RC24439 (W0711-202) November 28, 2007
Computer Science

IBM Research Report

Trading in Risk: Using Markets to Improve Access Control

Ian Molloy*, Pau-Chen Cheng, Pankaj Rohatgi
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

*and Purdue University

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Trading in Risk: Using Markets to Improve Access Control[*]

Ian Molloy
IBM T.J. Watson Research Center and Purdue University
immolloy@us.ibm.com

Pau–Chen Cheng   Pankaj Rohatgi
IBM T.J. Watson Research Center
pau@us.ibm.com   rohatgi@us.ibm.com

March 21, 2008

Abstract

With the increasing need to securely share information, current access control systems are proving too inflexible and difficult to adapt. Recent work on risk-based access control systems has shown promise at resolving the inadequacies of traditional access control systems, and promise to increase information sharing and security. We consider some of the core open problems in risk-based access control systems, namely where and how much risk to take. We propose the use of market mechanisms to determine an organization's risk tolerance and allocation. We show that with the correct incentives, an employee will make optimal choices for the organization. We also comment on how the market can be used to ensure employees behave honestly and detect those who are malicious. Through simulations, we empirically show the advantage of risk-based access control systems and market mechanisms at increasing information sharing and security.

1 Introduction

Current access control systems, especially in the intelligence community, are highly restrictive and inflexible. New access control models are constantly being developed, and each is designed to more closely resemble the practical limitations an organization may wish to place on their data, and simplify the specification and verification of policies for the system. Mechanisms that are too coarse do not have the expressive power to distinguish between some transactions that should be allowed, and those that should be denied. They will either introduce security holes or stifle progress, possibly forcing employees to work outside of the system, thus decreasing security [14]. Prudent security is the delicate balance between minimizing risks and maximizing benefit within the constraints of the organization.

We argue that any access control system is an attempt to model the organization's notion of *risk*; the more fine grained our access control systems become, the tighter we bind on the organization's unique notion of risk. This comes at the cost of complexity in designing, verifying, executing, implementing, and specifying the underlying mechanisms and policies for the data. A benefit- and risk-based access control system would directly address the goal of any access control system: manage risk of access to sensitive data.

We argue that a bounded laissez-faire system of access control modeled and implemented as a risk market is beneficial to traditional restrictive and rigid access control systems. The central issue in a risk–based access control system is to determine where and how much risk to take. In other words, it is a *risk allocation problem* where risk is being treated as a limited resource. It is well–known in the realm of Economics that

[*]This Research is conducted through participation in the International Technology Alliance sponsored by the U.S. Army Research Laboratory and U.K. Ministry of Defense and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

a properly set–up market tends to allocate resources in an optimal way [24][10]. We validate this argument by simulating a risk market and other risk allocation methods, such as centralized pre–allocation. We also simulate a multilevel security (MLS) access control system based on the Bell–LaPadula model. The results of our simulation show that the risk market outperforms other risk allocation methods in terms of delivering the best risk vs. benefit tradeoff when risk–taking is capped by an organization's risk tolerance. The risk market also outperforms the MLS system not only by delivering much better risk vs. benefit tradeoff but also by staying within an organization's risk tolerance when the MLS system has no notion of aggregated risk, let alone staying within any risk tolerance.

We also argue that a risk market can be made resilient against different kinds of attacks, such as collusion among malicious participants or espionage by employees with outside funding, by providing proper incentives for good behavior. A risk market also makes it easier to detect malicious behaviors since all participants must go through the market to access resources and their trading behaviors are logged. The participants' access to an organization's resources can also be logged. Analyzing these logs would uncover irrational and abnormal behaviors which could lead to identification of malicious participants.

Figure 1 provides an overview of an access control decision, illustrating where and how the risk market may impact a transaction. An employee will (1) make a request for information to the system, authenticating themselves and indicating what information they wish to access and how they would like to access it. The Quantified Risk Based Access Control System will evaluate the request based on known information, such as previous access history and behavior logs, and quantify the risk associated with the access. It will then provide the user with an access ticket (2) describing the access, and indicating the request price in risk tokens. In some instances the price may be infinite, indicating a `deny` decision has been made.

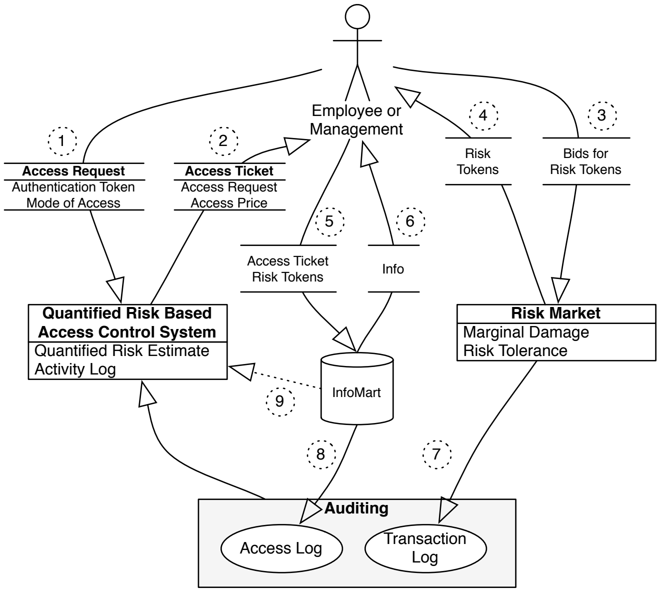

Figure 1: Overview of How the Market Fits into an Organization's Access Control

The user will evaluate the price. If they have enough risk tokens and agree to the price, they will purchase access to the information (5) from the InfoMart (database or other information storage mechanisms) which will return the information requested in the Access Ticket (6). If the price is too high, the user may reissue the request, possibly indicating a less risky method of access (such as soft-copy on instead of hard-copy). If the user does not have enough risk tokens, they may place bids on the risk market (3) and if successful, purchase additional tokens (4). They then proceed to step (5). When fulfilling the request, the InfoMart should ensure the Access Ticket is still valid and the price has not changed (9). To facilitate auditing, detect irrational or malicious behavior, and ensure more accurate risk calculations, the Risk Market will maintain transaction logs (7), and the InfoMart will maintain access logs (8).

Appendix C gives an example of using risk–based access control in the banking industry.

This rest of this paper is organized as follows: Section 2 discusses related work and the open problems we address. Section 3 highlights skepticism of risk–based access control systems and how we plan to address them. Section 4 provides a background on auction theory and describes the risk market, and Section 5 discusses the attacks on risk market and defenses to these attacks. In Section 6 we describe our simulation and experimental results. We conclude in Section 7.

2 Related Work

There have been several works proposing risk-based access control systems. The MITRE Jason Report [14] presents a history of *risk*, the probability of loss or damage in MLS systems in government settings, and identifies the many flaws, shortcomings, restrictions, and incompatibilities with the current classification system. They propose a three-phase solution based on risk. Phase-zero quantifies risk, phase-one places restrictions on the maximum amount of risk the organization is willing to assume for any given document, and phase-two uses the quantified risk from phases zero and one and allows the organization to bound the aggregate amount of *harm*, or expected damage, expended within the entire organization. Each transaction (a subject trying to access information via a given method) is assigned a cost in units of harm that is charged against the subject accessing the information. By placing bounds on the amount of risk in tokens that are generated, the organization can limit the amount of harm they are willing to assume.

A second work, Fuzzy MLS [6], calculates risk values for transactions based on standard MLS labels. Risks are composed of two components: *temptation* represented by the difference in classification and clearance labels between the subject and object, and *inadvertent disclosure* or *slip of the tongue*, represented by the difference in compartment membership between the subject and object.

There exists a soft boundary, below which all transactions are allowed, and a hard boundary, above which all transactions are denied. Between the hard and the soft boundaries, risk mitigation mechanisms are used to reduce the risk the organization must assume such as requiring the user to pay for the transaction similar to [14].

A third work from Zhang et al. [28] describes a benefit- and risk-based access control system where a static subset of transactions are allowed. Benefit and risk values are multidimensional vectors representing attributes, such as currency, intellectual property, physical property, or human life. The allowed transactions satisfy a weak-optimal (Pareto-optimal) state where the aggregate benefit outweighs the aggregate risk for each component of the vector. The state is largely static, allowing an agent to select an alternate weak-optimal subgraph and expend an amount of risk capital to perform the reorganization. It is an intractable problem to update the state, and the system provides no guarantees that a given sequence of accesses will result in a net gain.

2.1 Open Questions and Possible Solutions

All of the above risk-based access control systems use risk tokens as a main enforcement mechanism. Both [14] and [6] propose mechanism and strategies to distribute risk tokens, but neither provides details or analysis on how this should be done. We analyze how well token-based risk-based, access control systems perform in general, and attempt to answer several questions that are pivotal before they may be successfully deployed. First, how well do they perform compared to traditional access control systems? Second, how much risk should an organization expend, i.e., how should an organization quantitatively determine their risk tolerance? Third, how should the risk be distributed within the organization to maximize the organizations expected utility, i.e. how and where risks should be taken? Finally, how successful are hard and soft boundaries as risk mitigation measures, and how should these boundaries be determined?

The main research question we wish to answer is how to determine an efficient risk distribution method. We see the remaining questions as tangential, and their answers will depend on how we distribute risk. If risk can be quantified based on the models in Section 2, then risk must be a limited resource, and this is a resource allocation problem. Resource allocation is the main focus of microeconomics, and we consider some of the possible solutions in that area. We have determined several methods an organization may choose to use to distribute risk.

1. *Static Preallocation* - Management makes risk allocation decisions across the organization.

2. *Request Based* - Employees request risk tokens, and management determines if the requests should be filled.
3. *Hierarchical Allocation* - Risk tokens are allocated at the top of the organizational hierarchy, and reallocated by organizational units down the hierarchy until it reaches the end user.
4. *Market Mechanism* - Markets, such as auctions, allow employees to buy and sell risk tokens openly.

[14] proposed a combination of request based and hierarchical allocation methods, while mentioning the possibility of using markets and currency exchanges. We shall argue that market mechanisms and economies result in the most efficient allocations, and answer many of the tangential questions. In Section 4 we analyze how the market should be structured to obtain a good resource allocation.

2.2 Markets in Other Settings

While the proposal to use market economies of risk and damage seems outlandish, the concept is not foreign in other areas. Market mechanisms have been used successfully in areas such as the FCC auction of the electromagnetic spectrum [1], legistations for limiting greenhouse gases in the US [2] and UK [9], and CO_2 quota allotments by BP [17]. More details and examples are given in Appendix A.

3 Skepticism

Many readers may be skeptical about the concept of risk-based access control systems, especially using a market as the main distribution mechanism of access tokens. These mistrusts can be divided into three categories: risk calculation measurement, proofs of security, and mistrust in using markets for security applications.

1. *Risk Calculation Measurements* An often cited source against risk-based access control is Cybenko's "Why Johnny Can't Evaluate Security Risk" [8]. This is not a problem unique to risk-based access control systems; all access control systems are based on the same problem of guessing what is and what is not a risky transaction, and is identical in all access control systems: Bell-LaPadula, RBAC, Chinese Wall, DAC, and ORICON, etc. What is unique about risk-based access control systems is they make explicit use of the *educated guesses* that are a part of all access control systems; risk calculation enables counting risk in units of tokens, and therefore placing a numerical bound on the amount of damage by limiting the total number of risk tokens. The bound will be there even if the risk calculations are not accurate, but inaccuracy is inherent in all access control systems.
2. *Safety Analysis* Other critiques stem from a lack of safety analysis compared to more traditional access control systems. In a risk-based access control system, it does not make sense to ask the question "Is principal X allowed to do Y?" since, unless the cost of the transaction exceeds the risk budget, the answer is an unequivocal "yes," or at best "probably." This allows all users access to all resources, implying undesirable privilege escalation and causing many to naively assume unbounded escalation. Unlike traditional access control systems, it places bounds on aggregate access, providing a different safety analysis question; one which traditional access control systems are incapable of answering. Risk-based access control systems are also open to new attacks such as denial of service due to insufficient or inadequately allocated risk budgets.
3. *Mistrust in the Market* Due to the potential for denial of service and privilege escalation, risk allocation becomes an increasingly important, and difficult question that must be resolved. Since it is so pivotal to the security of the system, it seems naive to trust the task to a market. The risk market opens the system to increased collusion, and additional privilege escalation, becomes an attractive attack vector for espionage (employees paid from outside the system), and denial of service due to employees hoarding risk and refusing to sell, gaming the market for personal gain without performing constructive work, and inefficiencies due to employee incompetence.

We will discuss the implications and resolutions of all the above skepticism throughout the paper, and attempt to place empirical bounds when possible. Many concerns are solved by the correct incentives while others, such as accurate risk measurements, remain open problems and affect all aspects of information security.

4 Auctions and the Risk Market

4.1 Auction Theory

Adam Smith's *invisible hand* is the idea that in a free market, an individual acting selfishly will benefit the community as a whole [23]. The first fundamental theorem of welfare economics can be seen as a formalization of this idea [20]. It states that under ideal conditions, a competitive market will converge towards an efficient allocation of resources. These ideal conditions are known as *perfect competition*, where the buying decisions of agents cannot affect the market price. The second fundamental theorem states that by performing a lump sum redistribution and allowing the market to take over will result in a *Pareto-optimal* allocation. Any redistribution from a Pareto-optimal allocation that is beneficial to one individual is detrimental to at least one other [20].

The behavior of a market is determined by the relationship of buyers and sellers. A seller's willingness to sell at a given price is governed by their *marginal cost*, the change in total cost (TC) when the quantity produced (Q) changes by one unit,

$$MC = \frac{\partial TC}{\partial Q} \tag{1}$$

and dictates the price they are willing to accept for their goods at a given quantity. Similarly, buyers define a demand curve D, based on the *marginal benefit* (MB), dictating the price they are willing to pay for a given quantity. A supply and demand graph can be seen in Figure 2.

The point where the supply and demand curves cross is called the equilibrium, and dictates both the *equilibrium price* and the *equilibrium quantity* that can be exchanged in the market. Beyond this quantity, no sales are possible if all individuals are rational. A trader that is *individually rational* will always agree to a contract when beneficial given their current information. We can also observe that by increasing the marginal cost function we shift the equilibrium quantity to the left (and the price up). The equilibrium can also be changed by changing the buyer's demand. The gray area represents aggregate net gain that can be observed through a redistribution of goods using the market. For the marginal cost function $MC(\cdot)$, and demand curve $D(\cdot)$, the equilibrium quantity is Q^*, while if we increase the marginal cost to $MC'(\cdot)$, the equilibrium quantity decreases to Q'. A *market's efficiency* is a measure of the extracted value or utility versus an optimal allocation.

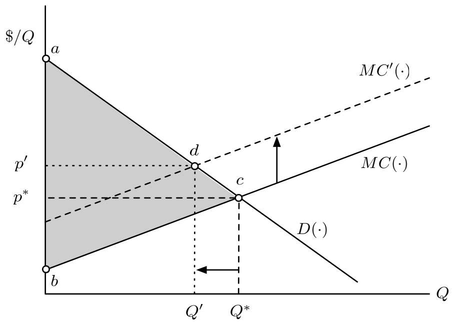

Figure 2: Supply and demand determine the equilibrium quantity Q^* and the equilibrium price p^*.

While the theorems of welfare economics indicate that free markets are capable of solving the risk-distribution problem, [16] proves that they are the only Pareto-optimal, individually rational, dominant strategy solution given imperfect, incomplete information. A *dominant strategy* will maximize an individual's utility regardless of the strategy played by other players. Perfect competition will not hold in the risk market due to collusion, espionage, agents *gnimag* the market, static risks, and a reliance on large numbers. Luckily, [16] further proves that there is an equivalence between the classes of economies that exhibit the Pareto-optimal, individually rational, and dominant strategy properties and those that are perfectly competitive. Using techniques attributed to Groves [11] we shall provide employees with the appropriate incentives to make optimal choices, regardless of the temptations to cheat the system.

4.2 Assumptions

There are several assumptions that we make when considering the risk allocation problem.

1. Demand for expending risk far exceeds risk tolerance.
2. Employees are rational and behave selfishly.
3. Employees may collude to increase their utility.
4. There are externalities such as espionage.
5. There are fixed risks regardless of the risk budget.
6. There is no correlation between the benefit an employee will obtain from one resource with the benefit another employee will obtain from accessing another (possibly identical) resource.
7. Risk tokens may be used immediately after creation and there are no risk token storage risks.
8. Benefits, risks, and damages can be quantifiably measured or estimated within a reasonable bound.
9. Risk tokens can only be used to access resources once.
10. New risk tokens are periodically released.

It should be noted that most of these assumptions are designed to more closely model the real world, and do not simplify the risk distribution problem. For example, assumptions 1, 3, 4, and 5 either make efficient allocations more difficult, or provide attackers with additional abilities. Assumptions 6 and 7 are either subsumed by the solutions to fixed costs and externalities or have known solutions in the economics literature [18] and are removed to shorten and focus this work. Assumption 9 is a security formality which prevents double spending risk and is easily covered by the cryptographic cash literature [4]. Assumption 10 is merely required for actual deployment and is presented as a technicality since risk tolerance is a time dependent variable.

We consider Assumption 2 to be reasonable at this stage of the research. Future work may consider other models such as prospect theory [15].

4.3 Risk Market

The risk market is based heavily on the supply chain internal markets of McAdams [19] and the incentives of Groves [11] to obtain a dominant strategy that maximizes the organization's profits and is different from a standard commodities market in a few ways. From the organization's point of view, risk tokens represent expected harm, and the marginal cost of production we shall term the *marginal damage*, $MD(\cdot)$. Production is managed by the security officers of the organization, and a maximum risk tolerance level R^* determines the organization's risk aversion and total risk token production. A regular employee estimates the benefit they will receive from trading risk tokens for access to resources. We term this the *marginal benefit*, $MB(\cdot)$. The marginal damage function defines the organization's risk supply, while the marginal benefit is the demand for risk. To minimize outside influences and facilitate auditing, a strictly internal currency—corporate dollars—are used in the market. If the employee's benefit function is correctly aligned with the organization's, then $\mathsf{AREA}(abc)$ in Figure 3 is directly proportional to the organization's net gains.

By the first fundamental theorem of welfare economics, a perfectly competitive market will result in an efficient allocation. Our work is to compensate for violations of perfect competition and attacks against the system such that we still obtain an efficient allocation.

For this work, we chose to use a double auction which is similar to a stock or commodities market, however there are many alternatives (see Appendix B). One such alternative is a Vickrey-Clarke-Groves (VCG) auction [26, 7, 11] which we will make use of in our simulations.

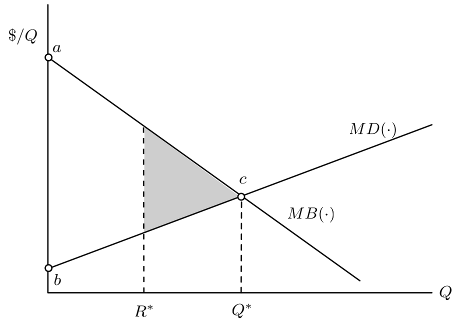

Figure 3: Risk Cost/Benefit. An Organization's Risk Tolerance R^* May Yield Suboptimal Gains.

4.3.1 Marginal Damage

The marginal damage is the expected amount of harm to result from releasing an additional risk token. This requires us to determine how much damage a given token may cause. The Jason Report pegs the value of risk as

> *1 token = risk associated with one-day, soft-copy-only access to one document by the average Secret-cleared individual* [14],

and calculates the cost of access to a resource by

$$\mathsf{Cost} = \mathsf{Risk} \times \mathsf{Damage}. \tag{2}$$

From this equation, we see that each risk token represents a constant amount of harm c. In the strictest model, this yields the marginal damage function

$$MD(Q) = c \tag{3}$$

for constant harm c. Organizations may wish to be more risk adverse and conservative due to uncertainty in the risk calculations, the aggregation problem, mistrust in the market, fear of agent collusion or espionage, or to model cognitive processes, such as weighting a loss more highly than an equal gain [15]. To allow such flexibility we only require that

$$\int_0^Q MD(x)dx \geq Qc. \tag{4}$$

4.3.2 Risk Tolerance

An organization's risk aversion may make them tolerant to a certain aggregate amount of expected damage. This could be due to a formal risk analysis, insurance, or other parameters. Howard [13] has developed guidelines for determining an organization's risk tolerance in terms of total sales, net income, and equity. While the monopoly optimal quantity of risk tokens is Q^* (see Figure 3), the organization's risk tolerance R^* may be different. The organization can appreciate the maximum benefits of information sharing when $R^* \geq Q^*$ and all beneficial accesses are appreciated.

Alternatively, when $Q^* > R^*$, there are lost opportunities and the organization looses the profit $\int_{R^*}^{Q^*} MB(x) - MD(x)dx$ represented by that shaded region in Figure 3. In these instances, it may be beneficial for the organization to consider risk transference through insurance to cover the expected damage gap $Q^* - R^*$ [3].

4.3.3 Fixed Risks

Fixed risks are threats to resources independent of any transaction that do not depend on the risk expenditure. These include miss-configured services such as databases, file servers or web servers, bugs in services or the operating system, weaknesses in encryption keys and algorithms, hardware failures, physical penetration, social engineering, and others. Fixed risks may potentially threaten all assets accessible by the system. The *fixed harm* (FH) is the expected damage to all resources due to such risks.

Fixed risks are typically analyzed at the risk management and not the access control level. By merging these two levels of security, we can adapt the access control system to directly handle these risks. This complicates risk-based access control systems since any surplus benefit for some level of risk production Q must compensate for the fixed harm. If this harm is large, the rational action may be to prevent all access.

McAdams [19] uses Grove's Mechanisms [11] to determine the optimal level of production Q'. We adopt this solution when $Q' \leq R^*$, but must take appropriate actions otherwise. An organization's profit function

$$\mathsf{Profit}(Q) = \int_0^Q MB(x)dx - \left(\mathsf{FH} + \int_0^Q MD(x)dx\right) \tag{5}$$

can be used to determine the quantity Q' that will maximize profits. In the case when $Q' \geq R^*$, the extra risk $Q' - R^*$ may be transfered, otherwise the profits $\mathsf{Profit}(Q') - \mathsf{Profit}(R^*)$ are lost, noting that $\mathsf{Profit}(R^*)$ may be negative, and in some instances, we may minimize the organization's losses.

4.4 Formal Model

4.4.1 Employee and Adversarial Model

An organization is comprised of *employees* who perform work. An employee is considered a *subject* when we need to make an access control decision, and employees are modeled as *agents* when performing simulations. We may use these terms interchangeably. An employee may be an adversary.

Denote the set of possible allocations of risk tokens as Δ such that for $\delta \in \Delta$, employee i receives δ_i risk tokens. Each employee i has a type τ_i and a utility function $u_i(\tau_i, \cdot)$ that determines the desirability of the current state. We will obtain an influence over u_i through the use of a wage W_i given to employee i based on the performance of the market. We assume there is a trusted, honest, rational employee, namely the security officer, in charge of risk assessments and risk token production. This is standard in the access control literature [12].

Regular employees and security officers buy and sell risk tokens using an internal currency and employees are allowed to carry a negative balance, but are subject to periodic audits and security reviews. For each time period T and employee i, we can define the difference in capital, λ_i, and the market profit $\pi_{i,T}$, though without loss of generality we will simply refer to the market profit as π_i. The difference in capital λ_i is defined as the amount of currency the employee received from selling risk tokens minus the currency spent purchasing risk tokens, $\lambda_i = \lambda_i^{in} - \lambda_i^{out}$.

The benefit a regular employee receives from expending their net purchased risk, δ_i, is $B(\delta_i, \tau_i)$, and the cost of producing δ' risk tokens for security officers is $C(\delta')$. We contrast B with u_i as defining only the benefit the employee and the organization will receive. We now define the internal market profit for regular employees as $\pi_i = \lambda_i + B(\delta_i, \tau_i)$ and for security officers as $\pi_i = \lambda_i - C(\delta')$. We define the vector or employee profits as $\overline{\pi}$, and the vector of employee profits excluding employee i as $\overline{\pi}_{\setminus i}$.

4.4.2 Incentives

In a standard free market, an individual's utility is based solely on their market profit, π_i, a cause of much of the instability. In perfectly competitive markets this is sufficient, however when the conditions for perfect competition fail, incentives may promote the convergence to Pareto-optimal solutions, maximizing the organization's profits. For the risk market, we shall use an incentive structure similar to [19, 11]. Incentives are some utility that we are able to bestow on employees based on their behavior. Internal markets are unique in the amount of control over an individual's utility function via incentives an organization has when compared to traditional markets, such as a stock exchange. These incentives can be realized monetarily though a bonus since we have direct control of an employee's salary, and can prevent many of the problems such as over or under speculation that may skew market prices that are present in traditional commodities markets. An employee's wage W_i can be modeled abstractly as a function of her market profit and the market profit of all the employees, plus some base salary Y_i.

$$W_i = Y_i + \mathsf{X}_i(\overline{\pi}_{\setminus i}) + \alpha\pi_i \tag{6}$$

Malicious agents may behave as rational agents with an external influence, ϵ_i, such as money from nefarious activities

$$\hat{W}_i = W_i + \epsilon_i(\delta). \tag{7}$$

The goal of mechanism design is to choose a function X_i and α such that it is in the employee's best interest to make optimal decisions for the organization. We further must compensate for the unknown payoff $\epsilon_i(\delta)$ that an employee receives from their malicious activities. [19] and [11] define several incentive structures:

1. *Fixed Wage* $W_i = Y_i$. Employees have no incentives.

2. *Market Wage* $W_i = Y_i + \alpha\pi_i$. Employees have a direct incentive to maximize market profits.

3. *Cooperative Wage* $W_i = Y_i + \beta(\pi_i + \sum_{j\neq i}\pi_j)$. Employees have a direct incentive to maximize the organization's profits.

4. *Peer Group Relative Wage* $W_t = Y + \gamma\left(\pi_t - \frac{\sum_{t'\neq t}\pi_{t'}}{|\{j|\tau_j=t\}|-1}\right)$. Employees of type t have an incentive to increase his own profit and decrease the profit of their peer group.

5. *Accurate Prediction Wage* $W_i = Y_i + \alpha\pi_i - \kappa(\pi_i - \tilde{\pi}_i)^2$. Employees have an incentive to bid accurately. See Section 5.3.

5 Attacks and Defenses

We now consider the various goals an attacker may have, and the mechanisms in place in the market that protect the organization from these attacks.

5.1 Decreasing the Utility of Other Agents

Malicious agents may wish to decrease the utility of honest agents. Unless the malicious agents yield exceptional market power, it is unlikely they will be able to adversely affect the organization's profit. This will require the collusion of a large number of employees; See Section 5.3. Employees with a cooperative wage will be most adversely affected by reducing their market profit. This can be done in one of two ways: increasing the sale price forcing honest employees to pay their maximum value or by a denial of service (DOS) attack (see Section 5.2).

Malicious employees attempting to force honest traders to pay their maximum price for risk tokens will need to deny profitable trades (otherwise the market is behaving correctly and there is no attack). A cooperative wage provides the incentives against such an attack. By decreasing the market profit for honest employees, an attacker decreases the organization's profit, and consequently his own wage. A similar argument may be made with DOS attacks. It is in the best interests of the malicious employees to make the profitable trades with a cooperative wage.

In double auctions, bid and ask prices are openly posted, providing attackers with additional information which can potentially make some of the above attacks easier. Thrope and Parkes [25] describe a scheme that allows a double auction to function with unconditionally hidden bid and ask commitments. While it is still possible for attackers to infer the private information of others, it increases the cost of the attacks and their visibility as actively attacking the system.

5.2 Denial of Service (DOS)

An attacker may wish to prevent other agents from accessing resources. When an organization determines its risk allocation quantity solely on the marginal damage and marginal benefit (supply and demand) of the employees, the organization will always produce the quantity Q^*. A malicious agent purchasing additional risk in an attempt to block other agents will simply increase Q^*. In this setting, risk is a limited resource in only the most cursory sense. The DOS will be unsuccessful, and we can use the techniques in Section 5.3 to identify the malicious agent.

When the organization produces $\min(Q^*, R^*)$ risk tokens, then a malicious agent can perform a DOS attack by purchasing a large quantity of tokens and refusing to sell to other agents. If they purchase an additional Q' risk tokens, this will cause the agent to lose $\alpha\int_{Q^*-Q'}^{Q^*} MB_i(x) - MD(x)dx$ in market profit. By definition of the attack, this will decrease the profit for the organization. If we define the vector of market wages where the agent behaves maliciously as $\overline{\pi}'$, then the agent loses the incentive $\beta(\sum_{j\neq i}\pi_j - \sum_{j\neq i}\pi'_j)$. When the rogue incentive ϵ_i is small, a rational agent can maximize their utility by behaving honestly and maximizing the organization's benefit. When ϵ_i is large enough to make this is the dominant strategy, we resort to techniques in Section 5.3 to identify the malicious agent.

5.3 Escalation of Privileges

One of the main criticisms with the market mechanisms is that they potentially allow malicious agents to escalate their privileges. First, we note that the aggregate amount of harm that may be caused can be restricted by the organization to R^*. The organization has the additional control of the risk assessment on all access requests, which could take the requester's trustworthiness, need, and access history into account. Intrusion detection techniques may be employed to detect employees who are behaving suspiciously and increase their risk rating, and consequently the cost of accessing damaging resources. One such solution is to define an employee's risk to be a monotonically increasing function with regards to the number of risk tokens they have retained (purchased or carried over) within a given time interval. Employees attempting to "stockpile" or horde risk tokens to make a damaging purchase will find the cost to access the resource increases with the progress of the attack.

Beyond these measures whose defenses reside in the fundamental design of the risk-based access control systems, we can use the market incentives and mechanisms to detect potentially malicious employees and punish them. In Groves' and McAdams solutions to fixed costs, a conditional internal market is used to determine if production will occur. Before production begins, each agent states an amount they are willing to pay, $\tilde{\pi_i}$, for there to be production. These values are used to compensate agents who may be adversely affected by the decision of the group. We refer the reader to [11, 19] for details on the conditional market.

McAdams [19] uses the prediction value, $\tilde{\pi_i}$, to create a strictly dominant strategy where each employee states their actual expected market profits by punishing inaccuracy

$$W_i = \alpha(\pi_i - p_i) - \kappa(\pi_i - \tilde{\pi}_i)^2. \tag{8}$$

In this case κ is a positive constant and p_i is a charge or a compensation value based on the market profit predictions of all other employees and π_i is the ex post observed marketprofit [11]. The intent of p_i is to make the allocation of Q^* risk tokens that maximizes the organization's utility to be beneficial for all employees.

Malicious employees that will have low observed benefits compared to their expended capital purchasing risk tokens, will incur a high penalty and identify themselves as either malicious or incompetent. In either case, increased audits, risk, or other mitigating measures can be applied to the employee. This is an advanced warning system that adequately identifies and handles suspicious behaviors and suppress the damage caused by the unbounded accesses.

Depending on the malicious incentive ϵ_i, this penalty may again align the organizations benefit and the employees', creating an incentive to behave honestly, or leak only the resources they had a valid need to access. In more extreme cases of espionage this may not create strong enough incentives, but the employees market profit deviation is enough to easily identify their malfeasance and even estimate their malicious incentive ϵ_i. In situations where employees do not directly state their expected market profit in advance, we can estimate this value by observing their trading patterns and their net loss in revenue. The market has unexpectedly becomes an integral part of an anomaly based IDS.

5.4 Collusion

Collusion among employees may increase the market profit within a group. By fixing prices, over/under–estimating costs, etc., a group of colluding employees can manipulate market prices. [19] notes that peer group relative wages will remove the advantages of colluding.

5.5 Greedy Employees and Market Prowess

One extremely attractive property of these incentives is that the more greedy, selfish, and rational an employee is, the better the organization performs. Thus an employee attempting to maximize their own utility will do so by making trades that maximize the organization's profit, leading to efficient allocation of risk and decreased losses. Furthermore, interaction with the market may be automated due to dominant strategies and experimental economics work indicates that even random yet rational (constrained zero-intelligence (ZI-C)) bids will converge to optimal distributions [10]. See Appendix B.1 for more details.

6 Simulation

To test our hypothesis that token-mechanism risk-based access control systems perform well, we simulated an organization's access to resources and market interactions. We compared our results to an MLS system such as Bell-LaPadula without compartments. The simulation is implemented in C++ on Linux using a machine with two 2.0 GHz, dual-core X86 processors and 4G bytes of RAM. It uses Monte Carlo simulation and allows specification of parameters such as number of agents, risk tolerance, the range of agent trustworthiness, resource damage and benefits, hard boundaries, marginal damage constant c, and the average slope of the marginal benefit function. For each set of parameters, the results were averaged over several runs of the simulation.

6.1 Simulation Design

We model a quantified risk-based access control system similar to the one described in [14], which assumes all risks are quantified and access control is governed by Equation 2. Since our results relate to the distribution of risks, it is applicable to both the region between the soft and hard boundaries in FuzzyMLS [6] and the redistribution capital in [28].

We simulate the access control patterns of an organization, which is a collection of employees that we term agents. The agents access resources, such as files and databases, that have associated benefits and damages. An agent obtains the complete benefit when they access a resource, and the complete damage only if the resource is harmed. We determine a priori which resources are harmed based on the risk value of the transaction (actual risk) and calculate the cost of the transaction based on estimated risk, which is bounded to be within one standard deviation of the actual risk. We assume all fixed costs are zero to prevent the need to additionally model Grove's mechanisms. Each agent is assigned a competence value which affects their ability to estimate benefits and their ability to bid on the market. Finally, each agent is charged the amount in Equation 2 and hard upper bounds such as those discussed in [6] were used to deny transactions.

To maintain consistency between simulations and distribution methods, the integrity of the risk and damage calculations, and simplify the simulation, we constrain agents to access resources in a predetermined order, thus forming a queue. The benefit values within an agent's queue trend down, but are not monotonically decreasing.

6.1.1 Preallocation

We consider three possible preallocation methods:

1. *Risk Level* - We use the inverse of an agent's risk value to obtain a trust value. A clearance level (similar to BLP) is $\lfloor \log \mathsf{Trust} \rfloor$. All agents at the same clearance level are given the same number of risk tokens, and agents at higher levels are given more tokens than agents at lower levels.

2. *Risk Proportional* - Risk tokens are distributed similar to the risk level distribution method in a continuous manner. i.e. we do not take the floor.

3. *Constant* - Each agent is given an equal share the the risk budget.

We do not consider the request or hierarchical distribution methods since they are difficult and impractical to simulate and are known not to provide efficient allocations [16]. In real instantiations, risks may be traded by management and deposited into employee accounts for simplicity.

6.1.2 Market Strategy

The incentives from Section 4.4.2 encourage rational individuals to behave honestly when interacting with the market. For our simulation, agents place bids and asks as ZI-C traders similar to [10]. To constrain the agents, we need to determine their marginal benefit function, and we investigate two different functions. First, the expected benefit for a single resource is dependent on the agent's competence value, and is taken from a uniform random distribution for the region $\mathsf{Benefit} * \mathsf{Competence}, \mathsf{Benefit}/\mathsf{Competence}$ where $\mathsf{Competence} \in [0, 1]$. A higher competence value will provide a tighter bound on the actual benefit for the resource and will bias as agent to bid higher rather than lower, a phenomenon often observed in experimental economics [5]. An agent's marginal benefit function is defined one of two ways

1. *Iterative* - An agent individually considers each resource independently from their queue. Resource $m + 1$ is only considered after the first m.

2. *Foresight* - An agent considers accessing the next m resource concurrently such that the amortized return on investment (ROI) is maximized.

Note the foresight method creates a monotonically decreasing marginal benefit while the iterative method may not.

6.1.3 Vickrey-Clarke-Groves Optimal

For comparative purposes only, we assume perfectly competent and rational agents playing the dominant strategy and use Vickrey-Clarke-Groves (VCG) mechanisms to determine three optimal distributions:

1. *Maximize Benefit* - This will maximize the benefit the organization obtains.

2. *Maximize Net Gain* - Using the ex post leaked documents, we determine the maximum profit the organization can obtain.

3. *Maximize Damage* - Using the ex post leaked documents, we determine the maximum damage the organization can obtain.

6.1.4 Simulation Parameters

Due to the lack of availability of real world data regarding information loss and breaches, we varied the distributions of all parameters and performed Monte Carlo simulations to obtain coverage of the search space. Log normal distributions were used for damage, risk, and benefit values while a normal distribution was used for competence. We varied the number of agents, the maximum number of resources in each agent's queue, the risk budget, the marginal damage (Equation 3), and the average rate of drop in the benefit.

6.2 Monte Carlo Simulation Results

6.2.1 Comparison with Bell-LaPadula (BLP)

As one of our alternative access control models, we simulated BLP. To accomplish this, we first assume that an agent's trust level and a document's damage are divided into clearance and classification levels on a log scale, using a base of two. To convert a risk level into a trust level, we simply assume a maximum trust level, and subtract an agent's risk. To allow a transaction, we require that an agent's trust level dominates the resource's classification level. We don't simulate compartments.

Since BLP does not manage risk we, do not impede access, but we do tally the expended risk. We use identical agents and resources to compare the market mechanisms and BLP and allow the market to regulate its own risk budget. We note several observations. The results are shown in Figure 4.

1. BLP yields greater benefits, but significantly greater harm than the risk-based systems.

2. BLP is more erratic, and more likely to yield net losses.

3. BLP has a lower ROI than the market.

4. For BLP, harm depends on the damage of the resources accessed, while with the risk-based system harm only depends on the amount of risk expended.

5. All information sharing stops when resources become too damaging in BLP at the point when most resources dominate the subjects accessing them, while access only slows for risk-based systems. This is discussed more in the context of marginal damage below.

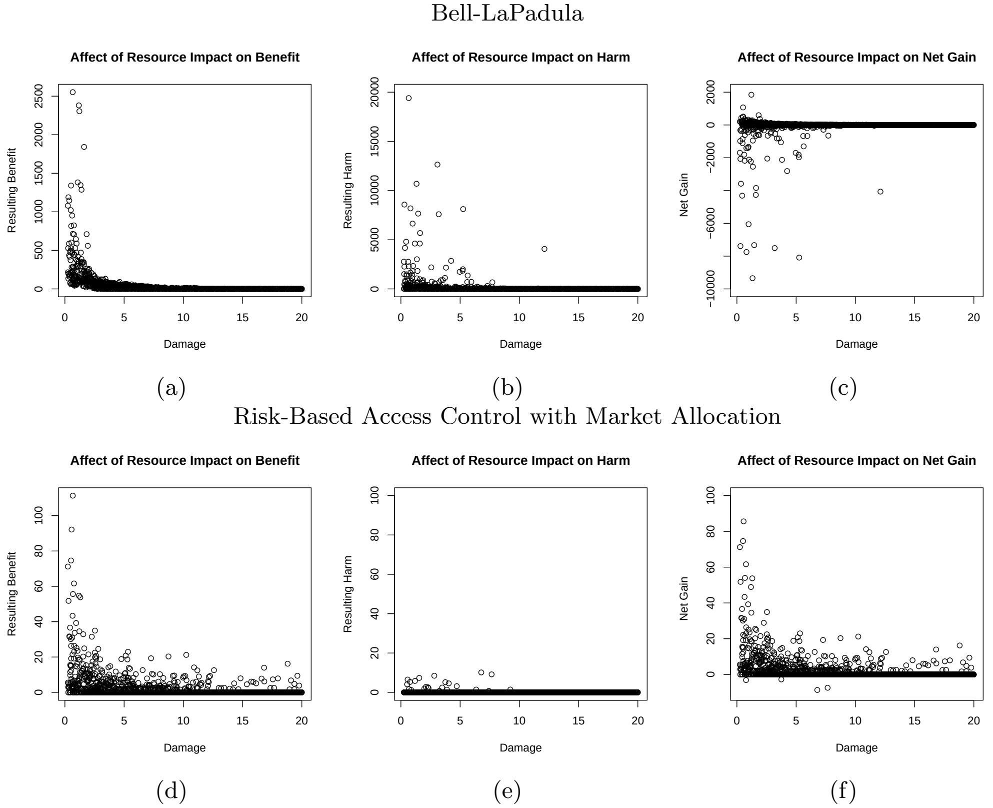

(a)

(b)

(c)

(d)

(e)

(f)

Figure 4: Bell-LaPadula (a-c) Versus a Risk-Based Access Control System (c-e) using a Risk Market

6.2.2 Effect of Risk Allocation on Expected Benefits

Next we compare the various preallocation schemes with the market based mechanisms and the VCG mechanisms to determine efficient allocations and risk utilization. Recall from Section 6.1.2 that an agent's competence bounds their ability to estimate the benefit they expect to obtain from purchasing risk, and efficiency is the ratio of extracted utility to maximum utility. Since the market mechanisms, even given the zero-intelligence traders described in Section B are known to lead to efficient allocations, the new results from our risk market shows how the agent's ability to estimate benefits affects the efficiency. In Figure 5, we compare an agent's competence to the efficiency of the market. Possible extracted utility is calculated from the VCG mechanisms. Our simulations showed the following:

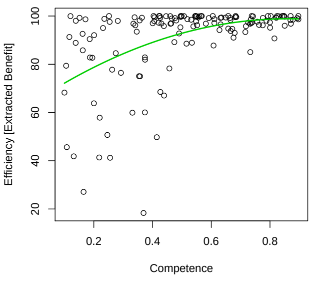

Figure 5: The Affect of Competence on Market Performance. A Linear Regression is Shown.

1. VCG was able to allocate around 98-99% of R^*.

2. The double auction market was able to use around 92-95% of the same allocation.

3. All preallocation methods were extremely inefficient, and only used around 15% of R^*.

4. The risk market was extremely efficient, in many cases exceeding 90%.

5. Agent's require a minimum average competency of around 42% to reach 90% efficiency.

The low utilization for preallocation is due largely to agents being incapable of purchasing access to resources due to insufficient allocations. When agents are given more tokens than they are capable of using, the remaining risk surplus is wasted and cannot be reallocated to other agents to increase the efficiency.

6.2.3 Determining Organization Risk Tolerance

Any organization should engage in a formal risk assessment before a risk-based access control system could be adequately implemented. During such an analysis, a risk tolerance R^* may be determined. Regardless of this quantity, our experiments suggest that, barring fixed costs discussed in Section 4.3.3 which we did not simulate, allocating no more than Q^* risk tokens is an optimal strategy. The organization has a level of control of the quantity Q^* by manipulating the marginal damage function. Figure 6(a) illustrates how an adequate amortized marginal damage constant c, can ensure the expected benefit remains positive. See Section 4.3.1 for details on the marginal damage. Using a lower marginal damage can result in greater net gains (due to increased information sharing and resource access), but constitutes a greater risk of a net loss. By increasing the marginal damage too high, employees are starved and are prevented from accessing valuable resources, and information sharing drops to zero. In Figure 6(b), we illustrate the ROI ratio of net gains per aggregate risk expenditure. Note that the change in plot density above $MD = 2$ is due to complete starvation where $Q^* = 0$, and the ROI is undefined.

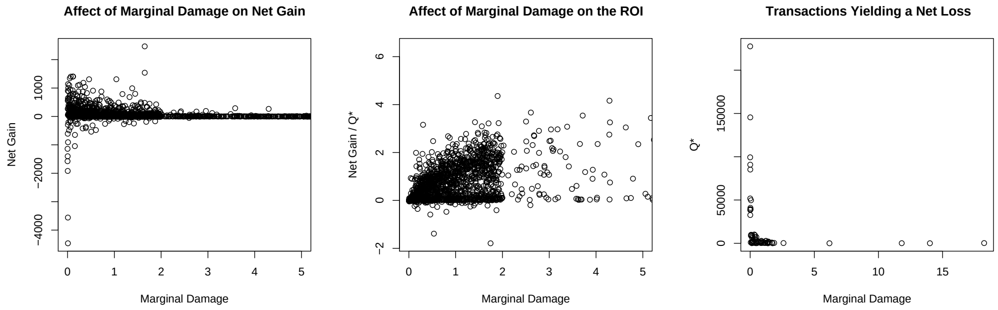

(a) (b) (c)

Figure 6: The Optimal Risk Allocation Q^* Can Be Determined by the Market by Modifying the Marginal Damage Function. Note That While a Lower Marginal Damage Increases the Net Gain, It Also Increased the Risk Budget and Variability, Increasing the Risk of Incurring a Net Loss.

6.2.4 Effect of Hard and Soft Boundaries

In our experiments, we did not model a soft boundary. When strictly using the market forces and always distributing Q^* risk tokens, a soft lower bound (below which all transactions incur a cost of zero) has no impact on the performance of the market; risk production shifts from Q^* to Q', where $Q' - Q^*$ represents the risk associated with transactions below the lower bound. We modeled an upper bound and classified transactions into two categories: net loss and net gain. For these experiments, all R^* risk tokens were released, corresponding to $MD = 0$. Figure 7(c) illustrates the class of transactions that produced a net gain, while Figure 7(d) are the transactions that resulted in a net loss. Observe that, compared to Figure 6(c), the distribution of net losses is uniform across risk budgets and hard boundaries. While hard boundaries are insufficient for ensuring net gains, they are beneficial for restricting the long negative tail that is present

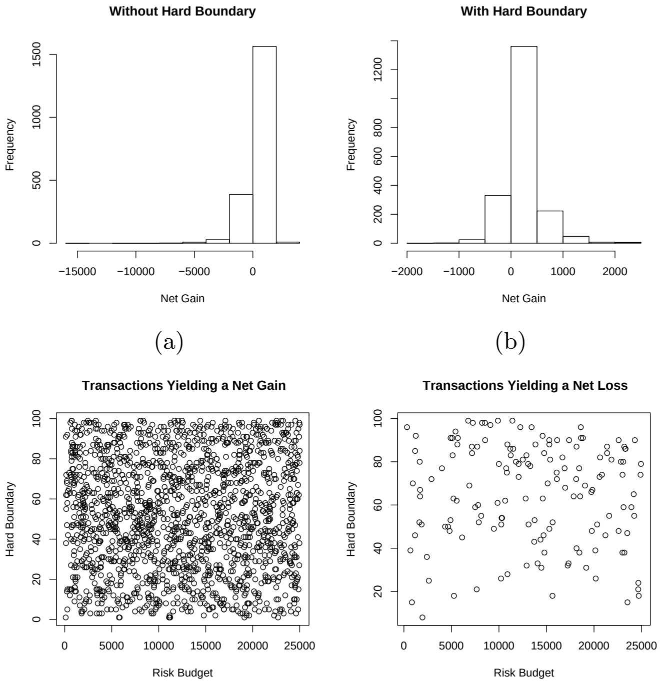

(a) (b)

(c) (d)

Figure 7: Hard Boundaries Limit the Long Negative Tail of the Net Gain Distribution, but Are Insufficient at Ensuring Positive Net Gains

when hard boundaries are not used. Figure 7(a) and 7(b) show the distribution of net gains without and with hard boundaries respectively.

7 Conclusion and Future Work

We have shown that markets, such as a double auction, are extremely effective at determining efficient risk allocation within an organization. When provided with the correct incentives, employees will directly benefit by making optimal choices and greedy behavior ensures a convergence towards the optimal distribution. We show that markets are also effective at allowing an organization to dynamically determine their appropriate risk allocation quantity that will maximize their returns for the given time period. Further, we illustrate how the market may be used as an effective IDS, where rogue employees are identified. This allows for highly dynamic and pertinent risk mitigation measures to be taken to limit the organization's expected harm.

Some issues in [14] need to be addressed. The risk market could be extended to allow multiple departments within an organization to independently make risk and damage assessments and allow efficient sharing of resources among them, such that no department is capable of increasing their profits by falsely inflating the impact (and thus the cost) of their resources. The risk market could also allow information sharing among entities such as governments. Each entity will produce its own risk tokens and internal currency. This introduces additional complications that need to be addressed such as inflation, liquid resale markets, and currency exchange markets.

References

[1] FCC: Wireless Telecommunications Bureau (WTB), October 2007. `http://wireless.fcc.gov`.

[2] J. Bingaman, A. Specter, T. Harkin, T. Stevens, L. Murkowski, and D. Akaka. *Low Carbon Economy Act of 2007*, 2007. 110^{th} United States Congress.

[3] R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. *WEIS 2006*, June 2006.

[4] S. Brands. Untraceable off-line cash in wallets with observers. *CRYPTO*, 1993.

[5] D. Brenner and J. Morgan. The vickrey-clarke-groves versus the simultaneous ascending auction: An experimental approach. *A1.133 WP 188*, 1997.

[6] P.-C. Cheng, P. Rohatgi, C. Keser, P. A. Karger, G. M. Wagner, and A. S. Reninger. Fuzzy MLS: An experiment on quantified risk-adaptive access control. *IEEE Symposium on Security and Privacy 2007*, 2007.

[7] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1), September 1971.

[8] G. Cybenko. Why johnny can't evaluate security risk. In *IEEE Security & Privacy Magazine*, volume 4, pages 5–5, Jan-Feb 2006.

[9] P. Erwin and J. Hardy. Draft climate change bill. Technical report, Department for Environment, Food and Rural Affairs, March 2007.

[10] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Markets as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1), 1993.

[11] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[12] M. A. Harrison, W. L. Ruzzo, and J. D. Ullman. Protection in operating systems. *Communications of the ACM*, 19(8):461–471, 1976.

[13] R. A. Howard. Decision analysis: practice and promise. *Manage. Sci.*, 34(6):679–695, 1988.

[14] JASON Program Office. Horizontal integration: Broader access models for realizing information dominance. Technical Report JSR-04-132, MITRE Corporation, 2004.

[15] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, March 1979.

[16] L. Makowski and J. M. Ostroy. Vickrey-clarke-groves mechanisms and perfect competition. UCLA Economics Working Papers 333, UCLA Department of Economics, July 1984.

[17] T. W. Malone. Bringing the market inside. *Harvard Business Review*, pages 106–114, April 2004.

[18] D. McAdams. Storage in internal markets. `http://www.mit.edu/~mcadams/papers/im/storage.pdf`, 2005.

[19] D. McAdams and T. W. Malone. Internal markets for supply chain capacity allocation. Technical Report 4546-05, MIT Sloan School of Management, 2005.

[20] R. P. McAfee. *Introduction to Economic Analysis.* 2006.

[21] B. C. on Banking Supervision. International convergence of capital measurement and capital standards. Technical report, Bank for Internaional Settlements, June 2006. Basel II.

[22] M. H. Rothkopf. Thirteen reasons why the vickrey-clarke-groves process is not practical. *Operations Research*, 55(2):191–197, 2007.

[23] A. Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations.* 1776.

[24] S. Sunder. *Experimental Asset Markets: A Survey*, chapter 6. Princeton University Press, 1995.

[25] C. Thorpe and D. C. Parkes. Cryptographic securities exchanges. *Financial Cryptography and Data Security*, February 2007.

[26] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, March 1961.

[27] G. Wearden. The biggest rogue traders in history. January 24 2008. http://www.guardian.co.uk/business/2008/jan/24/europeanbanks.banking.

[28] L. Zhang, A. Brodsky, and S. Jajodia. Toward Information Sharing: Benefit And Risk Access Control (BARAC). *Policy 2006*, pages 45–53.

A Markets in Other Settings

While the proposal to use market economies of risk and damage seems outlandish, the concept is not foreign in other areas. Since 1993 Congress has allowed the Federal Communication Commission (FCC) to use auctions to resolve license application conflicts for the electromagnetic spectrum. The FCC uses simultaneously ascending auctions to efficiently allocate the limited resource, and since 1997 the use of auctions has been required for such applications. Similar systems have been adopted by other countries and for other applications [1].

Several countries have proposed legislation to reduce the amount of greenhouse gases produced such as the United States Low Carbon Economy Act of 2007 [2] and the United Kingdom Climate Change Bill [9]. Businesses would be required to purchase allotments from the government, allowing them to produce a given amount of greenhouse gases, such as CO_2. Companies caught producing more greenhouse gases then they are allowed are fined—hopefully more than the competitive equilibrium—producing an incentive to comply.

BP [17] has experimented with market economies internally to reduce the amount of CO_2 and other greenhouse gases produced. By performing a lump-sum distribution and allowing entities within the organization to trade CO_2 allotments, they can determine the most economically efficient method to decrease their carbon footprint.

Intel, in conjunction with MIT, performed similar experiments where plant managers and sales representatives attempted to make efficient use of chip production capacity. Their experiments are similar to Sunder's [24]—which will be discussed in Section B—and provided each player with different information regarding the marginal cost of producing chips, supply and demand, and sales forecasts. Their experiments were extremely successful, and allocation efficiency rose from 86.6% to 99% by the third round [17, 19]. HP has performed similar experiments to forecast sales figures for their printer devision [17] and as efficient scheduling algorithms for utility data centers[1].

Markets have been proposed infamously for other security applications. The Defense Advanced Research Projects Agency (DARPA) began work on a project called FutureMAP Policy Analysis Market that was intended to be used as a futures market for potential terrorist attacks in the Middle East. After much furor, however, the project was abandoned[2].

B Market Mechanisms and Auction Theory

Much of the literature on dominant strategy mechanisms are on based on the works of Vickrey [26], Clarke [7], and Groves [11] and are termed Vickrey-Clarke-Groves (VCG) mechanisms. A Vickrey Auction [26] is a sealed-bid second price auction where bidders submit their bids without knowing the bids of others, and pay the amount of the best loosing bid. Under these conditions, the dominant strategy is to bid the average value when considering auctions for multiple goods. When considering multiple round auctions, such as quarterly distributions are fresh risk tokens, this strategy is no longer a weak equilibria, making it unattractive for our purposes. Clarke [7] later extended Vickrey's work to multiple item auctions where truth-telling is the dominant strategy. This is accomplished with a variable charge based on the difference in an individuals assigned output and actual output. Groves [11] considers the problem of determining compensation so that truth-telling is the dominant strategy and individuals behave optimally.

Combined, these are known as VCG mechanisms and work as follows. An individual submits bids for all combinations of goods that are being auctioned. A central authority determines the optimal distribution based on the bids, and each agent pays the highest amount that would have been bid for the objects had they not been present. VCG auctions are advantageous in that truth-telling is the dominant strategy, and they can determine not only the optimal distribution, but also the optimal number of resources to be distributed

[1] A. Byde, M. Sallé, and C. Bartolini. Market-based resource allocation for utility data centers. Technical report, Hewlett-Packard, 2003

[2] T. Daschle. Trading in death. Congressional Record, July 29 2003

and an optimal pricing policy. While VCG mechanisms are extremely attractive in theory, they do not work well in practice. Rothkopf [22] comments on thirteen problems with the VCG process that make them theoretically attractive yet impractical. These problems range from being NP-complete, the disclosure of valuable confidential information, possible collusion among bidders, to issues related to the dominant strategy being only a weak equilibrium, and may not be an equilibrium in multiple run auctions, such as a quarterly distribution of new risk tokens. Due to these problems, we must look into more practical alternatives to the VCG process. We do comment on the usage of VCG in our simulations, and use them to calculate the optimal distributions which we then compare alternative markets to.

B.1 Double Auction

A standard free market, such as a stock or commodities market, are known as double auctions. Sunder [24] provides a survey of double auction markets and their ability to disseminate information among players. In their experiments players had private information regarding possible states and values of assets. While the simplest markets converged to the competitive equilibrium rapidly, the parameters and circumstances such as the number of states, rate of information dissemination, ability to purchase information, futures markets, blinding, and others, affect the ability to converge and the rate of convergence. In general, asset markets were effective at providing efficient distributions of assets.

In some configurations when the number of possible states are large, or the amount of private information in the system is too low, the market may converge to a false equilibrium. This is often the result when a large enough number of the traders misinterpret the market and assume an incorrect state. Their trading behavior influences the beliefs of other traders, resulting in the false equilibrium. It is possible that false equilibrium could be eliminated by removing the short sale restriction [24].

In all of the above experiments human traders were used, and the convergence could naturally be attributed to their rationality, memory, motivation, and learning. Gode and Sunder [10] question this hypothesis by employing what they termed "zero-intelligence" traders. Constrained (ZI-C) agents were prevented from trading at a loss (individually rational), while unconstrained (ZI-U) agents were not. Both sets of agents place bids and asks taken from a uniform random distribution. While the ZI-C agents were unable to learn from past trading experiences, within each time interval the allocation efficiency of these markets approached 100 percent, while ZI-U agents did not. While the allocation efficiency of the human and zero-intelligence traders is indistinguishable, human motivation to maximize profits results in a lower price variability and profit dispersion among the agents [10]. These results are encouraging. When combined with the appropriate incentives, such as those discussed in [11, 19], agents need only be rational for the market to perform well.

C Real World Example

While information security does not yet have mature metrics for calculating risks, other industries do have well established risk metrics. For example, the Bank for International Settlements (BIS) has sought to standardize regulations and risk calculations for banks internationally as found in Basel II [21]. The long established use of risk calculating and management make financial institutions attractive for deployment of risk-based access control systems.

Several infamous instances of fraud illustrate where risk-based access control systems could have been advantageous to banks. By 1995, Nick Leeson bankrupt Britain's oldest merchant bank, Barings Bank, with US$1.4B is losses in futures trades. For more than a decade, Yasuo Hamanaka made fraudulent trades in copper, eventually losing Sumitomo Corporation US$2.6B. In 2007 and 2008 Jerome Kerviel, a trader at Societe Generale, France's second-largest bank, lost US$7.1B making fraudulent futures trades [27].

A risk-based access control system could be integrated into a bank that would allow them to practice proper risk management—preventing rogue traders from causing excessive damage. A fixed amount of *risk tokens*, commensurate with the bank's total risk tolerance, are released into the risk market. When an employee wishes to make a transaction, the financial risk is calculated, and the employee is charged the appropriate amount of risk tokens. The observable gains each quarter are easily determined; the employee's bonus can be determined based on the incentives in Section 4.4.2. If the employee does not have enough risk tokens to cover the transaction, he may purchase more on the internal risk-market with an internal currency. Unbounded fraud is not possible due to the limited amount of risk tokens in the market. A per–employee limit on risk–taking, expressed as total amount of risk tokens charged to the employee, may also be enforced. While risk-based access control systems may be beneficial for other industries such as the intelligence community described in [14] or the “gray area” between allow and deny such as Fuzzy MLS [6], risk token mechanism enforced risk-based access control systems are naturally suited to financial institutions initially. A traditional access control system could be used together with a risk–based system to ensure some undesirable actions are never allowed.