

Case Study: Predicting the Impact of a Physical Access Control Intervention

Tristan Caulfield
University College London
United Kingdom
t.caulfield@ucl.ac.uk

Simon Parkin
University College London
United Kingdom
s.parkin@ucl.ac.uk

ABSTRACT

We investigate a planned physical security intervention at a partner organisation site, to determine the potential individual cost of security upon employees when replacing a secure door with a turnstile. Systems modelling techniques are applied to model the lobby area of the site, and to guide data collection to situate the model. Managers at the site were consulted during preference elicitation to identify meaningful model parameters. Direct observation of regular employee behaviours from pre-recorded CCTV footage provided localised data: 1800 sequences of behaviour events were logged over one working day for approximately 600 employees and visitors. This included responses to security events, such as returning to the card reader or moving to a different turnstile. Model results showed that if one turnstile was implemented at the observed site, an average of 0.5 seconds would be added to individual entry times for employees, amounting to over sixty hours for the site as a whole over a year. Three turnstiles approach the time cost of a secure door.

CCS Concepts

•Security and privacy → Formal security models; Economics of security and privacy; Usability in security and privacy;

Keywords

Security Modelling; Security Interventions; Physical Security

1. INTRODUCTION

Large organisations may renew their technology as new options become available, to introduce new advances to the organisation but also as a means to positively influence the working culture [19]. Equally, the deployment of new technologies can be one way to influence the security culture of the organisation.

Security mechanisms, including security policies and technical controls, may have unintended consequences when deployed. Human factors of security research can help organisations to identify potential blockers to secure working introduced by the mechanisms themselves [1]. Non-compliance and the misuse or avoidance of security technologies can be a result of a mismatch with the way

people in the organisation work. There is then a need for organisations to prepare security *interventions* with a consideration of the *individual cost* of security for employees.

The individual cost of security – and of security compliance – in organisations can encompass both information security and physical security, e.g. creating a hard copy version of a file in order to continue working with it off-site. Here we build on prior modelling formalisations [9, 10] and related data-gathering techniques to characterise the impact of a physical security *intervention*, in this case a secured building entry system. Security interventions can be modelled, using appropriate data collected from a real environment [4].

Prior work has shown how interviews and surveys with a cross-section of employees in an organisation can identify potential *hot-spots* where the experience of security at work has the capacity for improvement [3]. For one collaborating partner organisation, the potential for unescorted visitors to enter the organisation's sites emerged as a concern (where tailgating is but one element of this). Here we focus on one of the organisation's main sites, where secured building access interventions were being considered. Specifically, this would mean replacing the main secure door in the entrance lobby with a series of turnstiles. Although the company has many thousands of staff, this particular site has approximately 500 members of site-related staff, and receives approximately 70 contractors and visitors a day. Human factors research informs how employee interactions with secured physical access systems are captured, as well as related security behaviours (for instance their response when access cards or card readers do not appear to function correctly).

A model of the lobby and employee activity within it was used to explore interactions with the existing secure door. Crucially for the accuracy of the model, access to Closed Circuit Television (CCTV) camera footage provided the means to conduct *direct observation* of activity within and beyond the entrance lobby. Unusable security systems may be misused or abandoned [13], but if security is unavoidable (as with a single means of entering a building), employees may be seen to be complying with policy, but bear any personal cost grudgingly [6] (such as effort or cognitive load).

We combine a number of exercises: preference elicitation, data collection, and model building. Both the site manager and the physical security manager at the observed site were consulted, to better understand regular behaviour and how to parametrise the model. By visiting another site where a candidate intervention – turnstile access – was already deployed, we were able to translate characteristics to the model of the secure door.

Model results showed that if one turnstile was implemented at the observed site, an average of 0.5 seconds would be added to individual entry times, amounting to over sixty hours for the site as a whole over a year. Three turnstiles approach the time cost of a

secure door, incurring almost an extra day for collective entry times over a year. The approach demonstrates that situated data collection and predictive modelling can be used to anticipate changes to the individual cost of security, but also calibrate interventions to local parameters.

2. RELATED WORK

Morisset et al. [17] describe a formal model of *soft enforcement*, where an ‘influencer’ exerts control upon individual ‘decision makers’. The influencer manages a set of rules dictating the behaviour of the decision maker, observes the environment in which the decision maker operates, and exerts an effect upon the decision maker. The decision-maker will make observations about the environment and make decisions accordingly. In our work we reflect observations of how an individual may respond to situations of non-compliance or costly compliance, including the presence of other people around them. The presence of multiple influencers is considered in the soft enforcement model, where these may for instance be managers of separate business functions. Here we aim to capture the influence of function managers in the modelled system, in part to explain the regular employee behaviours that are modelled and in turn the performance outcomes that are of importance.

Lenzini et al. [15] provide a modelling formalism to support planning around whether and how a malicious party can gain unauthorised access to resources in a complex physical location. The modelling approach the authors describe allows definition of people and objects within the locations that are being considered. Agents – including intruders – are able to act upon objects and move across locations. The approach is concerned with quantifying the cost and probability of a successful attack, where employee interactions with doors and other objects indirectly contribute to the outcome. Here we model employee interactions with physical entryways in terms of routine movement into and out of a secured building, to understand the impact of infrastructure decisions upon employee behaviour within an observed location (or a specific area within that location).

Beautement et al. [4] model employee use of USB storage devices across various locations and their related risk profiles. Beautement et al. use interviews with employees in an organisation to inform the model design, where we use consultation with system managers and direct observation of employee behaviour. A balance is considered between training, monitoring, and IT support costs; here support is represented by the reception staff, and in part by the physical security staff. A similar model is described in the work by Arnell et al. [2] to model password usage and support in a large organisation. The model leverages empirical password studies, including direct observation in a controlled environment, to develop measures of the time required to create and enter a password. A paradigm of Breaches, Productivity and Investment builds on the interrelated measures described by Beautement et al. Here we consider one location, focusing on the individual cost of security and the impact of non-compliant behaviour on the individual and those around them.

Rather than building a model, Wornton [20] describes how socio-technical and resilience frameworks can be used to examine various factors which contribute to a particular threat. The socio-technical perspective illustrates interrelationships between goals, people, physical infrastructure, technology, culture, and procedures. Here we explore the presence of similar factors and how they influence each other. Wornton proposes a resilience framework intended to avoid purely responsive approaches to system management, relating steps to anticipate, plan, implement, monitor, respond, and learn. Here we use predictive systems modelling techniques to forecast the im-

pact of potential system events on a wide range of interdependent factors in the observed environment.

3. BACKGROUND

This work demonstrates steps in a repeatable approach to engagement with large organisations and their security infrastructure, combining empirical data-gathering techniques with a systems model approach:

1. Identify the main security concerns of decision-maker stakeholders.
2. Interview a sample of employees at various levels and in a range of business functions, focusing on their interactions with IT security and physical security policies as part of primary work activities.
3. Survey a wider range of employees using scenario-based questions derived from interviews.
4. Identify core security concerns from survey results, identifying *hotspots* in the alignment of security and productivity.
5. Identify a limited number of promising candidate interventions through discussions with the partner organisation.
6. Build a model of interventions, using real-world observations and data from preceding steps in the process. Use the model to explore different configurations, to forecast their effect upon the organisation and its employees.
7. Deploy the intervention, and measure performance over time to understand if it is successful.

An application of the first 5 steps of this process is described in more detail in [3]. Here we focus on the 6th step through application of a systems modelling approach [9, 10], to explore the potential impact of an intervention identified at a partner company (the preceding 5th step). Being able to predict the impact of security interventions before deployment is especially important for protection measures; procurement and deployment may be a one-time cost, but also a sunk cost which cannot be recovered should a protection measure later need to be decommissioned [8]. If *security hygiene* is not practiced during design and deployment, it is not assured that employees will be able to effectively use the security mechanisms provided for them – if security is not doable, it creates additional burden for those individuals [14].

3.1 Security-related decision-making

The scenario-based approach of the survey exercise can capture how individual and organisational factors in the workplace impact security behaviours. A representative survey question related to physical building access is detailed below (reproduced from [3]). The topic of tailgating and unescorted visitors was investigated at a number of partner organisations:

Jessica is heading toward an unmanned turnstile and notices a man she does not recognise in front of her pass through the barrier by following close behind someone else unfamiliar. The two men are walking close together although they do not appear to obviously be in conversation. The second man is holding a cup of coffee in one hand and his laptop in the other. His ID badge is not immediately visible. Jessica decides to:

- Follow the man and ask to see his ID badge.
- Find a security guard at one of the manned turnstiles and tell them what happened.
- Return to her desk, she sees this sort of thing quite regularly and it is probably because his hands were full that he did not swipe through himself.
- Do nothing, if he is up to some mischief the security guards will catch him later on.

The scenario encapsulates a number of factors to investigate further: the visibility of ID badges; the presence of physical security staff in a lobby area, and; the ability an individual has to traverse a secure entryway (such as a secure door or turnstile) depending on what items they may be carrying about their person (such as a bag or suitcase, etc.). This then directs modelling efforts to: the design of the turnstile as an intervention measure; how easy it is for employees to use; both the expected and typical use of ID badges to gain access, and; the interactions not only between employees and the entryway but also between employees and lobby staff such as physical security and reception staff.

3.2 Modelling interventions

Security managers in organisations define security policies which dictate rules and expectations for both employee behaviour and how security controls should be used. A modelling framework based on mathematical systems modelling [10] is applied here, alongside simulation techniques to explore different kinds of interventions and their potential impact upon an existing organisation. Decision-making economics not only inform the model, but also guide the collection of information from stakeholders at the organisation under observation.

This approach can be used to explore the consequences of specific intervention design choices, towards more effective policy and deployment decisions. The models capture the logical and physical structure of systems, the choices and behaviour of *agents* within the system (such as employees), and managers' preferences about what constitutes a desirable outcome. Models are configured from observational data—in this case, the observation of employee behaviour at an one site operated by a partner organisation. The methodology follows familiar practices from classical applied mathematics and simulation (see Figure 1), where we focus on the *observations*, *models*, and *consequences*.




Figure 1: The classical modelling cycle

Following this approach we develop executable systems models of distributed systems [9], composed of *locations*, *resources*, and *processes*. Locations represent physical and logical locations, connected to become a location graph. A location can contain resources such as physical objects, people, or information. Processes then model changes in the system, acting upon defined resources.

For the purposes of coordinating real-world data collection, a representation of a real location, such as a room, is built up from basic atomic locations identified in the observed environment. For example, an office might have a breakout area, a number of desks, and aisles between them, all of which could be represented as different locations. Logical structures are also represented using locations: different locations in computer memory, or computers on a network, and so on.

This modelling approach, based on distributed-systems modelling, can be used to represent a wide range of systems – here we focus on different physical access control mechanisms. The approach can be used to model systems where information—rather than physical—security is the subject of interest, and even other, non-security-related systems. The modelling approach is a generic tool, that in this case is applied to help evaluate the *individual cost* for employees when interacting with the physical security mechanisms of a system.

3.2.1 Agents, decision-making, and manager preferences

Models contain *agents*, which act autonomously according to a set of decisions, interacting with other actors and resources. An agent will move around the location graph according to particular rules, and at the same time have resources associated with it which they may be carrying or wearing, or information that is known to the agent. These agents can then represent employees in an organisation, moving around a specific physical site. An agent will interact with security mechanisms in a given location, making decisions which (knowingly or unknowingly) relate to policy expectations. This is represented as a choice function, modelling a point in a process at which a decision is made amongst a number of alternatives. System managers can use a model to explore the consequences of different policy design decisions, where these are captured as *preferences*.

3.2.2 Example – tailgating

During the engagement with the partner organisation described in this paper, tailgating and unescorted visitors emerged from interview and survey analysis as a candidate security *hotspot*. In the time since the original engagement, the organisation has chosen to investigate physical access control interventions.

A prototype tailgating model was previously built [9], as a generic representation of building entry behaviour. Principally, the model explores the behaviour of both employees and attackers in the lobby near a secured area. Rather than modelling attacker behaviour, here we focus on employee behaviour and the impact that navigating security measures has upon employee time, as a measure of individual *security effort*.

In the generic model each employee has an ID card, needed to navigate a secure door or turnstile. Based on the behaviour of those around them, or the state of resources (such as whether a door is already open), an employee may respond in different ways. Their response is also governed by their own resources, for instance if they have forgotten their ID card. An employee may then ask for a temporary pass to facilitate continued access for the day, or choose to tailgate through the door by following another employee.

Similarly, an employee may observe another person tailgating, or – depending upon the culture of the organisation – hold a door open if they see someone approaching the door after them. Physical security staff in the model may also interact with employees, or by their presence (or lack of) influence behaviour near the secure entrance.

The basic model described in [9] is comprised of five locations:

Outside, Lobby, Reception, Entryway, and Atrium. Agents start in the Outside location and progress to the Lobby. Here we incorporate a simplified notion of an *Agent*, with a focus on modelling behaviours in the Lobby, which is connected to both the Outside and the Atrium (which in turn leads to the main building).

The basic model tells us that in order to explore interventions, we must determine the security policy and local rules at an observed site. This is achieved most immediately by consulting physical security staff. The prevalent cultural norms (such as holding doors open or challenging strangers) can also be determined through discussions with the site manager. An aim of this exercise is then to capture the preferences of these stakeholders with regards to employee security behaviour.

4. METHODOLOGY

Informed by the approach described in the Background section, the following elements are required for modelling a physical access intervention and identifying potential consequences of deployment:

Data collection

1. Capture characteristics of the existing entry mechanism and its surrounding environment, and potential intervention(s).
2. Document decision-makers' preferences.
3. Observe employee treatment of ID cards, actions in the Lobby area, and responses to events in the Lobby.

Modelling

1. Define Locations, Resources, Processes and Agents which represent the environment under observation.
2. Define Agent behaviours based on observed employee behaviour.
3. Define model rules for how different physical access mechanisms affect individual behaviour, resulting in changes to the individual cost—here, we consider time as a representative measure—to gain access to a secured area.
4. Run the model a large number of times, resulting in a distribution of outputs.

Models of organisation-based behaviours can be used to bring together knowledge and data from different sources [12], as we do here by combining manager perspectives, observations of the physical environment, and direct observation of employee behaviour within a specific building.

4.1 Physical access interventions

We focused on one specific site operated by the partner organisation. This site had the means to allow direct observation of employee activity via Closed Circuit Television (CCTV) footage, but was also large enough that footfall in the Lobby area of the building would provide sufficient data to demonstrate regular behaviours and interactions between individual employees, groups of employees, and the security apparatus in the Lobby.

The site under observation is a major site, but not the head office. In fact, this is the reason behind the physical security intervention in this instance: the head office implemented a particular system, which may be replicated at smaller sites, providing consistency across sites (where some staff may travel between sites).

It is the specifics of the deployed access control mechanism that are studied here, serving to scope the model. Hence, real-world

behaviour data gathered from the observed site is important – regular behaviour around existing access control mechanisms will be observed, and a sample of behaviour data for the 'new' system collected, to then be translated to the model of the observed site.

4.2 Preference elicitation

We met with policy decision-makers directly involved in the arrangement, resourcing, and staffing of the site. Decisions about security policy may filter down through an organisation to increasingly local teams (and in this case, sites), where policy is interpreted and implemented according to the preferences and capabilities of local decision-makers [5].

Discussions involved the site manager and the physical security manager; the site manager had a sense of the impact of security upon the general workforce and the organisation's working culture, and was also in a position to explain the norms of the site (for instance when staff break for lunch every day). The physical security manager characterised the security culture at the site, how policy is imparted to staff, and in essence how secure the regular observed behaviour of staff at the site is relative to policy and security expectations.

A *preference elicitation* exercise with each of these managers elicited: (i) which other elements of site operation relate to physical security (e.g., fire safety), and (ii) where physical security is positioned as a priority for the manager. Each manager's own activities around security may also influence employee behaviour – for instance, a manager may devote some of their own time to reaching out to employees to remind them of expected security behaviour on-site. If many employees find it difficult to use an access mechanism such as a door or turnstile, it could potentially become the subject of wider discussions with staff based at the site.

4.3 Data collection

The interview and survey techniques described in the Background collect self-reported behaviour data, where respondents may report that they behave in one particular way and yet potentially act differently in reality. Direct observation is then critical for understanding the reality at a specific observed site, as well as for calibrating the model.

Relating to the security culture at the site, challenging of strangers seen on-site was part of the culture – if we were to stand in the Lobby area of the building in order to conduct observations, we would then immediately be interfering with the routine behaviour of the employees being observed. As an alternative, the site under observation had a range of CCTV cameras, both within and outside the Lobby area. Crucially, these cameras provided image timestamps, and a collective view of:

- The approach to the Lobby from Outside.
- The card reader near the Secure Door, and the surrounding Lobby.
- The Secure Door leading to the Atrium.

We then chose a day of the week with relatively high footfall of both employees and visitors to the site, on the advice of managers. Timestamps of employee interaction with the secure door also indicated the *individual cost* of using the security mechanism and associated card reader; this would inform the creation of the model. Routine responses to events involving the door were also captured (e.g., if someone tried to open the door having not used the ID card reader as intended).

4.4 Employee observations

The two authors observed CCTV footage from a combination of camera viewpoints. Dedicated data collection software, created specifically for the study, is shown in Figure 2. It was necessary to capture the timestamp of each observed event manually, to the nearest second, with agreement by both researchers. When an event occurred, the video would be paused – or moved back and played again at a much slower speed – to capture a sufficiently precise timestamp and a detailed understanding of the observed behaviour. Researchers would watch 3-4 different viewpoints between them. One would manage the camera controls, the other data collection using the dedicated software. In the centre of Figure 2 it can be seen that a sequence of events is built up, to track each individual as they move through the Lobby.

As shown at the bottom of Figure 2 there is a function to capture additional notes – this was important when building up the dataset of observations, as a sense of regular and irregular events developed. This could include going through a secure door without incident, or needing to stop near the card reader and spend time locating a personal access card. Having a basic explanation for the duration of a sequence of security behaviour events informs modelling later on.

We observed existing pre-recorded CCTV footage from a day in the past, did not create any new system data and examined footage only while on-site. No personally identifying information, or details of personal physical characteristics, was logged in the data collection tool or discussed with staff at the site. Prior consultations with managers at the site was then useful for providing some explanation for the behaviours seen in the footage (e.g., visitors to the site will go to the reception to sign in, then go to the seating area to wait for their host). This removed the need to involve security staff during the data collection phase.

We were interested foremost in recording the time taken to get through a secure door for regular staff, as well as staff with visitors. We were also interested in capturing the footfall of staff and how staff interact with others around a secure door (for instance when a queue forms), where we then compare this with a turnstile and map the costs onto the original site. To do this, a model of the system is built, using observation data points, which then allows us to forecast the *individual cost* of physical security to enter a secured area.

4.5 Models of physical access control

Outcomes from the preference elicitation exercise and employee observations inform the design and calibration of a model of the Lobby area and surrounding parts of the building. We also visited another site with turnstiles already in place, to observe limited footage of the use of turnstiles. We captured timings of regular events, but also documented responses to those events which were comparable to observations relating to secure doors. For instance, if an employee's ID card does not appear to work and permit access, they may retry, go to reception, etc. Capturing timing and responses for one or more turnstiles can then allow us to model a series of variants overlaid on the secure door model, e.g. two turnstiles side-by-side.

5. RESULTS – DATA COLLECTION

Here we report data collected about: the Locations and Resources to be modelled (components of the Lobby); the Agents and their Processes (employees navigating the secure door); and variants to model (samples of Location timings for both secure doors and turnstiles). Collecting real-world data for these model components allows us to situate the models and provide meaningful information

Event	Time	Notes
From Atrium	12:31:17	Wearing Lanyard,
To Desk	12:31:22	
Leave Desk	12:31:52	
To Atrium	12:31:56	

Figure 2: The event-logging window within the data collection tool.

to support infrastructure decisions.

5.1 The observed site

As in Figure 3, the site has one main door, used for entry and exit by all staff — other doors exist on the site, but are used in emergencies only. Once inside the Lobby, there is a Reception Desk to one side, and a Security Reception to the other side (for obtaining ID badges). As part of the Lobby there is a Visitor Area, where visitors wait for hosts who come through the Secure Door to collect them. Visitors spend time at the Reception Desk signing in and getting temporary (restricted access) passes. The card reader is located close to the Reception Desk. Once beyond the Secure Door, staff enter the Atrium and from there can reach other parts of the site.

5.2 Decision-maker preferences

By talking to decision-makers at the site, we are able to discern between company policy and local provisioning decisions. Security was amongst the site manager's priorities, as was being able to plan cleaning, catering and grounds maintenance around employee movement within the site. For instance, staff at the observed site were seen to move in large numbers toward the on-site cafeteria at a fixed time, where the secure door can become a bottleneck with staff.

The physical security manager discussed security behaviours, but also supporting procedures; their team's duties included creating visitor sheets at least a day in advance and managing visitors once they are on the site. We then observe interactions with reception staff at the Reception Desk location, but also Security Administration staff at the Security Reception. Visitors can include staff from other sites, new starters, and contractors who need different but routine arrangements for their ID badges. To maintain site security, there is a baseline level of staff in the Lobby, with at least one member of physical security staff by the secure door and at least one member of reception staff at the desk at all times (so visitors may queue at the desk, but will not be unattended).

5.2.1 Organisation security culture

Both the site manager and the physical security manager noted typical behaviours of staff at the site. Staff, for instance, may wear their ID badge in different ways, which affects both visibility but

also their ability to locate the badge when they need to pass through the secure door.

Physical security managers would point out lapses upon sight (where staff clearly do not have their badge on their person). Discussions with the two different decision-makers informed the distinction between regular behaviour and ‘persistent offenders’, individuals who routinely forget to wear their badges (approximately 1 percent of people on-site on a given day). In most cases badges were on the individual’s person or in their car, meaning at most a short trip to the car park and back to the Lobby. However, the vast majority of offenders were seen to change their behaviour gradually over time and routinely wear their ID badges.

Considering the challenging of strangers (as unescorted visitors), the physical security manager noted that it was not straightforward to change a normally ‘polite’ attitude. Observation and modelling of interactions between individuals in the Lobby was then important (e.g., holding a door open for someone following behind), as physical security staff may for instance join regular team talks to highlight recent issues and provide reminders of secure working behaviours. Anticipating infractions or confusions around how to use a physical access mechanism can help decide how best to use that limited time.

5.3 Behaviour observations

Observation of employee activities identifies individual patterns (Processes) as employees move between Locations within the building (e.g. Outside to Lobby to Atrium) and across sub-Locations (such as from the Reception Desk to the Secure Door). Normally an employee enters the Lobby from Outside and walks straight to the Secure Door via the Card Reader — routinely it was seen to take 10 seconds to walk across the Lobby from Outside and one second to go through the Secure Door. We did not see any suspicious activity in the footage we observed: ID badges were worn, displayed, or used successfully with the card reader in all cases of building entry.

Observations also identified other regular routines, such as employees collecting visitors and then returning them to the Lobby after a meeting, and visitors and staff from other sites signing in (entering the Lobby, moving straight to the Reception Desk, then going to the Visitor Area or the Security Window). In some instances an individual may take a few seconds to retrieve their card (even when wearing their ID badge, in the case of those wearing coats or carrying luggage).

Groups of people may enter the Lobby from either Outside or the Atrium; we do not capture whether individuals in these groups are known to each other or not, however holding doors open for others was one regular observation (the ‘politeness’ mentioned by one of the managers). This may be a combination of politeness and practicality — most staff wore their ID badge about their person so that it was clearly visible, and the card reader made a noise when a card was tapped against it, which could be heard by others walking ahead.

5.4 Physical access interventions

Not all interventions are wholly tailored to a site, and may be based on known solutions and existing knowledge from larger or main organisation sites, at least for a large organisation such as the one observed in this study. We then describe each access mechanism according to: the physical actions it allows (e.g., one-way or two-way movement, if and how tailgating is possible); how individuals and groups respond to the mechanism, and; in combination, the *individual cost* of traversing the access mechanism.

We do not model the interplay between ID badges and an ac-

cess mechanism, for instance whether they track if a specific individual is on-site or not. Although this may influence the regular behaviours of staff at the site, this would constitute a separate intervention. It was also observed that the vast majority of staff wore their ID badges and used them to tap the card reader (even when a door was open in front of them).

5.4.1 Secure door

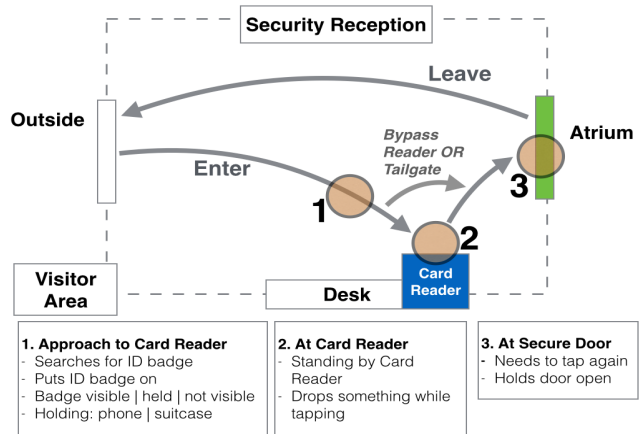


Figure 3: Access-related employee traits and behaviours of interest in the lobby area.

The representation of the Secure Door is derived from observed behaviours and their timings – a combination of our observations about how the mechanism operates, but also how employees interact with and respond to it. Regular and exceptional behaviours of interest are shown in Figure 3.

- **Regular use.** It was seen to take approximately one second to move through the door.
- **Failed read.** If the card reader was not properly activated, there was a walk of approximately two seconds between the Card Reader and the Secure Door, where the Card Reader was on the path between the Lobby and the Atrium. If someone does not tap the reader correctly, they in effect have to retrace their steps, adding two seconds to walk back, and two seconds to return to the Secure Door. There is also a delay of at least one second as an individual realises that they cannot get through the Secure Door.
- **Blocked access.** If visitors queue up at the Reception Desk, they may block the Card Reader, adding a few seconds of delay while an individual moves around them to access the reader.
- **Bottleneck.** Individuals may approach the door from both sides. Queues then form on either side, where the door may be held open until all staff have moved through it. Moving through the door takes the same amount of time, but those queuing on either side remain idle until they are at the door.

5.4.2 Turnstile

A turnstile was seen as a way to prevent tailgating, especially in combination with the consistent presence of physical security staff. We observed a limited amount of footage of turnstiles in active use at another site, but also spoke to the site manager and physical security manager there.

The turnstiles were two-way, and it took staff approximately two seconds to get from one side to the other. Of note is that the card reader is affixed directly on each turnstile, meaning that if reading fails, there is no need to retrace steps to try again. Having a wheeled suitcase may add two more seconds, as turnstiles restrict movement between both the walls of the turnstile and the bars that turn and close behind an individual as they walk through it.

We identify exceptions to normal behaviour for the Turnstile, to then model how they may manifest in the observed site. The intention is to transpose behaviours from the head office to the original site using the model. Observations of the Secure Door then constitute a ‘before’ picture, and the modelling of the Turnstile the ‘after’ picture. Exceptions to normal behaviour include:

- **Wrong turnstile.** An individual swipes the reader for the Turnstile next to the one they are stood in front of. They would then have to go through that Turnstile, even if someone else is already approaching it to go through. It takes one second for an individual to realise the error.
- **Change turnstile.** If a card read fails for some reason, an individual may try another turnstile. If instances are located side-by-side, rather than try the same turnstile again, an individual may instead try the one next to it.
- **Cutting across.** If a person moves across the path of other people on the other side of the turnstile, it could add another second to each of those people. That person is then also blocking the exit of that specific turnstile until they are out of the way of the adjacent turnstile. There are then implications for modelling more than one instance of a security mechanism as being side-by-side in the Lobby.
- **Failed read.** A person may fail to tap their badge correctly, while also not being able to watch or listen out for the reader to confirm that the badge has been read (which can occur where there are many turnstiles being used at busy times). The person may step into the bar, then step back, and tap again; the process then takes two seconds longer than normal. Similarly, a short delay is incurred by the person behind as the Turnstile is blocked.
- **Bottleneck.** Two people approach the Turnstile from different directions at the same time (where similar behaviour was seen for the Secure Door). One person will block the exit, then step back a little while waiting for the other person to start moving through (showing that they’re going through), then the “blocker” will move to an adjacent Turnstile.

6. MODEL AND RESULTS

We created a model to represent the main behaviours captured during data collection, and to enable exploration of the possible effects of a change from a secure door to turnstiles at the entrance to the building. We model employees entering and leaving the building, and interacting with the physical access controls. The structure of the model is relatively simple, as we are primarily interested in the performance of the two different kinds of access control, in terms of individual and collective time cost. The model however has the capacity to be expanded to include other parts of the building, where the compositionality of the modelling framework [10] intentionally supports this.

We use data gathered from our observations to design and parametrise the model (approximately 1800 sequences of actions). Figure 4 shows the arrival and departure times of employees from the

building for the day we observed. In the model, employees follow similar patterns. This is important, because differences in performance between the door and the turnstile appear when the entrance is congested during peak periods, such as the morning, evening, and lunch break.

The model contains four locations: outside the building, the entrance lobby, a location where the access control — be it a secure door or turnstile(s) — takes place, and the atrium inside the building. In the model, an employee typically enters the building, moves towards and then proceeds through the door (or turnstile, depending on the scenario), and then progresses into the atrium.

From our observations, we know the typical distribution of times that employees take crossing the lobby, moving through the secure door, and so on. We use these distributions in the model, sampling from them to determine the time taken for each employee movement. We also consider the non-standard behaviours discussed in the previous section, such as when an attempt to tap in with an ID badge fails, or when a turnstile is blocked by a person moving toward the entryway from the opposite direction, again using the times from our observations.

For the door, we model the time it takes to open and close. If there is a group of employees who have arrived at the same time, then the first one will take a few seconds to open the door, and the others can then follow behind with minimal delay while the door stays open. For turnstiles, this is not the case. For these, we model a mixture of queueing behaviours, representative of the behaviours identified during data collection. Some employees move to whichever turnstile is available (in scenarios where there is more than one) while others queue for a particular turnstile waiting for it to become available.

We also model the possible delays that can occur for both door and turnstile, based on the regularity of events seen in the data sample we collected. For example, attempting to tap in with the ID card fails approximately 1.5% of the time, requiring the employee to tap again, which takes around two seconds on average to do. Variation is incorporated into regular events such as moving through an entryway, by way of small random time costs added to the actions derived from sampled data — this accounts for a range of behaviours, such as stopping to locate an ID badge about one’s person.

6.1 Results

We use the model to explore how the switch from doors to turnstiles would affect the time it takes for employees to enter or leave the building. We look at four cases: the first as it currently is, with a door, and the others with one, two, and three turnstiles replacing the door. In the different cases, the only element which differs is the access control. The rest of the environment, such as the time it takes for an employee to walk from the entrance of the lobby to the secure door or turnstile, remains constant. We run each different scenario 1000 times to estimate the average performance, which accounts for different possible random variations in timing. For example, on an average day, there might only be a few cases where more than three people arrive simultaneously; however, it is still possible that there will be days when there are many instances of simultaneous arrivals, causing larger queues. By running the model a large number of times with different random arrival times (based on observed distributions), the overall performance of the system can be estimated.

The results from the simulations are shown in Table 1. For each simulation run, we record the time taken by employees to enter and exit the building, passing through the different access control mechanisms. The values we present are based on the average times taken by employees over the 1000 simulation runs for each scenario.

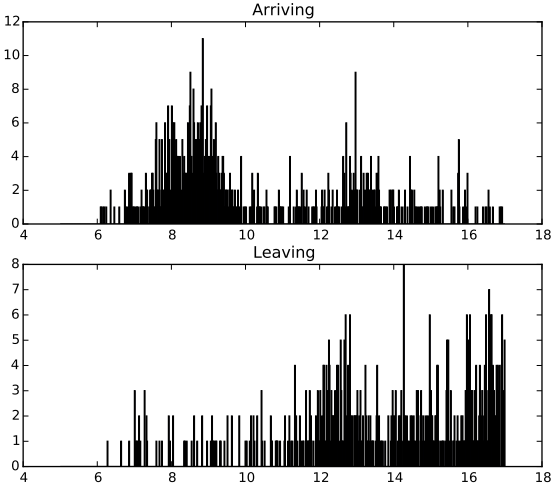


Figure 4: Histograms of employees’ entry and exit times, with time of day (in 24 hour format) on the x-axis and entry/exit events per minute on the y-axis. Busy entry times can be seen at 7:30–9:30am. The lunch break can be seen 12–1:30pm where there is a mix of entry and exit events as employees leave and re-enter the site. From 4pm staff leave the building. Other times illustrate individuals moving to and from the lobby as part of their work.

First are the average enter and exit times for each of the scenarios. The secure door has the lowest times, with employees taking on average 10.9 and 7.8 seconds to enter and leave, respectively. For one turnstile, the values increase to 11.4 and 8.5 seconds, but as additional turnstiles are added, the time drops closer to that of the door: 11.1 and 8.1 seconds for two turnstiles, and 11.0 and 8.0 seconds for three.

The additional cost for an individual employee of even the worst, one-turnstile case is not particularly high. However, this site has a fairly large number of employees, which means that the total amount of additional time spent by employees can be significant. For one turnstile, each day, employees collectively spend an additional 875.2 seconds entering and leaving the building compared to the door. For two and three turnstiles it is 354.5 and 310.2 seconds, respectively.

It is interesting to note the difference in performance between the different turnstile scenarios. With one turnstile, the time spent is much greater than the other two, which are fairly close in comparison. Having just one turnstile causes bottlenecks at peak times, when large numbers of employees are arriving or leaving. With two turnstiles, the delays at these times are reduced by more than half. Adding a third turnstile further reduces the time spent, but by a much smaller amount, as three turnstiles are only rarely used simultaneously. With a door as the access mechanism, even though employees are still required to tap in, entry times are reduced if the door is opened by one person then held open by others who arrive and keep the door open. This minimises delay, but also demonstrates the separation between interacting with the card reader and interacting with the door. Data collection demonstrated high levels of both successful interaction with the card reader and wearing of ID badges.

Over larger amounts of time, the differences in performance add

	Door	1 TS	2 TS	3 TS
Enter time (sec/person)	10.9	11.4	11.1	11.0
Exit time (sec/person)	7.8	8.5	8.1	8.0
Total time (sec/day)	12423.7	13298.9	12778.1	12734.0
Inc. over door (sec/day)		875.2	354.4	310.2
Inc. per week (min)		72.9	29.5	25.9
Inc. per year (hours)		63.2	25.6	22.4

Table 1: Model outputs for configurations of one Secure Door, and a number of Turnstiles (TS). These are average values over 1000 simulations, showing the time taken per person to enter and exit, the total employee time spent per day, and the increase in time over the door when using turnstiles.

up. One turnstile takes an additional 72.9 minutes per week over the door, with two turnstiles at 29.5 minutes and three taking 25.9 minutes. Over the course of a year (considering 260 working days per year), this means that 63.2 additional hours of collective employee time are spent with one turnstile, 25.6 hours with two, and 22.4 hours with three.

When considering whether or not to install turnstiles, or how many of them to use, decision makers have to balance the costs of this additional time against the benefits that the turnstiles provide. In terms of security this would be a reduced chance of tailgating, and in terms of the experience of moving through the building an increase or decrease in time spent queueing to enter or exit. Although three turnstiles does not provide a much greater reduction in time than two, it might still be worth using three, if possible, as it provides redundancy. The difference in performance between one and two turnstiles is large; if only two were installed, and for instance one of these developed a fault, it would cause a large increase in the employee time spent entering and exiting. If there were three turnstiles and one was taken out of use, then the increase in time would not be large.

7. DISCUSSION

The physical security presence at the observed site was maintained at a high level, with at least one member of physical security staff present in the lobby at all times. Following this rationale, we modelled cases where there was an over-provisioning of turnstiles to maintain security and orderly entry and exit. Organisations with fewer resources or a different security culture may make different choices about the access mechanisms they deploy. Similarly, the routines at a site may govern the individual cost of security; a site and employees may follow routines that are unlike what we observed for e.g., morning arrival onto the site, or an organisation may have very few visitors.

Should turnstiles be introduced at the observed site, managers may opt to provide redundancy in the number of turnstiles to limit queueing. In this sense, support provisions, security and productivity are managed together – a balance of these three considerations has been seen to resonate with information security managers [18], and here we see a similar paradigm shared by both a site manager and a physical security manager. A similar approach of using CCTV analysis to investigate business targets has previously been demonstrated in the retail sector, when tracking the removal of stock from stores [7].

The model results show a potential increase in collective entry time for staff at the site if turnstiles are deployed; approximately one additional day is added over a year for three turnstiles, going up to more than sixty hours if only one turnstile is deployed. Her-

ley [11] frames various information security mechanisms in terms of the costs and benefits (both potential and actual) for the individual. Extra time may be added with the introduction of turnstiles, but further study may determine if the characteristics of a turnstile are experienced or perceived as benefits by employees, e.g., having the card reader affixed directly to the security mechanism, where most staff were seen to incorporate tapping their ID badge into a fluid movement from Outside directly to the Atrium. These factors would then need to be framed in terms of perceived and actual benefits.

The observed site illustrated a kind of *everyday security* [16], where employees and managers collectively practiced secure behaviour. The reality of everyday behaviour is accommodated, where workable solutions then emerge – this was subtly demonstrated in the footage we reviewed, which was taken during cold months. Many staff entered the building wearing coats which partially or completely obscured their ID badges, but were not challenged as long as they could produce an ID badge upon reaching the card reader.

7.1 Limitations

We relied on manual logging of events, such that although effort was taken to capture accurate timing data, the repeatability of the approach was reduced. The CCTV cameras were positioned for security, but provided an incomplete view of some behaviour events, such as whether a person's lanyard was clearly visible to physical security staff in the room (even if it may not have been visible from the view of the camera). Repeated or slowed playback was used to verify observations in such cases, and prior discussions with managers identified the various signs of routine patterns of behaviour.

Data collection was expensive in terms of the time required to capture individual events. Three full days were required to capture relevant events during the footage of one day of site activity in and around the lobby, involving the combined effort of two researchers with dedicated event-logging software. This limited our capacity to also log events on a quieter day of the week, for instance. However, by observing a busy day of the week we captured the implications of increased footfall, which was critical for modelling the impact that potential interventions would have on the individual cost of security.

We do not model staff learning a new system, such as a turnstile. Individuals may require time to change existing habits or learn new ones. This was seen at the site with turnstiles already in place, where there were instances of individuals tapping the reader for an adjacent turnstile. The rate at which new employees join an organisation may then also influence the occurrence of irregular behaviours, a factor which in future could be investigated with policy stakeholders.

8. CONCLUSIONS

Systems modelling techniques were applied to model the potential replacement of a secure door with a turnstile at a partner organisation site, and the impact this might have on the *individual cost* of security for employees and visitors. The modelling approach also guided data collection at the site to situate the model. Preference elicitation exercises were conducted with the site manager and physical security manager at the observed site, to identify meaningful model parameters and regular behaviours. Direct observation of employee behaviours from CCTV footage provided localised data to support the model. This included responses to security events, such as returning to the card reader after a failed read or moving to a different turnstile.

Model results showed that if one turnstile was implemented at

the observed site, an average of 0.5 seconds would be added to individual entry times, amounting to over sixty hours for the site as a whole over a year. Deploying three turnstiles together would approach the individual time cost for each person of having a secure door in place. The approach demonstrates the application of situated data collection and predictive modelling to anticipate changes to the individual cost of security, but also in calibrating interventions.

Future work will involve more direct co-design of interventions with partner organisations, informed by findings from engagement with employees through interviews and surveys and consideration of human factors of security principles. An intervention can involve the complete replacement of an existing technology solution, or otherwise the recalibration of controls which are already in place. The means to effectively measure the performance of workplace security over time will also be explored; here we manually examined CCTV logs, and in future it may be that appropriate data can be collected automatically within an observed system to provide an indication of employee satisfaction with security.

Acknowledgments

The authors wish to thank the participating organisation for their generous cooperation, and thank Kat Krol for comments. This work was supported by UK EPSRC, grant nr. EP/K006517/1 (“Productive Security”).

9. REFERENCES

- [1] Anne Adams and Martina Angela Sasse. 1999. Users are not the enemy. *Communications of the ACM* 42, 12 (1999), 40–46.
- [2] Simon Arnell, Adam Beutement, Philip Inglesant, Brian Monahan, David Pym, and Angela Sasse. 2012. Systematic decision making in security management modelling password usage and support. In *International Workshop on Quantitative Aspects in Security Assurance*. Pisa, Italy. Citeseer.
- [3] Adam Beutement, Ingolf Becker, Simon Parkin, Kat Krol, and M. Angela Sasse. Productive Security: A Scalable Methodology for Analysing Employee Security Behaviours. In *Symposium on Usable Privacy and Security 2016 (SOUPS)* (2016). USENIX.
- [4] Adam Beutement, Robert Coles, Jonathan Griffin, Christos Ioannidis, Brian Monahan, David Pym, Angela Sasse, and Mike Wonham. 2009a. Modelling the human and technological costs and benefits of USB memory stick security. In *Managing Information Risk and the Economics of Security*. Springer, 141–163.
- [5] Adam Beutement and David Pym. 2010. Structured Systems Economics for Security Management.. In *WEIS*.
- [6] Adam Beutement, M Angela Sasse, and Mike Wonham. 2009b. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 workshop on New security paradigms*. ACM, 47–58.
- [7] Adrian Beck and Andrew Willis. 1999. Context-specific measures of CCTV effectiveness in the retail sector. *Surveillance of public space: CCTV, street lighting and crime prevention, crime prevention studies series* 10 (1999), 251–269.
- [8] Rainer Böhme. 2010. Security metrics and security investment models. In *International Workshop on Security*. Springer, 10–24.

- [9] Tristan Caulfield and David Pym. 2015a. Improving Security Policy Decisions with Models. *IEEE Security and Privacy Magazine* 13, 5 (2015), 34–41.
- [10] Tristan Caulfield and David Pym. 2015b. Modelling and simulating systems security policy. In *Proceedings of the 8th International Conference on Simulation Tools and Techniques*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 9–18.
- [11] Cormac Herley. 2009. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*. ACM, 133–144.
- [12] Helen PN Hughes, Chris W Clegg, Mark A Robinson, and Richard M Crowder. 2012. Agent-based modelling and simulation: The potential contribution to organizational psychology. *Journal of Occupational and Organizational Psychology* 85, 3 (2012), 487–502.
- [13] Iacovos Kirlappos, Adam Beautement, and M Angela Sasse. 2013. Comply or Die Is Dead: Long live security-aware principal agents. In *International Conference on Financial Cryptography and Data Security*. Springer, 70–82.
- [14] Iacovos Kirlappos, Simon Parkin, and Martina Angela Sasse. Learning from “Shadow Security”: Why understanding non-compliance provides the basis for effective security. In *Workshop on Usable Security (USEC 2014)* (2014).
- [15] Gabriele Lenzini, Sjouke Mauw, and Samir Ouchani. 2015. Security analysis of socio-technical physical systems. *Computers & electrical engineering* 47 (2015), 258–274.
- [16] Harvey Molotch. 2013. Everyday Security: Default to Decency. *IEEE Security & Privacy* 11, 6 (2013), 84–87.
- [17] Charles Morisset, Iryna Yevseyeva, Thomas Groß, and Aad van Moorsel. 2014. A formal model for soft enforcement: influencing the decision-maker. In *International Workshop on Security and Trust Management*. Springer, 113–128.
- [18] Simon Parkin, Aad Van Moorsel, Philip Inglesant, and M Angela Sasse. 2010. A stealth approach to usable security: helping IT security managers to identify workable security solutions. In *Proceedings of the 2010 workshop on New security paradigms*. ACM, 33–50.
- [19] Edgar H Schein. 2010. *Organizational culture and leadership*. Vol. 2. John Wiley & Sons.
- [20] Katharine E Worton. 2012. Using socio-technical and resilience frameworks to anticipate threat. In *2012 Workshop on Socio-Technical Aspects in Security and Trust*. IEEE, 19–26.