Beyond Risk-Based Access Control: Towards Incentive-Based Access Control

Debin Liu[1], Ninghui Li[2], XiaoFeng Wang[1], and L. Jean Camp[1]

[1] School of Informatics and Computing
Indiana University, Bloomington, Indiana, US
[2] Department of Computer Science
Purdue University, West Lafayette, Indiana, US

Abstract. In recent years, risk-based access control has been proposed as an alternative to traditional rigid access control models such as multi-level security and role-based access control. While these approaches make the risks associated with exceptional access accountable and encourage the users to take low-risk actions, they also create the disincentives for seeking necessary risky accesses. We introduce novel incentive mechanism based on Contract Theory. Another benefit of our approach is avoiding accurate estimate of the risk associated with each access. We demonstrate that Nash Equilibria can be achieved in which the user's optimal strategy is performing the risk-mitigation efforts to minimize her organization's risk, and conduct human-subject studies to empirically confirm the theoretical results.

Keywords: Insider Threat; Access Control; Risk Management; Incentive Engineering; Human-Subject Experiment.

1 Introduction

Access control is used pervasively for security and privacy protection in computer and information systems. Traditionally, access control policies are encoded as rigid rules. One example is the multi-level security (MLS) policy [1] used in the United States government. A user is assigned a security clearance and a set of compartments, a document is assigned a security level and a set of compartments, and access is allowed only if the clearance of the user dominates the level of the document and the user's set of compartments is a superset of the document's. A 2004 report by the JASON defense advisory group [2] discussed many of the shortcomings of such policies, which are often too restrictive to meet modern needs.

Risk-based access control. Several proposals for risk-based access control [2–5] have been proposed to address these limitations. In these approaches, initially each user is assigned certain amount of risk tokens, called risk budget. For each access, the access control system estimates the amount of risk for that access, and the user has to pay in the form of risk tokens in order to gain access. The user's aggregated risk is controlled by the risk budget, and the user is incentivized to choose low-risk accesses because it costs less risk tokens and the surplus of the budget brings her benefits. Such an approach, however, suffers from several limitations. First, it unintentionally discourages the accesses that involve spending risk tokens, as the user has an incentive to avoid accesses and save her tokens for personal gains. Essentially, this creates a situation in which the user has a conflict of interest between personal rewards and accesses necessary for better completing her job, and as a result, discourages necessary risk-taking behaviors. This goes against the intention of enabling more accesses. Second, this places an additional burden for users, as it creates a new critical resource, risk budgets, the users need to manage for fulfilling their duties. They have to keep track of the amount of risk tokens they have, and be constantly debating whether to save them or spend them. Finally, this approach requires as a first step, a relatively precise quantification of the risk of accesses, which is notoriously difficult.

Incentive-based access control. We believe that these problems can be mitigated if the user is willing to use exceptional accesses in the best interest of her organization. In this paper, we propose incentive-based access control (IBAC) as a solution. Our idea is to separate the two objectives of limiting aggregated risks and incentivizing users to adopt risk-mitigation methods to reduce risk, and achieve them through different mechanisms.

2 Overview of Incentive-Based Access Control

We had three goals while designing our incentive-based access control. First, it should provide access flexibility to respond to unforeseeable situations by allowing exceptional accesses to bypass controls. Second, it should enable the organization to manage aggregated risks, and identify not only potentially malicious insiders but also those who are simply risk-seeking. Third, users should face consistent incentives to adopt risk-mitigation efforts reducing organizational risk whenever feasible.

2.1 Basic Concepts

Exceptional access When a user requests an access, the organization first roughly estimates the risk, k, for that request depend on the users' profile, such as security clearance and roles, where applicable. Based on the quantified estimate, the access request falls into one of the following categories: "*allowed access*" when the risk is low enough; "*denied access*" when the risk is too high; or "*exceptional access*" when the risk is in-between.

Risk budget Organization set risk budget for each individual employee. Risk budget enables its owner to obtain exceptional access to override the control when necessary, at the expense of risk points. This provides the flexibility an access model needs to work in a practical environment. On the other hand, risk budget restricts the number of exceptional accesses and prevents exception abuse.

Risk-mitigation efforts We assume that for each exceptional access, there is a set E of risk mitigation methods available for a user to adopt. Different $e \in E$ may represent different access method; they may also represent different additional technical and nontechnical efforts for mitigating risks. In the rest of this paper, we use the term "access method" and the term "risk-mitigation effort" interchangeably. For ease of discussion, we assume that a special method $e_0 \in E$ represents the case that user takes an exceptional access without making any risk-mitigation effort, implicitly the highest resulted risk.

Contract-based rewards We consider two types of incentive contracts to provide users rewards and align their interests with their organization. A contract of the first type is designed for the situations where the organization can observe and verify the risk-mitigation effort e taken by the user. The contract specifies a *reward* function $r(e)$ which determines the reward that the user will get for choosing e. The second type of contracts are used when the organization is unable to determine the user's risk-mitigation effort. This happens when the effort cannot be directly observed by the organization, or is too expensive to monitor. Instead, such a contract is specified over a set of observable outcomes T that indirectly reflect the effort e the user takes, and therefore called *consequence-based contracts*.

Personal costs Given an exceptional access a, each risk-mitigation effort e will incur some cost for the user; this is determined by the cost function $c(e)$ which reflects the expenses incurred in terms of usage of storage/computational resources, opportunities, interference in the normal work flow, and others. We assume that $c(e_0) = 0$. In this paper, we measure the costs of risks, personal costs of effort and incentive rewards using the same metric, the *IBAC token*. We use the domain of real numbers R to denote this domain. Therefore the risk function is denoted as $k : T \to R$; the cost function is denoted as $c : E \to R$; and the reward function is denoted as $r : E \to R$.

2.2 Putting Things Together

When a user requests an access, IBAC makes access decision to allow any low-risk request and deny any request that has too high a risk. If the risk is in-between, that request is considered as an exceptional access and priced at the risk estimate. User can obtain an exceptional access only if his risk budget is sufficient to cover the expense of risk points for the price. The risk budget deduction is conducted spontaneously without any requirement of human response.

Once an exceptional access is granted, IBAC will launch an contract-based mechanism to incentivize users to perform risk-mitigation efforts. When multiple risk-mitigation efforts exist, a discrepancy of incentives emerges between the organization and the employee: the former prefers the effort that minimizes its risk, while the latter wants to choose what minimizes her cost. This problem is tackled by the incentive contracts.

3 Effort-based Contract

Contract game. When the user's effort is observable, the contract is determined by a reward function $r : E \rightarrow R$ that specifies the reward the organization offers to the user based on her observable or verifiable risk-mitigation effort. Under a contract r, the user attempts to choose the mitigation effort e that incurs the minimum amount of total personal cost and the maximum amount of the reward, by solving the following optimization problem:

$$\min_e [c(e) - r(e)]$$

On the other hand, the organization seeks a low risk and a low reward associated with an access. Thus, its combined optimization goal is described as follows:

$$\min_r [k(e) + r(e)]$$

Here we call $c(e) - r(e)$ an *adjusted cost* for the user and $k(e) + r(e)$ an *adjusted risk* for the organization.

From the above equations, it is evident that none of these parties can unilaterally achieve its optimization objective: the user's selection of e depends on the organization's choice of r, and vice versa. This essentially forms a game, which we call *contract game*. Solving this pair of optimization problems gives us one or more *Nash Equilibria*[3] in which neither player (the user or the organization) has an incentive to deviate from its equilibrium strategy (choice of e or r), as this does not gives it a better payoff (a lower outcome for its optimization target). Here we denote an equilibrium contract by r^* and an equilibrium effort by e^*.

3.1 Cooperative user and Deterministic Cost

We assume users are *cooperative* and *rational*. A cooperative rational user will choose among all e's such that $c(e) - r(e)$ is minimum, and when there exists more than one such e, choose the one that results in the least perceived risk: i.e., $k(e)$ is smallest. Again, such a perception can be inaccurate, based upon, for example, some qualitative concepts such as high or low risks.

Let E be the set of risk-mitigation efforts for an access. The organization constructs its contract $r(e)$ in the following way: for $e^* = \arg\min_{e \in E} [k(e) + c(e)]$, we have $r(e^*) = c(e^*)$; for other $e \in E$ except e_0, $r(e)$ can be set to any value in $[0; c(e))$. We call this contract *compensation contract*, as it makes up for the user's cost for choosing the effort in the best interest of the organization, i.e., e^*. Here we show that this contract is optimal.

Theorem 1. *The compensation contract and selection of e^* form a Nash Equilibrium in the contract game.*

A Nash Equilibrium is a strategy pair in which neither party (the organization or the user) can be better off by switching to another strategy, given that the other party sticks to its equilibrium strategy. This outcome is optimal in the sense that when the users are cooperative and rational, the equilibrium strategy (i.e, the compensation contract) achieves the lowest adjusted risk $k(e) + r(e)$ for the organization. We prove this result as follows.

Proof. Let us first show that under the compensation contract, a cooperative and rational user will choose e^*. The adjusted cost of $e \neq e^*$ is greater than $c(e_0)$, which is zero. Comparing it with the adjusted cost of e^*, we have: $c(e) - r(e) \geq 0 = c(e^*) - r(e^*)$. Note that the equality holds only when $e = e_0$, which makes $c(e) = 0$. When this happens, a cooperative user will still choose e^*, which incurs a lower risk, according to its definition.

Consider a contract under which a cooperative and rational user chooses e' and gets a reward x. This only happens when e' is no less attractive than e_0, that is $c(e') - x \leq c(e_0) = 0$. Therefore, $x \geq c(e')$. For the organization, this choice gives it an adjusted risk at least $k(e') + c(e')$, which is no better than what is offered by the compensation contract, according to its definition.

[3] The contract game is actually sequential in which the organization moves first and the user moves next, and the Nash Equilibrium we describes throughout the paper all refers to the sub-game perfect Nash Equilibrium [6].

We note that the best option from the organization's point of view minimizes the combined cost of the organization and the user $k(e) + c(e)$, rather than simply the organization's risk $k(e)$. Hence even from minimizing the organization's adjusted risk, the best option maximizes the combined social welfare.

Example 1. Consider a bank that utilizes the effort-based contract, which encourages its tellers to reduce organziation's risk when accessing sensitive customer information. For a read access, we assume the access-control mechanism offers the options with or without the copy-paste function. The one with it incurs a risk at 1000 IBAC tokens while the one without incurs a risk at 200 tokens. On the other hand, the operational cost the teller has to undertake in the absence of the function is quantified as 20 tokens. This example is described in table 1.

Table 1. Example 1

Effort	$r(e)$	$c(e)$	$k(e)$	$(c-r)$	$(k+r)$
e_0	0	0	1000	0	1000
e_1	r	20	200	$20-r$	$200+r$

Organization computes $k(e_0) + c(e_0)$ to be 1000, and $k(e_1) + c(e_1)$ to be 220, and determines that e_1 is preferred, and set $r(e_1) = 20$, and $r(e_0) = 0$. The compensation contract the bank offers to the teller becomes: *"You can receive 20 IBAC tokens as reward if you access the database without copy-paste functionality. Do you want to accept this offer?"*

3.2 Considering Non-deterministic Cost

The deterministic model above assumes that $c(e)$ is a constant across all users and observable to the organization, which has several limitations. First, different users may have different personal costs. Second, under different circumstances, the same person's cost may vary. Third, a user may not be able to accurately estimate her cost, which leads to the seemingly randomness in effort selection, and alternatively, the organization's estimate of her cost can be inaccurate.

We model all these by using a random cost $c(e)$. Specifically, we assume that $c(e)$ is a random variable with a Gaussian distribution. Such a distribution can be estimated in practice through performing public surveys or sampling. In our model, an individual is viewed as being randomly drawn from a population that produces that cost distribution. Her personal cost, therefore, is also treated as a random draw from the distribution.

The user's objective, again, is to minimize her adjusted cost. The organization, however, cannot observe her specific cost, but the distribution of the cost, and thus intends to minimize the expected adjusted risk, as follows:

$$\sum_e \Pr[e \text{ is chosen}](k(e) + r(e))$$

Due to the uncertainty of the cost, design of an optimal contract becomes very complicated. For simplicity, here we only consider two risk-mitigation efforts: the low effort e_0 and the high effort e_1. Since the low-effort e_0 represents no risk mitigation effort, we assume that $c(e_0)$ is constant at 0, and $c(e_1)$ follows a Gaussian distribution $N(\ ;\ ^2)$.

Because changing both $r(e_0)$ and $r(e_1)$ by the same amount does not affect the user's incentive to choose e_1, and the organization would like to minimize r. The effort-based contract should set $r(e_0) = 0$, and we need to determine $r(e_1)$.

The user chooses e_0 whenever the following inequality holds:

$$c(e_1) - r(e_1) \geq c(e_0) - r(e_0) = 0$$

As $c(e_1)$ follows a Gaussian distribution $N(\ ;\ ^2)$, this happens with the following probability:

$$p = 1 - \quad \left(\frac{r -}{\quad}\right);$$

where is the cumulative distribution function (cdf) of the standard Gaussian distribution.

The organization intends to minimize the adjusted cost, and hence would set $r(e_1)$ as follows:

$$r^*(e_1) = \arg\min_r [p \cdot k(e_0) + (1-p) \cdot (k(e_1) + r)]$$
$$\Rightarrow \arg\min_r \left[\left(1 - \left(\frac{r - }{\quad}\right)\right)(k(e_0) - k(e_1) - r) + r \right]$$

The above can be solved using standard function minimization techniques. We note that when the risk difference between e_0 and e_1, i.e., $k(e_0) - k(e_1)$, is large, the reward for taking e_1 should also grow, to reduce the probability p. However, the value $r(e_1)$ is not sensitive to inaccuracies of estimating $k(e)$.

Theorem 2. *The above contract and the user's choice of e minimizing its adjusted cost forms a Bayesian Nash Equilibrium in the contract game.*

In a Bayesian Nash Equilibrium [6], the strategy of each player (the organization or the user) maximizes its expected payoff against the other's strategy, given its belief about the other player (probability p of choosing e_0). The correctness of the proof directly follows the above analysis.

Example 2. In Example 1, consider that an access without the copy-and-paste function incurs a probabilistic cost to the user, which follows a Gaussian distribution with a mean of 20 and a stand deviation of 10. Then, the optimal contract is built as follows:

$$r^*(e_1) = \min_r \left[\left(1 - \left(\frac{r-20}{10}\right)\right)(1000 - 200 - r) + r \right]$$

where is the cdf of a standard normal distribution.

Using the Matlab built-in function of *fminbnd*, r^* is found to be 46. The effort-based contract for this example is then constructed as: *"You can receive 46 IBAC tokens as reward if you access the database without copy-paste functionality. Do you want to accept this offer?"*

4 Consequence-based Contract

Effort-based contract can be used only when it is possible for the organization to observe the action taken by the user. We now show that even when the organization is unable to observe the actions, the organization can nonetheless compute an optimal effort-based contract and translate that into a consequence-based contract.

We assume that there is a set T of possible outcomes, and the organization has a damage estimation function $k : T \to R$, that maps each outcome to a quantitative estimate of damage. A consequence-based contract is given by a function $r : T \to R$, which gives the reward the user gets for each of possible outcome. Given such a contract r, the user will choose an action that minimizes her expected cost. The user's cost for each action $e \in E$ if the the personal cost $c(e)$, minus the expect reward the user gets. Let $T = \{t_1; t_2; \cdots; t_n\}$. Suppose that the probability each t_i when taken action e is p_i^e, then the user's adjusted cost for action e is given by:

$$uac(e) = c(e) - \sum_{1 \le i \le n} p_i(e) r(t_i)$$

The user will choose an e that minimizes the above cost function. The organization needs to generate a contract such that the user's optimal choice will minimizes its expected cost as well. When a user chooses action e, the organization's adjusted cost is given by

$$oac(e) = \sum_{1 \le i \le n} p_i(e) k(t_i) + \sum_{1 \le i \le n} p_i(e) r(t_i)$$

We note that if we use $k(e)$ to denote the risk to the organization when the user takes action e, then

$$k(e) = \sum_{1 \le i \le n} p_i(e) k(t_i)$$

We further note that for every consequence-based contract, one could construct an equivalent effort-based contract, by computing

$$r'(e) = \sum_{1 \le i \le n} p_i(e) r(t_i);$$

and setting

$$r(e) = r'(e) - min(r'(e)):$$

The goal of the organization is to choose $r(t_i)$ such that the same e minimizes both $uac(e) = c(e) + r(e)$ and $oac(e) = k(e) - r(e)$, which is exactly the same as in the effort-based case. The only difference is that in the effort-based case, the organization could directly set $r(e)$, but in the consequence-based case, the organization can only *indirectly* set $r(e)$ by choosing $r(t_i)$'s.

Hence if there exists a consequence-based contract C such that its equivalent effort-based contract is optimal for the effort-based case, then C is the optimal contract for the consequence-based contract case as well. Suppose that C is not, and another contract C' gives better utility for the organization, then its corresponding effort-based contract must also be better, contracting the assumption that C's corresponding effort-based contract is optimal.

Given an optimal effort-based contract r, one can try to solve for the corresponding consequence-based contract by first solving the following systems of linear equations to compute $r'(t_i)$, which may be negative:

$$r(e_j) = \sum_{1 \le i \le n} p_i(e_j) r'(t_i):$$

And then set

$$r(t_i) = r'(t_i) - \min\left(r'(t_i)\right):$$

Example 3. An employee wants to download a file from the Internet. His browser can be adjusted to high security setting or low security setting. We assume that the browser security setting upgrade has a cost of 20 IBAC tokens to the user, and a drive-by-download attack will cause the organization a loss of 1000 tokens. The system cannot detect what security level the user's browser has. Assuming the statistics that a high security setting browser can prevent 90% of drive-by download, and a low security setting browser can only prevent 40% of the same attack, this example is presented in the table 2.

Table 2. Thrid Example

	Consequence k		$c + r$
	attacked	unattacked	
Low setting	60%	40%	$60\% r(a) + 40\% r(u)$
High setting	10%	90%	$20 + 10\% r(a) + 90\% r(u)$
$r(k)$	$r(a)$	$r(u)$	

From previous section, we have the optimal effort-based contract is

$$r^*(e_0) = 0;$$
$$r^*(e_1) = 20:$$

We can construct the corresponding consequence-based contract by solving the following linear equations:

$$r^*(e_0) = 0 = 60\% r'(a) + 40\% r'(u);$$
$$r^*(e_1) = 20 = 10\% r'(a) + 90\% r'(u):$$

We further set $r(t_i) = r'(t_i) - \min\left(r'(t_i)\right)$, and provides the employee the following contract: *"You can receive 40 IBAC tokens if there is no virus detected after your file is downloaded. Upgrade your browser setting to high security level will greatly reduce the probability of being infected."*

5 Human-Subject Evaluation

We performed a human-subject experiment to evaluate the efficacy of our access control model. The main goal was to understand how users, who may not be perfectly rational, choose their access choices under our contract-based incentive mechanism. We are also interested in understanding how well our model helps suppress the risk the organization undertakes.

5.1 Subject Recruitment

We recruited 36 volunteers. More than 90% of the participants use computers more than 5 hours per day. An interesting finding from their background survey is their different attitudes towards the security protections for their organization's and their own computing systems: 61% of the participants chose to scan their personal computers immediately upone seeing a virus warning, while only 52% of them did so to their organization's computers. This echoes our hypothesis about the existing misalignment between employees' incentives and their organizations' interests.

5.2 Experimental Design

The participants were randomly and equally divided into three groups. Group 1 was treated as the benchmark. Group 2 and 3 worked under IBAC. All groups were given the same tasks of sending ten documents, each of which was attached to a different email. They were told that with a certain probability, these emails could be intercepted by untrusted parties. In order to reduce the risk of information leaks, we suggested, but did not require, that they encrypt the emails or the documents, or both. The participants did not have any obligation to make any of these mitigation efforts. We treated encrypting both email and document as the high effort (Level 3), encrypting only the document as the medium high effort (Level 2), encrypting only the outgoing email as the medium low effort (Level 1), and no encryption as the low effort (Level 0).

Document classifications. These documents were classified into four categories: Secret, Confidential, Restricted, and Unclassified. The risk associated with the document in one category was deemed as one order of magnitude higher than that of a document in its immediate lower category. For example, a secret document was assigned a risk magnitude of 1000, and a confidential one was given 100, and so on. Among these ten documents, two were unclassified, three restricted, two confidential, and three secret.

Experiment settings. In the experiment, Members in Group 2 and 3 need to interact with IBAC model: the budget-based mechanism kept records of their risk points deduction, and the contract-based mechanism incentivized them to select a risk-mitigation effort. Their effort choices determined their reward. In order to make the compensation scheme easy to understand, we combined the budget-based control and the contract-based incentive mechanism together, and set the set of interactions as a scenario of purchasing: we first gave each of them 1,000 points. They had to pay a price using their points for sending a document. The leftover points were later reimbursed as every 300 points for $5. Therefore, the prices must be carefully designed so that the rewards generated follow an effort-based contract for Group 2, and a consequence-based contract for Group 3.

Before starting the experiment, we gave the participants an excercise of encrypting email and document. We took that chance to measure the time each of them used. Given the time factor in determining compensation scheme, we translate the obtained time measurement into that participant's personal cost of mitigation effort. For example, we pay Alice $10 compensation if she completed the experiment in 10 minutes. Thus, if Carol needs 1 minute to perform a document encryption, we set $c_{encryption} = \$1 = 60$ points as Carol's personal cost of encrypting document.

The participants in Group 1 (the benchmark) always received compensation when they completed the tasks. The amount of the compensation was determined solely by the time they spent on the tasks: the less time was used, the higher it became. The participants in the other two groups got compensation, which also depended on time, plus rewards: the more risk points they spent, the less rewards they got.

5.3 Reducing Organizational Risks

Figure 1(a) shows the average risks incurred by each group. Both incentive contracts helped reduce the organization's risk exposure. Particularly when secret documents were transmitted, our contracts cut the risk by 30%.

5.4 Encouraging Risk-mitigation Effort

To understand how well our model could incentivize the individuals to perform risk-mitigation efforts, we first sorted the tasks from the lowest rank to the highest one, i.e., unclassified to secret, and then calculated the average risk-mitigation effort level for each task and each group. Figure 1(b) presents the results, in

which red dots represent the average levels chosen by the benchmark group for all ten tasks. The green dots describe these by Group 2 (effort-based contract) and the blue dots by Group 3 (consequence-based contract). We can see that both contracts successfully encouraged users to improve their effort levels, which is in stark contrast to those in the benchmark group. We also performed paired t-statistical tests for the effort choices in these three groups. Paired with the benchmark group, Group 2 and 3 generated t-values of 3.72 and 7.93 respectively. The outcomes indicate that both groups made statistically significant improvements on risk-mitigation effort, with a 99% confidence level.

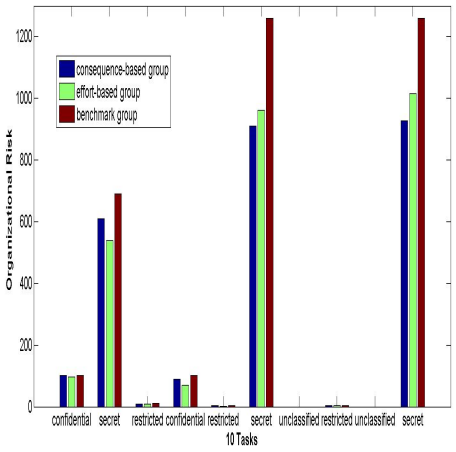

(a) Organization's Risk Postures

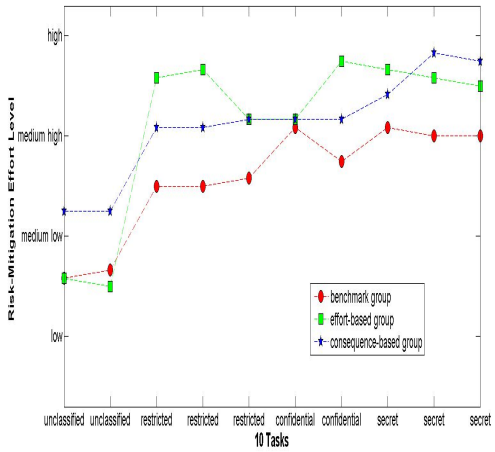

(b) Average Personal Risk Control Effort Levels

6 Conclusions

In this paper, we designed a new incentive-based access control mechanism that encourages the users to make necessary accesses, while discouraging them from taking unnecessary risks. This has been achieved through a novel contract-based incentive mechanism that rewards the users for the access that is in the best interest of their organization. We analyzed the IBAC mechanism using game theory and identified the optimal contracts that motivate both the organization and the users to play Nash-Equilibrium strategies. We also performed human-subject experiments that demonstrate the effectiveness of our approach in mitigating the organization's risk.

Acknowledgements

This work is supported in part by CNS-0716292.

References

1. D. Elliott Bell and Leonard J. LaPadula. Secure computer systems: Unified exposition and Multics interpretation. Technical Report ESD-TR-75-306, Mitre Corporation, March 1976.
2. MITRE Corporation. Horizontal integration: Broader access models for realizing information dominace. Technical Report JSR-04-132, JASON Defense Advisory Panel Reports, 2004.
3. I. Molloy, P. Cheng, and P. Rohatgi. Trading in risk: Using markets to improve access control. In *New Security Paradigms Workshop*, Olympic, California, September 2008. Applied Computer Security Associates.
4. A. Yemini, D. Dailianas, Florissi, and G. Huberman. Marketnet: Market-based protection of information systems. In *The 12th Int. Symp. on Dynamic Games and Applications*, 2006.
5. Debin Liu, XiaoFeng Wang, and L. Jean Camp. Mitigating inadvertent insider threats with incentives. In *The proceeding of Financial Cryptography and Data Security*, February 2009.
6. Martin J. Osborne and Ariel Rubenstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.