}} \phi$. Note that the semantics of

preference update depends on the definition of preference-based model update, which is explained below. The necessary model update, as defined in Definition 6, requires only the strictly smaller formula ϕ .

The modality $\Diamond_i\phi$ denotes the existence of a history (trace) that is preferred by agent i in some possible preference relation and in which ϕ holds. The modality $\bigcirc\phi$ denotes that ϕ holds at the next time point. The modality $\phi\mathcal{U}\psi$ denotes that ϕ holds until some time point, at which ψ holds.

In order to reason about the impact of a policy, it is helpful to be able to modify the preferences of the principals in the logic. This can be modelled using preference updates, which can remove (but cannot add) preferences between pairs of histories. A preference update is performed using a *distinguishing formula*, ϕ . Given two histories H, H' , if $H, t \models_{\mathcal{M}} \phi$ but $H', t \not\models_{\mathcal{M}} \phi$, then we call ϕ a ‘distinguishing formula’ for (H, t) and (H', t) . In this case, preference update for agent i leads to a new preference relation $\prec'_i$ such that $H \not\prec'_i H'$. The notion of preference update in history-based models that we use in this paper was introduced in [5].

Definition 5 (Preference relation update). Let $\prec$ be a preference relation and $\mathcal{M} = (\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \Pi)$ be a history-based model. The preference relation updated according to a formula ϕ at time t , $\prec^{\phi, \mathcal{M}, t}$, is defined as

$$\prec^{\phi, \mathcal{M}, t} := \prec \setminus \{(H, H') \mid H, t \models_{\mathcal{M}} \phi \text{ and } H', t \not\models_{\mathcal{M}} \phi\},$$

Lemma 1. If $\mathcal{M} = (\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \Pi)$ is a history-based model, $\prec$ is a preference relation over histories $\mathcal{H}$, ϕ is a formula, and t is a time-point, then $\prec^{\phi, \mathcal{M}, t}$ is a preference relation over histories $\mathcal{H}$.

Proof. Establishing this amounts to checking that the given relation is transitive. Suppose $H \prec^{\phi, \mathcal{M}, t} H' \prec^{\phi, \mathcal{M}, t} H''$. If $H, t \models_{\mathcal{M}} \phi$, then $H', t \models_{\mathcal{M}} \phi$, so that $H'', t \models_{\mathcal{M}} \phi$. Therefore $H \prec^{\phi, \mathcal{M}, t} H''$.

We extend updates of preference relations pointwise to updates of preference structures. We can then use preference relation update to update a model using a finite set of distinguishing formulae.

Definition 6 (Preference-based model update). Let $\mathcal{M} = (\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \Pi)$ be a history-based preference model. The updated preference model $\mathcal{M}![\Phi, t]$ (with respect to a finite set of distinguishing non-updating formulae Φ and time-point t) is defined as $\mathcal{M}![\Phi, t] = (\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \{\pi^{\phi, \mathcal{M}, t} \mid \pi \in \Pi \text{ and } \phi \in \Phi\})$.

We represent a preference update within the logic via the $[\Phi]$ modality. Given a model $\mathcal{M}$ and a finite set of distinguishing non-updating formulae Φ , a preference update modality is satisfied by history H at time t in model $\mathcal{M}$ (i.e., $H, t \models_{\mathcal{M}} [\Phi]\psi$), if and only if ψ holds in the model updated by Φ at time-point t (i.e., $H, t \models_{\mathcal{M}![\Phi, t]} \psi$).

Proposition 1. The logic CBL as defined in Figure 2, together with the supporting definitions, is a conservative extension of the temporal fragment (classical propositionals and temporal modalities) without the preference fragment (preference modalities, $\Diamond_i\phi$, $\Box_i\phi$, $[\Phi]\phi$).

Proof. Consider that all of the satisfaction clauses for the temporal modalities are independent of the structures required to define the preference modalities.

Example 1. Suppose a set of events $\mathcal{E} = \{c, d\}$, denoting compliance and deviation from management-desired behaviour, and a set of histories of all traces over the events of at most length two, that is $\mathcal{H} = \{\epsilon, c, d, (c; c), (c; d), (d; c), (d; d)\}$.

We consider only one agent in the system, an employee (that is, $\mathcal{A} = \{1\}$). The employee always prefers to deviate; that is, $\pi = (\prec)$ is given by the transitive closure of

$$c \prec d \quad c; c \prec c; d \quad c; c \prec d; c \quad c; d \prec d; d \quad d; c \prec d; d.$$

Let p_c be a proposition that holds for a history when the last event in its sequence is c ; that is, $h \in \mathcal{V}(p_c)$ if and only if h is of the form $h'; c$. Let p_d be defined similarly.

Let $\mathcal{M} = (\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \{\pi\})$. We can use the logic CBL to say that the employee prefers to deviate from the behaviour desired by the manager at the first opportunity; that is, $(c; c), 0 \models_{\mathcal{M}} \Diamond_1 \bigcirc p_d$.

Suppose the policy-maker introduces incentives to encourage greater compliance with policy. In CBL, this is modelled as a preference update with the formula $\phi = \bigcirc p_c$. Updating the preferences using this formula results in $\prec^{\phi, \mathcal{M}, 0}$, consisting in just

$$c; c \prec^{\phi, \mathcal{M}, 0} c; d \quad d; c \prec^{\phi, \mathcal{M}, 0} d; d.$$

This update removes the employee's preference to deviate at the first opportunity, but not at later opportunities; formally, $(c; c), 0 \models_{\mathcal{M}} [!\{\phi\}] \neg \Diamond_1 \bigcirc p_d$.

To deal with the second opportunity to deviate from policy, let $\psi = \bigcirc \bigcirc p_c$. Updating the original preferences using this formula results in $\prec^{\psi, \mathcal{M}, 0}$, given by

$$c \prec^{\psi, \mathcal{M}, 0} d \quad c; c \prec^{\psi, \mathcal{M}, 0} d; c \quad c; d \prec^{\psi, \mathcal{M}, 0} d; d.$$

This update removes the employee's preference to deviate at the second opportunity, but not at other opportunities; formally, $(c; c), 0 \models_{\mathcal{M}} [!\{\psi\}] \neg \Diamond_1 \bigcirc \bigcirc p_d$.

In some situations, the policy-maker may have less fine-grained control over the employees. For example, they can prevent one deviation, but have no control over which deviation is prevented. This is represented by updating the preferences using the set of formulae $\Phi = \{\phi, \psi\}$, resulting in the two possible preference relations above; that is, $\mathcal{M}_{[\Phi, 0]} = (\mathcal{E}, \mathcal{A}, \mathcal{H}, \mathcal{V}, \{\prec^{\phi, \mathcal{M}, 0}, \prec^{\psi, \mathcal{M}, 0}\})$. This update removes the employee's preference to deviate twice. However, there is now uncertainty about the preferences, and properties that hold for updates according to ϕ and ψ no longer hold. Indeed,

$$(c; c), 0 \models_{\mathcal{M}} [!\Phi] \Diamond_1 \bigcirc p_d \quad \text{and} \quad (c; c), 0 \models_{\mathcal{M}} [!\Phi] \Diamond_1 \bigcirc \bigcirc p_d.$$

We do, however, have the weaker property, that the employee does not prefer to deviate at both opportunities; formally,

$$(c; c), 0 \models_{\mathcal{M}} [!\Phi] \neg \Diamond_1 ((\bigcirc p_d) \wedge (\bigcirc \bigcirc p_d)).$$

To see this, note that the only histories preferable to $c; c$ are $c; d$, from the update for ϕ , and $d; c$, from the update for ψ , and $(c; d), 0 \not\models_{\mathcal{M}_{[\Phi, 0]}} ((\bigcirc p_d) \wedge (\bigcirc \bigcirc p_d))$ and $(d; c), 0 \not\models_{\mathcal{M}_{[\Phi, 0]}} ((\bigcirc p_d) \wedge (\bigcirc \bigcirc p_d))$. $\square$

Building on this set-up, we now introduce a logical model of the compliance budget. To this end, we let $\text{load}_n(\phi, i)$ denote that agent i has at least n distinct situations in which it would prefer to violate the policy ϕ .

Definition 7. *Let i be an agent and ϕ be an update-free formula. The load formulae of these parameters are defined by*

$$\begin{aligned} \text{load}_0(\phi, i) &\triangleq \top \\ \text{load}_{n+1}(\phi, i) &\triangleq (\phi \wedge \Box_i \phi) \mathcal{U} (\phi \wedge \Diamond_i \neg \phi \wedge \bigcirc \text{load}_n(\phi, i)). \end{aligned}$$

Given a load formula $\text{load}_n(\phi, i)$, we refer to n as the load value, and ϕ as the policy. The intuition for this is that if $\text{load}_n(\phi, i)$ holds, agent i has complied with policy ϕ , but would have preferred to deviate on n occasions, so expending compliance budget.

If we perform a preference update according to the formula $\text{load}_n(\phi, i)$, we will remove the preference to deviate from the policy ϕ at the first n opportunities. We can represent a bound on uncertainty on an agent's compliance budget — that is, uncertainty on how much more the agent will comply with policies — by updating according to a set of load formulae with a range of load values:

$$\text{load}_m(\phi, i), \text{load}_{m+1}(\phi, i), \dots, \text{load}_{n-1}(\phi, i), \text{load}_n(\phi, i).$$

4 An access control example

We illustrate the facility to model ideas from the compliance budget using an example concerning remote access policy. We suppose a business setting, with an internal (local) network that contains some core systems resources (for example, databases, workflow tools, or email servers). This core system can be divided into high security and low security components, where high security components are more valuable or more vulnerable.

Principals using devices on the local network can access the entire core system, both high and low security. Principals using devices on remote networks can access the core system, but with certain restrictions based on the type of connection that is used to access the resources.

The system is technologically configured so that a connection using a virtual private network (VPN) can access the whole core system, but a connection using a secure shell (SSH) can only access the low-security component of the core system. This is an attempted implementation of the (currently informal) policy that the high-security component of the core system should only be remotely accessed via VPN. Principals can, however, use SSH to connect to locally owned machines, and then connect directly to the high component of the core system. Hence, the policy can be circumvented by a determined principal.

Fig. 3. Access control from a remote network

This scenario is depicted graphically in Figure 3. To model this, we consider the policy-maker M , and the employee (principal) P . We assume a set of events $\mathcal{E}$, comprising: cc_{loc} , cc_V , and cc_S , connecting to the core system via the local network, VPN, and SSH; dc_{loc} , dc_V , and dc_S , disconnecting from the core system via the local network, VPN, and SSH; cd_S and dd_S , connecting to and disconnecting from an employee-controlled device on the local network via SSH; and a_L and a_H , accessing the low- and high-security component of the core system.

The technological configuration places various constraints on the behaviour of the system, represented by the set of histories that we consider, $\mathcal{H}$. An access event e occurs within the scope of a connection event e' within a history h if and only if there exist histories h_1, h_2, h_3 such that $h = h_1; e'; h_2; e; h_3$ and h_2 does not contain any connect or disconnect events. For example, the a_L event does occur within the scope of the cc_S event in the history $cc_S; a_L$, but does not occur within the scope of cc_S event in the history $cc_S; dc_S; a_L$.

The set of histories contains all finite sequences of the events, except for those where some a_L event does not occur within the scope of a connection cc_{loc} , cc_V , or cc_S , and those where some a_H event does not occur within the scope of a connection cc_{loc} or cc_V . For example, the history $cc_S; a_H$ is *not* included in $\mathcal{H}$, but the histories $cc_V; a_H$ and $cd_S; cc_{loc}; a_H$ are included in $\mathcal{H}$. The history $cc_V; a_H$ conforms to the informal policy that the high security component of the core system should only be remotely accessed via VPN. The history $cd_S; cc_{loc}; a_H$, however, does not conform to the informal policy, but it is included in our set of histories $\mathcal{H}$ as the a_H event *does occur* within the scope of a connection cc_{loc} .

There are various costs and benefits to the employee for their different actions/events that they can choose. Working locally (on site) gives direct, secure access, but comes with the possibility, for example, of being interrupted by a colleague. Working remotely removes the possibility of being interrupted, but requires accessing the core system via some additional security mechanism. Using a VPN to connect remotely to the core system gives full, secure access, but has the costs that the VPN is harder (than SSH) to operate, is more faulty (than SSH), and removes the ability, for example, to use a printer on the remote net-

work. Using SSH to connect remotely to the core system gives secure access and is easier (than the VPN) to operate, is less faulty (than the VPN), and retains the ability to use a printer on the remote network, but has the cost that it has limited access only to the low security component of the core system.

We demonstrate how to model the imposition of a policy that explicitly guides against using SSH to access the core system. In the remainder of this section, we overload our syntax slightly and use an event e as a proposition which is satisfied by a history if and only if the last event in the history is the given event. The histories that comply with this policy are those that satisfy

$$\phi \triangleq (\text{cd}_S \rightarrow ((\neg \text{cc}_{loc}) \mathcal{U} \text{dd}_S))$$

at every time step. Note that a finer-grained policy that only prohibits the use of such connections to access high-security resources could be described similarly.

Consider a user working at a remote site, engaged in a task which requires two accesses to the high-security resources at base, with intervening access to remote-site resources that are not available when connected via VPN. As described above, we endow the user with a preference relation favouring SSH connections above VPN: $\text{cc}_V; \mathbf{a}_H; \text{dc}_V \prec \text{cd}_S; \text{cc}_{loc}; \mathbf{a}_H; \text{dc}_{loc}; \text{dd}_S$ (and similarly for longer histories containing these as subsequences). The user may complete the task with any of the following three histories:

$$\text{cc}_V; \mathbf{a}_H; \text{dc}_V; \text{cc}_V; \mathbf{a}_H; \text{dc}_V \tag{1}$$

$$\prec \text{cc}_V; \mathbf{a}_H; \text{dc}_V; \text{cd}_S; \text{cc}_{loc}; \mathbf{a}_H; \text{dc}_{loc}; \text{dd}_S \tag{2}$$

$$\prec \text{cd}_S; \text{cc}_{loc}; \mathbf{a}_H; \mathbf{a}_H; \text{dc}_{loc}; \text{dd}_S. \tag{3}$$

Consider a model $\mathcal{M}$ that embodies this scenario. To model the imposition of the access control policy, we perform preference update according to the set of formulae $\Phi \triangleq \{\text{load}_i(\phi) \mid i = 1, 2\}$, arriving at model $\mathcal{M}' \triangleq \mathcal{M}[\Phi]$. This update reflects the policy-maker's inevitable uncertainty in the compliance budget of the user. Because of the user's prior preference, compliance with policy ϕ comes at a cost. Some of this cost is directly observable: witness the disconnections and reconnections in history (1), the least-preferred, most-compliant behaviour. However, other costs may not be observed, for instance the possible failure of attempts to access resources at the remote site while connected via VPN. Not only is the policy-maker unable to judge the effort available for compliance, but also the effort required to comply is uncertain. In our model, updating preference with $\text{load}_n(\phi)$ reflects the willingness of a user to deviate from prior preference in favour of compliance up to n times. Model $\mathcal{M}'$ contains a preference structure for each n , that is, for each possible value of the (unmeasurable) compliance budget. A highly compliant user ($n = 2$ in this example) becomes indifferent between histories 1–3. A user with low compliance budget ($n = 1$) retains just the preference for history 2 over the fully-compliant 1. Thus for the highly compliant user, preference and policy are aligned, so there is no reason to violate the policy. For the less compliant user, after the first VPN connection the budget is exhausted and the user prefers to revert to SSH, contravening the policy.

The scenario we have modelled ascribes at least some propensity for compliance to the user: we do not include $\text{load}_0(\phi)$ in the set of preference update formulae. As a result, we are able to draw some conclusions about the preferences of the user under the policy ϕ . For instance, each of the histories 1–3 satisfies

$$H, 0 \models_{\mathcal{M}} [!\Phi] \Box (\text{cd}_S \rightarrow \neg(\neg \text{dd}_S \mathcal{U} \mathbf{a}_H \wedge \bigcirc(\neg \text{dd}_S \mathcal{U} \mathbf{a}_H)));$$

that is, the user would never prefer to adopt a behaviour incorporating two accesses to high-security resources via SSH.

5 Further work: incomplete information reasoning

Our model of the compliance budget has been designed to account for the fact that the ‘budget’ is not a quantifiable value, and the rate at which it is depleted is unknown, as explained in [6,7]. This has led to a model in which we have, for each agent, a set of *possible* preference relations over histories. That is, our model incorporates *uncertainty* about the preferences of the agents: we know that eventually the compliance budget will be exhausted, but we do not know how long that will take. The impact of imposing a new policy ϕ is modelled by updating the agents’ preferences with $\text{load}_n(\phi)$ for an uncertain value of n .

Uncertainty over preferences is the qualitative analogue of uncertainty over payoffs. Harsanyi [13] demonstrates that *all* uncertainty over the structure of a game can be reduced to uncertainty over payoffs. Our model is therefore a simple qualitative setting in which to study situations of incomplete information. Security policy decisions are typically incomplete information situations because of uncertainty over the compliance budget of agents. As Harsanyi’s reduction suggests, this uncertainty subsumes lack of knowledge of the consequences of compliance on productivity. In the VPN example, the policy-maker insisting on VPN for remote access is not aware of the implications for individual agents, who may have difficulty accessing local resources (e.g., network printers) while connected to a VPN. Such issues may or may not be the reason that compliance is reduced, but, in our model, it does not matter: uncertainty in the compliance budget accounts for uncertainty over the details of agent behaviour, allowing us to model behaviour at an appropriate level of abstraction.

Much work remains, including: the metatheory of the logic and the theory of load formulae (e.g., for the interaction of multiple policies); other logics, such as epistemic variants to internalize uncertainty (note that history-based semantics supports epistemic constructions [19]); decision- and game-theoretic considerations such as optimality and equilibria; consideration of richer and larger models in order to explore the value of the approach for security policy-makers.

References

1. T. Alpcan and T. Başar. Network Security: Decision and Game-Theoretic Approach. Cambridge University Press, 2010.
2. R. Anderson. Why Information Security is Hard: An Economic Perspective. *Proc. 17th Annual Computer Security Applications Conference*, 358–265, IEEE, 2001.

3. R. Anderson and T. Moore. The Economics of Information Security. *Science* 314, 2006, 610–613.
4. C. Baskent and G. McCusker. Epistemic Game Theoretical Reasoning in History Based Models. *Proc. Workshop on Strategic Reasoning*, Oxford, 2015. Manuscript: <http://www0.cs.ucl.ac.uk/staff/D.Pym/sr2015-proceedings.pdf>.
5. C. Baskent and G. McCusker. Preferences and Equilibria in History Based Models. *Proceedings of the 12th Conference on Logic and the Foundations of Game and Decision Theory*, 2016, <http://loft.epicenter.name>.
6. A. Beautement, A. Sasse, and M. Wonham. The Compliance Budget. *Proceedings of the New Security Paradigms Workshop (NSPW '08)*, 47–55, 2008. doi:10.1145/1595676.1595684. ACM, 2008.
7. A. Beautement and A. Sasse, The economics of user effort in information security. *Computer Fraud & Security* 2009 (10) 8–12. doi: 10.1016/S1361-3723(09)70127-7.
8. A. Beautement et al. Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. In *Managing Information Risk and the Economics of Security*, M. Eric Johnson (editor), Springer, 2009, 141–163.
9. M. Collinson, B. Monahan, and D. Pym. *A Discipline of Mathematical Systems Modelling*. College Publications, 2012.
10. H. van Ditmarsch, J. Halpern, W. van der Hoek, and B. Kooi, eds. *Handbook of Epistemic Logic*. College Publications, 2015.
11. L.A. Gordon and M.P. Loeb. The Economics of Information Security Investment, *ACM Transactions on Information and Systems Security* 5(4) 438–457, 2002.
12. L.A. Gordon and M.P. Loeb. *Managing Cybersecurity Resources: A Cost-Benefit Analysis*. McGraw Hill, 2006.
13. J. Harsanyi. Games with Incomplete Information Played by ‘Bayesian’ Players, Part III. *Management Science*, 14(7):486–502, 1968.
14. M. Huth and M. Ryan. *Logic in Computer Science: Modelling and Reasoning about Systems*. Cambridge University Press, 2004.
15. C. Ioannidis, D. Pym, and J. Williams. Investments and Trade-offs in the Economics of Information Security. *Proc. Financial Cryptography and Data Security 2009*. LNCS 5628:148–162, 2009.
16. C. Ioannidis, D. Pym, and J. Williams. Fixed Costs, Investment Rigidities, and Risk Aversion in Information Security: A Utility-theoretic Approach. In *Economics of Security and Privacy III*, Bruce Schneier (editor), Springer 2012. 171–192.
17. C. Ioannidis, D. Pym, and J. Williams. Information Security Trade-offs and Optimal Patching Policies. *European Journal of Operational Research* 216(2):434–444, 2012. doi:10.1016/j.ejor.2011.05.050.
18. C. Ioannidis, D. Pym, and J. Williams. Is Public Co-Ordination of Investment in Information Security Desirable? *Journal of Information Security (Special Issue on Cybersecurity Investments, L. Gordon and M. Loeb, editors)*, 7, 60–80. doi: 10.4236/jis.2016.72005.
19. E. Pacuit. Some Comments on History Based Structures. *J. of Applied Logic*, 5(4):613–624, 2007.
20. R. Parikh and R. Ramanujam. A knowledge-based semantics of messages. *Journal of Logic, Language, and Information* 12(4):453–467, 2003.
21. A. Pnueli. The temporal logic of programs. *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (FOCS)*, 1977, 46–57. doi:10.1109/SFCS.1977.32.
22. M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.