# Systems Modelling for Economic Analyses of Security Investments: A Case Study in Identity and Access Management

Adrian Baldwin, Marco Casassa Mont, David Pym, Simon Shiu

HP Laboratories
HPL-2009-173

**Keyword(s):**
security analytics, identity management, economics

**Abstract:**

Identity and Access Management (IAM) is a key issue for systems security managers such as CISOs. More specifically, it is a difficult problem to understand how different investments in people, process, and technology affect the intended security outcomes. We position this problem within the framework of optimal control models in macroeconomics, and use a process model to understand the dynamics of the utility of possible trade-offs between investment, access, and security incidents (breaches). A utility function is used to express the security manager's IAM preferences, and the functional behaviour of its components is described via a process model. Executing our process model as Monte Carlo simulations, we illustrate the behaviour of the utility function for varying levels of investment and threat, and so provide the beginnings of a decision-support tool for systems security managers.

External Posting Date: July 21, 2009 [Fulltext] Approved for External Publication
Internal Posting Date: July 21, 2009 [Fulltext]

Presented in Trust Economics Workshop, London, June 2009

# Systems Modelling for Economic Analyses of Security Investments: A Case Study in Identity and Access Management

Adrian Baldwin Marco Casassa Mont David Pym Simon Shiu
HP Labs, Bristol
England, U.K.

**Abstract**

Identity and Access Management (IAM) is a key issue for systems security managers such as CISOs. More specifically, it is a difficult problem to understand how different investments in people, process, and technology affect the intended security outcomes. We position this problem within the framework of optimal control models in macroeconomics, and use a process model to understand the dynamics of the utility of possible trade-offs between investment, access, and security incidents (breaches). A utility function is used to express the security manager's IAM preferences, and the functional behaviour of its components is described via a process model. Executing our process model as Monte Carlo simulations, we illustrate the behaviour of the utility function for varying levels of investment and threat, and so provide the beginnings of a decision-support tool for systems security managers.

## 1 Introduction

Since CISOs have finite budgets, security investment strategy involves choices between risks and outcomes. Moreover, many of the outcomes and choices are intuitively correlated; that is, they trade off against one other. We are interested in how to help stakeholders (decision makers) better understand these trade-offs, and how to form a better-shared understanding of their preferences.

In this macro-economic style modelling approach, following the style outlined in [2] we can identify the components and preferences between them through utility functions such as:

$$U(C, A, K, t) \quad = \quad w_1(C - \bar{C})^2 + w_2(A - \bar{A})^2 + w_3(K - \bar{K})^2 \tag{1}$$

in which $C$, $A$, and $K$ and $\bar{C}$, $\bar{A}$, and $\bar{K}$ represent, respectively, the actual and target levels of confidentiality, availability and investment, and $w_i$s represent the appropriate weightings; $t$ denotes time. Thus the utility function is a weighted function of the deviations from target the three economic magnitudes whose mutual trade-off is of interest to us, with the weights expressing the decision-maker's preferences among the magnitudes. It should be noted that the quadratic form of the utility function is not the only choice available. It is, however, a convenient first step, derived from the basics of utility theory and portfolio theory, and provides a simple account of diminishing marginal utility. Richer choices are available, such as the asymmetric Linex functions employed in the work of Barro and Gordon [1], Nobay and Peel [8], and Ruge-Murcia [9]. Asymmetries in the components the utility function — in contrast to the symmetric form of the quadratic case — express the extent to which the security manager is relative more or less concerned about deviating above or below target.

The general setting for optimal control models is given in [6]. Having set up such an account of utility, our objective is to maximize it over the space of control variables which govern its dynamics

over time.[1] That is, the security strategy problem for a given organization is expressed as a utility maximization problem with respect to the organization's preferences.

The dynamics of $C$, $A$, and $K$ is explored through appropriately constructed system equations in [7]. In line with previous work by some of us and others [3, 2], however, this paper describes how to use a systems modelling approach to explore the components of interest. Systems models are intended to directly capture and represent (potentially multiple) stakeholders comprehension of a system, as such we believe they represent a more meaningful and trusted exploration of the utility components.

In the work presented in this paper, which analyzes trade-offs between security incidents, access, and investment in identity and access management, we do not have a set of system equations for the magnitudes that are of interest to us. Rather, we have a process model: an executable mathematical model of the system in which we are interested that is based on mathematical concepts of environment, location, resource, and process. The process model, when executed as a discrete-event simulation, produces as output numerical and graphical representations of incidents and access in a given investment context. Thus we can illustrate the desired utility function (see Section 3) in this case. Establishing an analytic connection between process models and the system equations required to drive utility maximization, as described above, is a topic for further research. Nevertheless, the illustrations of the utility function of interest that we are able to obtain prove to be a valuable guide in information security investment decision making for IAM.

IAM is a complex 'people, process, and technology' problem. It challenges CISOs on how to authenticate and authorize users; whether to centralize and automate processes (such as provisioning); and how to influence and reflect reliance on application and infrastructure security. With risks such as segregation of duties (SoD) IAM is also directly related to business level security and productivity concerns and so is a rich and relevant example for studying security strategy.

In our case study, we use a systems modelling approach to construct a mathematical model of a typical IAM system that can be tailored to fit different business and threat environments. We assume two investment instruments: configuration which covers provisioning and SoD, and enforcement which covers authentication, authorization, and general infrastructure and application security. We use Monte Carlo-style simulation to show the effect different investment choices will have on the predicted state of the system, and the predicted protection provided against different threat scenarios. Predictive modelling represents explicitly the causal dependencies within deployed security and IAM processes and provides a way to contextualise and calculate the metrics we then use for estimating the utility function.

## 2 The IAM Systems Model

Our previous security models [3, 10] have been written and run using the Demos2k [5] toolkit. In parallel with these studies, we have been designing and building a new tool chain, Gnosis — which captures the mathematical theory presented in [4], where a prototype implementation is also discussed — for executing the process models, its associated experiment manager, GXM, to manage the execution of Gnosis models as Monte Carlo simulations and to support the data and statistical analysis. The case study described here has been implemented and run using Gnosis. Separate work will describe the advantages and implications of the Gnosis tool-set. The description of the model provided should enable the reader to reconstruct an equivalent study using Demos2k.

The modelling idiom within which we work decomposes systems into four key conceptual facets: the collection of *processes* that characterize the behaviour of the system, the *resources* that are manipulated by the processes as they execute. and the *locations* around which the system is distributed, logically or spatially; finally, we consider the *environment*, described stochastically, within which the system exists. This idiom is supported mathematically by the work presented in [4] and is also discussed in [10].

[1] Alternatively, we may consider a loss (the opposite of utility) function and seek to minimize it.

### 2.1 Basic Structure of the Model

We construct a model with 2 investment instruments (control variables), one to set the level of investment on access configuration, and one to set the level of investment on enforcing this configuration. The investment of each ranges between 1-10, providing 100 different experimental runs, with (1,1) representing minimal investment in both instruments, and (10,10) the maximum. We also have the ability to vary the threat environment, although for the purposes of this study, we limited ourselves to two scenarios, one 'mild' and another 'full' which assumes many more internal and external 'attacks'. There are no assertions about the cost of moving the configuration instrument from, say 5 to 6, or of whether an investment of 4 is equivalent to a typical investment in enforcement. At this stage, the aim of the model is to explore the kinds of situations when diminishing returns are reached on a particular investment instrument, and to differentiate strategies based on assumptions about the threat environment.

The basic components of the model are a series of externalities that trigger IAM relevant processes, and each of the processes are affecting aspects of configuration and security state that we are interested in tracking.

Each of the processes models the effect these processes have on the IAM state. For example, new starters, staff leaving, job and organizational changes, introduction and retirement of applications, and applying automation to provisioning for an application all affect the configuration state, and so are included in the model. Similarly, application introduction, upgrade, retirement and migration projects each affect the overall enforcement state and so are included in the model.

We are interested only in the effect each of the processes has on the state, and so each of the processes are only defined in these terms. For example, this means the 'new starter' process does not capture any of the steps in the provisioning workflow; instead it simply records the expected side effect this process has on the access configuration. Moving the configuration instrument changes the assumed effect this process has each time it runs.

In parallel with the IAM processes are a series of threat processes such as internal fraud attempts, external hack attempts, former staff accessing application, and so on. These processes are also triggered by externalities. We assume more attacks will be thwarted by a good IAM state, and this is reflected in the way threat processes are modelled. Essentially, a good configuration and enforcement state will lower the chances that a threat process will succeed.

In the executable model, each of the processes are spawned in parallel, and relevant probability distributions are randomly sampled to determine when the external triggers fire and cause an instance of a process to run. Each of these affects the state models (configuration, enforcement and incidents) which we can sample to build a picture of what is going on.

There is not sufficient space to describe the model indetail, but, loosely, the investment instruments adjust the way each process affects relevant state. For example, if there is heavy investment in configuration then we expect the configuration state to be better than if there is little investment.

The architecture of the model is summarized in Figure 1

## 3 Simulation Results

Experiments exploring the states reached for each combination of the instrument for a 'benign threat environment (i.e., threat instruments all set low), and one 'high-risk environment (i.e., where the threat instruments all set high) have been performed. Each execution simulates a year of IAM activity, and we report the state recorded at the end of the year. The model was executed 100 times for each setting of the instruments, which is sufficient to bring the standard error to acceptable levels. Since there are $10 \times 10$ instrument combinations, the total number of executions for the two experiments was 20,000.

For each setting of the instruments we derived the expected number of breaches and business access, where

$$BusinessAccess \quad = \quad 1000 \times (nonaccess(cross))/(nonaccess(cross) + access(tick)) \qquad (2)$$

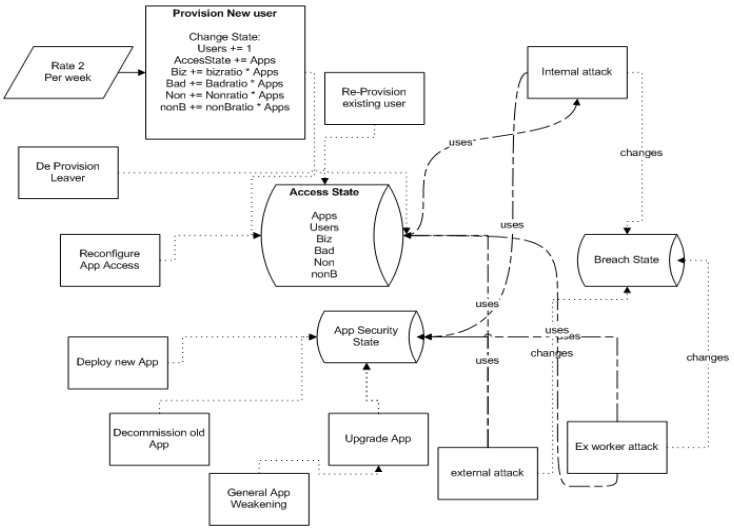


Figure 1: Basic Components of the IAM Model

For each setting of the instruments, we derive the expected number of incidents/breaches (IC) and denied business accesses ($A$), where

> IC = number-external-attacks-incidents + number-internal-attacks-incidents + number-ex-workers-attacks-incidents

$$A \quad = \quad 1000 \times (nonaccess(cross))/(nonaccess(cross) + access(tick) \tag{3}$$

The number of breaches/incidents (IC) is determined by modelling internal, external and ex-workers' attacks and the likelihood of success based on investments in 'enforcement'. $A$ is calculated by taking into account the proportion of 'denied accesses' to legitimate business users from the overall number of accesses.

More empirical work is required to establish an appropriate cost function for the investment Instruments ($K$), but to show how this could proceed we set the cost function as

$$K \quad = \quad 50 + 2^x + 1.8^y \tag{4}$$

where $x$ represents the enforcement instrument, and $y$ the configuration instrument. The (simple) intent here is to capture the exponential cost of achieving more with enforcement or configuration. The constant value represents the fact that there will be operational costs associated with these activities even if there is zero emphasis placed on them.

When these results are graphed they show (fitting with intuition) that it makes no sense to under-invest in either instrument as the breaches will be too large, but conversely also shows that over-inesting will be punished (by the prohibitive cost). It is also interesting to explore the various points of diminishing returns, e.g. when it makes no sense to invest in configuration, without first investing in enforcement.

# 4 Discussion

Equation 1, in the introduction, describes the utility of a generic dynamic model illustrating the chosen desired trade-off between confidentiality, availability, and investment. Figure 2 expresses the effect of IAM choices as the aggregated weight between incidents, access, and investment

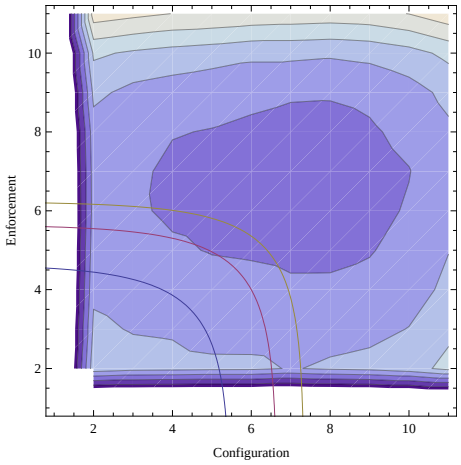


Figure 2: Iso-investments with iso-utilities

plotted against investments in configuration and enforcement. The corresponding dynamic utility model would be

$$U(C, A, K, t) \quad = \quad w_1(IC - \bar{IC})^2 + w_2(A - \bar{A})^2 + w_3(K - \bar{K})^2 \tag{5}$$

where $IC$, $P$, $K$ and $\bar{IC}$, $\bar{P}$, $\bar{K}$ represent, respectively, the actual and target levels of incidents, access and investment, and the $w_i$s express the weighted preferences given to variance from each of the targets; $t$ is time. Here the targets represent levels of the given quantities that are acceptable to the security manager in the context of the organizational priorities and policies. Generally, the aim of such policies is to minimize disruption of the business process whilst maintaining acceptable levels of security and cost.

More generally, each of the assumptions in the construction of the model, and each of the relationships explored through simulation, are approximating the actual relationships between $\bar{IC}$, $\bar{A}$, and $\bar{K}$, the incident count, productivity, and investment parameters in Equation 5. For example, the illustration that there are clear points where investing in one is severely handicapped without investment in the other, directly feeds the way we assume $\bar{K}$ affects both $\bar{IC}$ and $\bar{A}$.

We intend to continue to use the systems model of the IAM strategy, to guide parameterization of the dynamic model, and to explore further the relationship and role of these two styles of model based decision support.

**Acknowledgements**. We are grateful to Christos Ioannidis, Julian Williams, Brian Monahan, and Matthew Collinson for their advice on various aspects of this work.

# References

[1] R. Barro and D. Gordon. A Positive Theory of Monetary Policy in a Natural Rate Model. *Journal of Political Economy*, 91:589–610, 1983.

[2] A. Beautement, R. Coles, J. Griffin, C. Ioannidis, B. Monahan, D. Pym, A. Sasse, and M. Wonham. Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. In M. Eric Johnson, editor, *Managing Information Risk and the Economics of Security*. Springer, 2008. Preliminary version available in Proc. WEIS 2008: `http://weis2008.econinfosec.org/papers/Pym.pdf`.

[3] Y. Beres, J. Griffin, S. Shiu, M. Heitman, D. Markle, and P. Ventura. Analysing the performance of security solutions to reduce vulnerability exposure window. In *Proceedings of the 2008 Annual Computer Security Applications Conference*, pages 33–42. IEEE Computer Society Conference Publishing Services (CPS), 2008.

[4] M. Collinson, B. Monahan, and D. Pym. A Logical and Computational Theory of Located Resource. *Journal of Logic and Computation*, 2009. To appear (preprint available as HP Labs Technical Report HPL-2008-74R1).

[5] Demos2k. http://www.demos2k.org.

[6] M.P. Giannoni and M. Woodford. Optimal Interest-Rate Rules I: General Theory. Working Paper Series 9419, National Bureau of Economic Research, 2002. ISSU 9419, ISSN 0898-2937.

[7] C. Ioannidis, D. Pym, and J. Williams. Investments and trade-offs in the economics of information security. In Roger Dingledine and Philippe Golle, editors, *Proceedings of Financial Cryptography and Data Security '09 (to appear)*. Springer, 2009. Preprint available at `http://www.cs.bath.ac.uk/~pym/IoannidisPymWilliams-FC09.pdf`.

[8] R.A. Nobay and D.A. Peel. Optimal Discretionary Monetary Policy in a Model of Asymmetric Bank Preferences. *Economic Journal*, 113(489):657–665, 2003.

[9] Francisco J. Ruge-Murcia. Inflation targeting under asymmetric preferences. *Journal of Money, Credit, and Banking*, 35(5), 2003.

[10] M. Yearworth, B. Monahan, and D. Pym. Predictive modelling for security operations economics (extended abstract). In *Proc. I3P Workshop on the Economics of Securing the Information Infrastructure*, 2006. Proceedings at `http://wesii.econinfosec.org/workshop/`.