Trading in Risk: Using Markets to Improve Access Control*

Ian Molloy
IBM T.J. Watson Research Center
Hawthorne, NY, USA
and Purdue University
West Lafayette, IN, USA
imolloy@cs.purdue.edu

Pau–Chen Cheng and Pankaj Rohatgi
IBM T.J. Watson Research Center
Hawthorne, NY, USA
{pau, rohatgi}@us.ibm.com

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.

Categories and Subject Descriptors

D.4.6 [**Operating Systems**]: Security and Protection—*Access Controls*; K.6.0 [**Management of Computing and Information Systems**]: General—*Economics*

General Terms

Security, Economics, Management

Keywords

risk, market, risk-based access control

*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.

NSPW'08, September 22–25, 2008, Lake Tahoe, California, USA.

1. INTRODUCTION

The majority of information security research has been, at its core, about classifying actions into two classes (e.g., allow and deny, good and bad, or safe and dangerous), and ensuring this division is not violated. This is seen in virus scanners, intrusion detection, firewalls, and spam filtering. One of the primary mechanisms used to ensure a safe division between the two classes is access control. Access control systems make the distinction between allowed and denied at many levels. Lampson proposed an extremely general model [38] that was later refined and formalized by Harrison, Ruzzo, and Ullman [28]. The state of the system is represented by a matrix M where each row is a subject s and each column is an object o. An action a is allowed if the required right $r \in M[s, o]$. Harrison, Ruzzo, and Ullman also proved that answering the safety question of the model, "Can the system enter a dis–allowed state ?" is Turing Undecidable.

This model is too general and not very practical for specifying access control policies, and new models, such as multilevel security, role-based access control, or the Chinese wall, provide abstractions between subjects, objects, and rights that provide the access control system with a set of properties. While these abstractions ease administration, they produce errors in the classification; some actions that were once denied are now allowed, and some that were allowed are now denied. Neither is good in practice [30].

These systems have other inherent problems. Any action, whether allowed or denied, represents a risk to the system and its resources and presents possible benefits or gains. Each benefit is obtained by the user who performed the action, while the risks are shared between all users. Each user will attempt to maximize their gains by performing allowed actions with little consideration for the risks shared among all users. The users and the system as a whole may take too much risk.

It is our view that the field of information security should be viewed as a problem of *risk management*, where risk is roughly defined as the *expected values of damages* and treated as a *countable and finite resource*; the damages are the possible outcomes of security decisions and actions. When

risk is a limited resource, taking too much risk as a whole is an example of the "the tragedy of the commons" where individual actions over–tax a limited common resource and result in a very un–desirable outcome. This concept has been applied to resource management in other fields, including population growth and pollution [27], and in information technology storage systems [56]. If an organization does not treat risk as a limited resource, then it is willing to assume an unbounded amount of damage.

The field of information security does not typically view risk as a finite resource; we tend to bound the amount of risk we are willing to assume for each action, ignoring the aggregate effects that cause the tragedy. Failures to appreciate and understand the tragedy of the commons as it applied to access control has been seen in infamous espionage cases such as Aldrich Ames [61] and Robert Hanssen [59] and in the financial sector by rogue traders Nick Leeson and Jerome Kerviel [60].

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. 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 a properly set–up market tends to allocate resources in an optimal way [55, 25]. We validate this argument by simulating a risk market and other risk allocation methods, such as centralized pre–allocation. We also simulate a multilevel access control system which resembles a Bell–Lapdula model without categories. 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 multilevel system not only by delivering much better risk vs. benefit tradeoff but also by staying within an organization's risk tolerance when the multilevel 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 and access behaviors are logged. The benefit or loss that result from risk–taking will also be logged. Correlating among these logs could provide insight into the reasoning and decision making process of the employees as it identifies a minimum value the employees placed on the resources. In other words, the market provides information that enables a cost-benefit analysis to gain insight into *why* resources are accessed while a traditional access log can only show what resources are accessed.

The remainder of this paper is organized as follows: Section 2 provides an overview of how the risk market will fit into an organization's access control mechanisms. Section 3 discusses related work and the open problems we address. Section 4 highlights skepticism of risk–based access control systems and how we plan to address them. Section 5 provides a background on auction theory and describes the risk market, and Section 6 discusses the attacks on risk market and defenses to these attacks. In Section 7 we describe our simulation and experimental results. We conclude in Section 8.

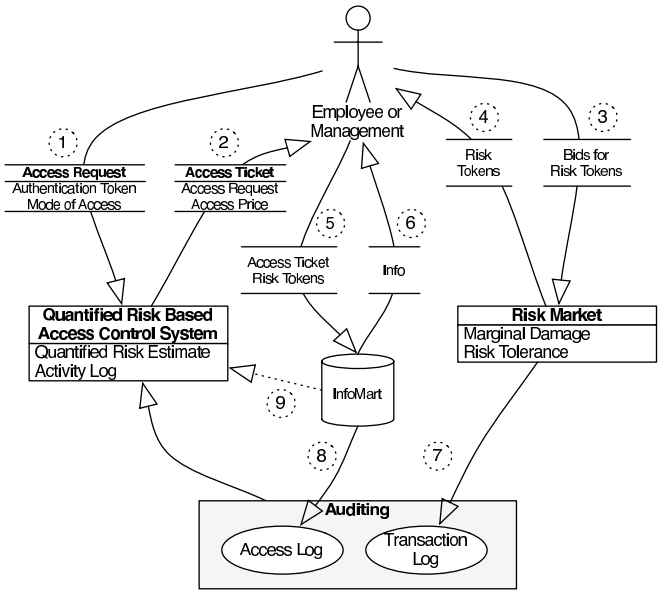

Figure 1: Overview of How the Market Fits into an Organization's Access Control

2. OVERVIEW OF SOLUTION

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.

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 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).

When interacting with the risk market, employees will use a strictly internal currency; *corporate dollars*, similar to the currencies used within virtual worlds. This helps minimize external influence, yet does not eliminate them. An employee may receive periodic sums of internal currency with which to trade on the market, or may be provided a line of credit, such as a credit card. Additional currency may be provided by the organization, and employees may receive additional payments for tasks they complete.

2.1 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 [3]. 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 £827B in losses in futures trades[1]. 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 Société Générale, France's second-largest bank, lost US$7.1B making fraudulent futures trades [60].

Financial risk is often measured as a value at risk (VaR), and managed from the top down. As the VaR increases beyond acceptable limits, risks are transfered and sold off, such as mortgage backed bonds [36]. This manages risk at the highest, aggregate level. Access control is managed by allowing transactions below a risk threshold, or with a desirable cost-benefit ratio, often made desirable using risk mitigation like higher interest rates [37]. While these institutions manage their risk, the management is quite different than what we propose here.

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 (or other time interval) are easily determined; the employee's bonus can be determined based on the incentives in Section 5.4.2. 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 [30] or the "gray area" between allow and deny used in Fuzzy MLS [14], 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.

[1]In 1994 Barings Bank had £5.9B in assets yet was sold to ING in 1995 for only £1. This implies that the damages were greater than the sum of the losses, similar to the inference problem in access control.

3. RELATED WORK

There have been several works proposing risk-based access control systems. The MITRE Jason Report [30] 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 [14], 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 [30].

A third work from Zhang et al. [66] 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 satisfying a weak-optimal (Pareto-optimal[2]) 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.

MarketNet [63, 62] proposes the use of domain specific currencies to control access to resources. Each domain sets the price for its resources in their own currency. Users wishing to access information must obtain enough money in the correct currency to trade for access. No details are given on how each domain should set pricing or currency limits. It is unclear how the authors handle violation of perfect competi-

[2]See Appendix A.

tion (see Section 5.1), arbitrage, or currency exchanges. It is also unclear how each user is incentivized, and prices are not tied to risks, making it unclear what invariants the system operates under or what security properties are ensured.

One may distinguish these access control models by the tight integration of their exception models. The need to bypass access controls has been noted in the medical industry [22, 24, 65] and is a main motivation behind [30]. Most exception models are accomplished through delegation, such as role delegation in RBAC [64, 49, 65], or delegation in trust management, such as RT [40]. Nissanke and Khayat [50] suggest a way to reorganize an RBAC hierarchy using risk analysis for given delegation models. Jøsang [31, 32] developed a form of subjective logic based on uncertain information that has been used to dynamically determine role membership in RBAC [19, 33, 9]. We do not believe any of these systems provides the flexibility of the risk-based access control models proposed here, and none bound the amount of aggregate risk. These systems also only consider predefined exceptions, defining a small number of "gray area" exceptions and maintain a binary view of access control.

Economic and game theoretic models, similar to the ones proposed in this work, have been applied to other areas of computer science. For example, scheduling [39, 15] or storage utilization [8] in multi-agent systems, in peer-to-peer networks to handle bandwidth usage in BitTorrent [17] or reputation [51, 41], privacy-preserving algorithms [48], and many others. The area of research into the design of games and rules that achieve desirable properties, such as efficiency, truthfulness, or collusion resistance, is known as mechanism design. Mechanism design is a popular tool to illicit desirable behavior from users in a wide range of applications.

There has also been work on security systems that make multiple, non–binary decisions, such as *automated response* in intrusion detection and prevention systems [2]. Such systems can be viewed as responding differently to different attacks based on perceived levels of risk from the attacks beyond a traditional allow/deny system. This is similar to our idea of moving away from binary decisions [13] for access control, but the work in automated response does not calculate quantified risk estimates or make risk vs. benefit tradeoffs dynamically as a risk market would do.

3.1 Open Questions and Possible Solutions

All of the above risk-based access control systems [30, 14, 66] use risk tokens as a main enforcement mechanism. Both [30] and [14] 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– and 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 organization's expected utility[3], 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?

[3]A unit–less measure of desirability. See Appendix A.

3.2 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 [4] and UK [21], and CO_2 quota allotments by BP [43]. More details and examples are given in Appendix B.

4. 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.

4.1 Risk Calculation

An often cited source against risk-based access control is Cybenko's "Why Johnny Can't Evaluate Security Risk" [18]. 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. Traditional access control models, such as Bell-LaPadula, RBAC, DAC, ORCON, and the Chinese Wall, all require policy makers to make implicit assumptions regarding future risks with no regard or concept of risk as a shared, limited resource. Such systems result in a "tragedy of the commons" where the benefits and costs are unevenly distributed, and everyone loses [27]. These traditional access control systems make one of two tradeoffs; either the risks are low enough, or the benefits offset the risks. None consider the aggregate effect that small risks may have on an organization. The aggregate effect has been seen in many of the previous examples of Leeson, Kerviel, Ames, and Hanssen.

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 quantification allows an organization to bound the amount of aggregate risk they are willing to assume in a manner similar to the value at risk calculation performed by any financial institution. The damage limiting bounds will be present even if the risk calculations are not accurate; the uncertainty is present in both risk-based and traditional access control systems

Risk-based access control systems provide a general model for controlling access control decision making for any level of granularity. Access control is not unique to information security, allowing risk-based systems to be deployed in domains where risk metrics are well established, such as the financial sector and securities exchanges. Information security risk metrics is an open problem with ongoing research [54, 13, 12] and automated methods for calculating risk will be required for deployment of risk–based access control. Models such as FuzzyMLS [14] that compute quantified risk estimates are already used in coarse–grained risk–based access control; more accurate models and methods are required for finer–grained control.

4.2 Safety Analysis

Safety analysis is another criticism of risk-based access control systems. The standard safety analysis question, formulated by Harrison, Ruzzo, and Ullman [28], asks whether the access control system may transition into a state where

a given principal will be allowed to perform a given action. In risk-based access control systems, this may only be answered in the negative if the cost of the transaction exceeds the organization's risk tolerance. For extremely sensitive operations, damages may be defined as infinite, implying infinite costs, thereby unconditionally denying the access.

We argue, however, that this is the wrong safety analysis question to ask. Traditional safety analysis considers whether a single transaction represents excessive risks while ignoring the small risks associated with allowed transactions. Many adversaries leak information for the resources they have the required privileges [59, 61]. Switching metaphors, this is a "death of a thousand cuts."

At the most basic level, risk-based access control systems can ensure that, with a given confidence, the organization will observe a limited amount of harm, whether it is monetary, breach of confidentiality, integrity, availability, or other undesirable outcomes. In this work, we will illustrate how markets may add robustness and recovery or identification of malicious users to risk-based access control.

4.3 Mistrust of Markets

Since users are granted all requests for which they have enough risk capital, the allocation of risk becomes an increasingly important and difficult question that must be resolved. Markets open the system to new types of attacks, such as denial of service (refusal to sell or irrational purchases of risk), privilege escalation (irrational purchases of risk), collusion (coordinated attacks), arbitrage (purchases of risk with intentions for resale only), and inefficiencies due to employee incompetence or lack of market foresight.

We will address each of these issues in turn, commenting on the ease of the attacks, their impact, and the resources the market provides for detection. Many financial institutions have made a risk management decision to focus resources on detection and recovery of incidents over prevention to obtain desirable bottom lines [47]. We believe the efficiency with which the market distributes risk, the tools it provides for accessing user competence and malice, and the simplicity with which the users interact with the market offset any risks it represents, which, if desired, may be modeled, quantified, and accounted for.

5. AUCTIONS AND THE RISK MARKET

We now provide a short background on economic theory and definitions applicable to our risk-market approach. We will attempt to provide the intuition behind each concept and refer the reader to the references for more information. For a more complete discussion on the background economic theory, terms, and definitions, see Appendix A.

5.1 Auction Theory

The behavior of a market is determined by the relationship between 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. The intuition behind the marginal cost is producing larger quantities of items typically has lower amortized costs. For example, producing a single widget will cost five dollars, while producing a second will only cost four dollars more. Similar to a seller's cost curve, buyers define a demand curve D, based on the *marginal benefit* (MB), dictating the price they are willing to pay for a given quantity. Supply and demand curves are the sums of the marginal cost and benefit curves for all producers and consumers respectively.

A supply and demand graph can be seen in Figure 2. As the buyer's marginal benefit (D) increases from right to left, so does the demand and the price the buyer is willing to pay for a more limited quantity. As the seller's cost (MC) decreases from right to left, so does the price the seller is asking. 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. By allowing buyers and sellers to openly place bids and asks, as we approach the equilibrium quantity, we will approach the equilibrium price. 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.

Modern economics rarely gauges the efficiency of a market in terms of currency, but in terms of net gain in *utility*, a unit-less measure of desirability. The gray area in Figure 2 represents aggregate net gain in utility that can be observed through a redistribution of goods using the market. The seller's gain is the lower triangle bp^*c which is the difference between the sale price p^* and the production cost (line bc); the buyer's gain is the upper triangle ap^*c, which is the difference between the buyer's value of the goods (line ac) and the actual price p^*. 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'.

Beyond the equilibrium quantity, no sales are possible if all individuals are rational. A trader that is *individually rational* will always agree to a sale when it is beneficial and reject any sale that is detrimental given their current information[4].

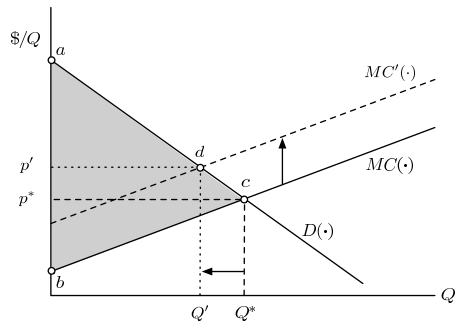

Figure 2: Supply and demand determine the equilibrium quantity Q^* and the equilibrium price p^*.

There will be a competition for the goods if the demand

[4]They *believe* it is beneficial regardless of whether it *is* beneficial.

for the goods exceeds the supply. If all traders are rational, the goods would go to the buyers who can get the most benefit from the goods since they will offer higher bids. In a risk market, the goods will be the limited amount of risk the organization is willing to take. The organization's goal is to get the most benefit by encouraging its employees to be prudent and take calculated risk where and when the perceived benefit is the greatest. It could do so by providing incentives to align the employees' interest with that of the organization, so employees who perceive greater benefit for the organization (and themselves) would acquire the risk tokens to pursue worthwhile opportunities. Of course, the risk market must be governed to address incompetence and malice. More details are provided in Sections 5.2, 5.4 and 6.

The ability of a market and the forces of supply and demand to efficiently distribute resources are well known. Many refer to this phenomenon as the *invisible hand*, a phrase coined by Adam Smith [53]. Under a strong assumption known as *perfect competition* (see Appendix A) the first and second fundamental theorems of welfare economics prove that a market will result in a *Pareto–optimal* allocation of resources [46]; no individual can increase their utility without decreasing the utility of another rational individual in a Pareto–optimal allocation.

Perfect competition will not hold in the risk market due to collusion among players, espionage or external utility and incentives, static risks, and a reliance on large numbers. Luckily, [42] 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. A *dominant strategy* maximizes an individual's utility regardless of the strategy played by other players. Using techniques attributed to Groves [26] we shall provide employees with the appropriate incentives to make optimal choices, regardless of the temptations to cheat the system. Individuals interact with the market using a strategy to determine how to place bids and asks, and the incentives should provide the employees with a *dominant strategy*. In the following sections, each of the above violations of perfect competition will be discussed.

5.2 Risk Market

Our risk market is based heavily on the supply chain internal markets of McAdams [45] and the incentives of Groves [26] 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. By limiting the number of risk tokens in circulation to fewer than R^* for each time interval, the organization can be assured they are operating within their risk tolerance.

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 AREA(abc) in Figure 3 is directly proportional to the organization's net gains.

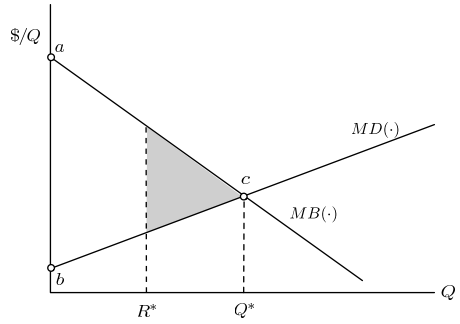

Figure 3: Risk Cost/Benefit. An Organization's Risk Tolerance R^* May Yield Suboptimal Gains.

The marginal damage and benefit functions allow the organization to adapt rapidly to new information and changing circumstances. For example, if new threats emerge the organization may change the shape of their marginal damage function and purchase risk tokens back from employees, and employees may realize some resources are less valuable than thought and sell surplus risk tokens.

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 commodity market, however there are many alternatives (see Appendix A.1). One such alternative is a Vickrey-Clarke-Groves (VCG) auction [58, 16, 26] which we will make use of in our simulations.

5.2.1 Initialization and Currencies

A risk-based access control system using a risk market will require a nominal setup for each organization, and each new employee. To limit the amount of liquid currency that may be used to purchase risk tokens on the market, a strictly internal currency will be used. This is similar to the "Linden Dollars" used in Second Life[5], or "Ithica Hours" used in Ithica, NY[6]. Physical monies do not need to be printed. Electronic cash [6, 11] may be used, but anonymity is not required, and not desirable, for the risk market. The organization may act as a centralized bank through which all transactions must filter.

An open question is how employees receive an initial sum of currency to trade on the market. Employees may be given a sums of money at given time intervals, or they may be extended a line of credit with which to trade. The balance of each employee's line of credit, or the amount they receive from the organization, may be based on the employee's job function, trustworthiness, clearance level, or any other attribute. We make no assumptions regarding how the currency is distributed to the employees. In the remainder of this work. We assume each employee receives an infinite line

[5] `http://lindenlab.com/`
[6] `http://www.ithacahours.org/`

of credit, and will be held accountable for their expenses at the end of each time period.

There are also many organizational structures governing who interacts with the market and what actions each employee may perform. Some organizations may choose to allow all employees to buy and sell risk tokens and exchange their tokens for access to resources. Other organizations may choose to designate market interactions as a management function. Managers will buy and sell risk tokens and deposit them into the accounts of the employees they manage. Market and resource access functions may also be kept distinct: employees who interact with the market (especially the security personnel) are prevented from spending risk tokens. Such separation of duty may help deter over bidding on the market and adds additional oversight to the access control system.

5.2.2 Marginal Damage

The marginal damage is the expected amount of damage to result from releasing an additional risk token into the market. To determine the marginal damage curve, we must be able to determine how much damage is expected for each risk token when used in different circumstances. The MITRE Jason Report [30] tokenizes both *risk*, the probability of compromise, in Phase 1 and *harm*, the expected amount of damage, in Phase 2. For example, risk is tokenized as

> *1 token = risk associated with one-day, soft-copy-only access to one document by the average Secret-cleared individual* [30],

and harm as

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

For example, assume an employee working at a bank wishes to credit an account $100. If the transaction is fraudulent, the damage is the $100 the bank loses. If there is a 0.1% chance that the transaction is fraudulent the bank can expect to lose $0.10 on the transaction. Damage values can be determined for any resource by an appraiser and is commonly done for insurance underwriting and litigation to determine damages.

By producing a limited number of risk tokens per document based on the potential impact, Phase 1 manages the risk on a per document basis. Phase 2 harm tokens manage aggregate risk across the entire system. This work may be applied to both concepts[7], allowing for different token types, however we focus on the single token, multiple resource model of Phase 2. Either way, we see that each token represents a constant amount of risk or 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 risk calculations, aggregation and inference, mistrust in the market, fear of agent collusion or espionage, or to model bounded rationality. To allow such flexibility securely we only require

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

[7]We acknowledge this results in ambiguity in the term "risk"

5.2.3 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, transference (insurance), or other parameters. Howard [29] has developed guidelines for determining an organization's risk tolerance in terms of total sales, net income, and equity, and the Bank for International Settlements has released standards for risk tolerance for the banking industry [3]. 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 loses 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 management techniques such as mitigation or transference of $Q^* - R^*$ through insurance [5]. Note that by controlling production and retaining the ability to buy back tokens, the organization can ensure it is always operating within its risk tolerance R^*.

5.2.4 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 hardware or software, 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 [45] uses Grove's Mechanisms [26] 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, and not maximize gains.

5.3 Assumptions

There are several assumptions that we make when considering the risk allocation problem. For simplicity we divide the assumptions into three categories.

Our first set of assumptions make the risk allocation problem more difficult.

1. Demand for expending risk far exceeds risk tolerance.

2. Employees may collude to increase their utility.

3. There are externalities such as espionage.

4. There are fixed risks regardless of the risk tolerance.

5. There is incomplete and imperfect information regarding opportunities and benefits.

Each of the above assumptions adds complication to the problem, and will be addressed when we describe our risk market implementation. If assumption 1 is false, then the risk allocation problem may be solved by the degenerate solution to allow all transactions. Allowing employees to collude and externalities such as outside incentives to cheat (espionage) encourages and rationalizes attacks. Fixed risks are risks to resources independent of resources access requests. Finally, in systems with complete and perfect information, the risk allocation problem may be solved optimally by a single entity, such as management. It is the incomplete information regarding the opportunities present in the system which makes risk allocation a difficult problem, and is one of the main advantages of markets.

Our second set of assumptions simplify the risk allocation problem.

6. Benefits, risks, and damages can be quantifiably measured or estimated within a reasonable bound.

7. Employees are rational.

8. There is no correlation between the benefit the organization observes for multiple transactions.

9. Risk tokens may be used immediately after creation and there are no risk token storage risks.

Assumption 6 is required at this stage of research, and solving it is outside the scope of this work. The Computing Research Association has identified information security risk metrics as one of four grand challenges in computer security [54].

We consider Assumption 7 to be reasonable at this stage of the research. Alternative models of human behavior may depend on the domain (medical, intelligence, financial, etc.), employee education level, experience, or presentation of information [35], adding complexity to the model. Additionally, there is no single accepted or agreed upon model for human behavior. Rational behavior is logical from a security standpoint since it makes any attack more difficult and less effective. Models that compensate for human irrationality may make attacks more effective for rational adversaries by compensating for human behavior[8]. Future work may consider other models of human behavior and reasoning.

Rationality is an interesting assumption in the setting of the risk market and access control and deserves additional discussion. Rationality is defined based on an individual's current information and her utility. Based on the information available to a *rational* individual, she will behave in ways to increase her utility and avoid behaviors that decreases her utility function. One's rational behaviors may seem irrational to others with different information or utilities. An organization usually trusts its employees to behave in a non–malicious way by giving them privileges. It has been suggested that such trust is a belief that an employee will behave irrationally in the view of the organization if her utility is purely based on measurable selfish gain [10], but will behave rationally if her utility also considers less measurable, more subjective concepts, such as norms, honesty, love, loyalty or friendship. By explicitly assuming employees are rational an organization may reward honest employees for their loyalty and remove an incentive that encouraged dishonest behavior in traditional systems. If cheating is the economically rational choice, traditional access control systems base their security in a belief (hope) of the users' benevolence. That is, most access control systems base their security on a belief that the subjects are benevolently irrational, while we assume they are maliciously rational.

Assumption 8 should be contrasted with assumption 2 in one key way: in assumption 8 we consider the benefit the organization obtains from their employees' actions, and in assumption 2 we consider the benefit an employee obtains from her actions, i.e., the organization does not benefit from collusion, but an employee may.

In assumption 9, risk token storage risks are risks to the integrity of the risk tokens. For example, theft or forgery are not considered in this work. Assumptions 8 and 9 have known solutions, and are often solved by the same mechanisms that address externalities [44]. Externalities are a more tangible and pertinent security problem.

A real implementation would require several properties which we assume in this work.

10. Risk tokens can only be used to access resources once.

11. New risk tokens are periodically released.

The first property is easily covered by the cryptographic cash literature [6], and the second is presented as a technicality since risk tolerance is a time dependent variable. An organization is willing to assume a given amount of risk per period (day, week, quarter, year, etc.) and will produce additional risk tokens as necessary within these bounds.

5.4 Formal Model

5.4.1 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. An organization can 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 [28].

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 each employee i, we can define the difference in capital, λ_i, and the market profit $\pi_{i,T}$, though without loss of generality we

[8]For example, risk aversion as modeled by prospect theory [35].

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, non–adversarial employee receives from the organization for expending his 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 will receive from the organization, where as u_i may include externalities if i is an adversary. Thus B could be used as an incentive to align the employee's benefit with the benefit of the organization. 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 $\overline{\pi}_{\setminus i}$.

5.4.2 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, aligning the organization and the employees, maximizing the organization's profits. For the risk market, we shall use an incentive structure similar to [45, 26]. 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 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. [45] and [26] 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 their own profit and decrease the wages of their peer groups. Groups of employees could be defined by many means, such as department or job function. Groups of employees who behave irrationally (refuse to sell or sell at a loss) will decrease their own wage, and increase the wage of other groups.
5. *Accurate Prediction Wage* $W_i = Y_i + \alpha\pi_i - \kappa(\pi_i - \tilde{\pi}_i)^2$. Where $\tilde{\pi}_i$ represents a postmortem analysis of employee i's estimate of her own market profit given the actual benefit she received and the price she paid for the resources (risk tokens). Thus $\tilde{\pi}_i$ should be closely aligned with the bids i placed. In Groves, $\tilde{\pi}_i$ is an agent provided estimation of the value of the resources, yet we may use this as an ex post value to determine the agent's competence. Employees have an incentive to bid as accurately as they can. Employees that consistently bid too high would operate at a net loss, with a large $(\pi_i - \tilde{\pi}_i)^2$ value that may be considered malicious. This type of incentive structure is useful both to encourage employees to make accurate estimates and bids on the market (simplifies bidding strategies) and as a tool to identify malicious activity.

6. 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.

6.1 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 6.3. Employees with a cooperative wage will be most adversely affected by reducing their market profit. The attack can be done in one of two ways: increasing the sale price of risk tokens to force honest employees to pay higher prices or by a denial of service (DOS) attack (see Section 6.2).

A cooperative wage provides the incentives against this 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. The technique discussed in Section 6.3 could be used to counter this attack if the incentive ϵ for malice (the externality) is large enough to offset the decrease in the malicious agent's 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 [57] 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.

6.2 Denial of Service (DOS)

An attacker may wish to prevent other agents from accessing resources by executing a denial of service attack. This can be accomplished by purchasing unnecessary risk

tokens reducing the liquid risk supply. This has the effect of preventing the least beneficial and most risky transactions because the risk tokens become too costly. If a unique risk token type is created for each resource (see Section 5.2.2), an attacker can target a single resource while their attack is undirected in the single risk token type model. In general an attacker is unable to target a given subject or resource.

A denial of service attack will have little impact and will be easy to detect. 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 6.3 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 and holding onto a large quantity of risk tokens. If the malicious agent purchases and holds on to Q' risk tokens, this will cause the organization to lose $\alpha \int_{Q^*-Q'}^{Q^*} MB(x) - MD(x)dx$ in profit. Other agents' market profits and incentives will also be reduced since fewer risk tokens will be available to them for making gains. To accomplish such an attack, the malicious agent must overbid sufficiently for the risk tokens and put himself in danger of being exposed (see Section 6.3); therefore ϵ must be large enough. If ϵ is small and the malicious agent has no incentive to overbid too much, rational agents can overcome the malice by behaving honestly and maximizing the organization's benefit and their incentives.

6.3 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, which are the same as fixed risks in our model, each agent states their predicted gains, $\tilde{\pi}_i$ for the resources. To ensure accurate predictions, McAdams deducts an amount proportional to the error in the agent's estimations $\kappa(\pi_i - \tilde{\pi}_i)^2$. Agents in the market with external incentives $\epsilon_i(\delta)$ may rationally purchase additional risk tokens to access unnecessary resources, escalating their privileges. For small incentives ϵ_i, the accurate prediction wage may discourage and compensate for such behavior. For large external incentives, the agent may still behave rationally by making excessive bids on the market.

While the agent's final wage $\hat{W}_i$ may be maximized by making such purchases, their market wage will reflect their losses (the amount spent purchasing risk tokens exceeded the amount of benefit their received from them). When we require the employees to provide gains estimates (for example, when dealing with fixed costs) we can determine their accuracy; otherwise it may be estimated based on the equilibrium price or market logs and the benefit received. Malicious agents with external incentive ϵ_i will have large discrepancies between their estimates and actual market wages. By observing these discrepancies, the malicious agents may be identified. We illustrate how the market logs may be used to identify the malicious agents in Section 7.2.5.

6.4 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. [45] notes that peer group relative wages will remove the advantages of colluding.

6.5 Greedy Employees and Market Prowess

One 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. Experimental economics work indicates that even random yet rational (constrained zero-intelligence (ZI-C)) bids will converge to optimal distributions [25]. See Appendix A.2 for more details.

6.6 One Time Attack

Risk-based access control systems actually limit the extent of damage a successful attack may have compared to traditional access control systems. Consider a privilege escalation problem in a risk-based access control system and a traditional access control system. Privilege escalation in a traditional access control system may result from errors in policies or employees changing job roles without old privileges being revoked. This is a common oversight in practice [59].

Upon a successful privilege escalation in a risk-based system, a user obtains single resource access which they would normally be denied, and the attack must be executed in full for each additional access. In a traditional system, once the adversary has obtained the desired privileges, they may access any and all resources for which their escalated privileges are sufficient. Thus the attack (or policy error) need only be executed once.

6.7 The Winner's Curse

Humans cannot be expected to behave rationally (one of our assumptions). Experiments by [34] found that real players often overbid in second-price sealed bid auctions[9], while

[9]The dominant strategy is to bid the actual value of the good being auctioned.

[23] showed that even though the players did not converge to the game theoretic optimal strategy, the efficiency still improved with successive experiments. The phenomenon where players often overbid (and thus overpay) for an item is known as the *winner's curse*. Overbidding often results from overestimations for the value of the good being sold, or irrational incentives to win, such as the emotional pleasure from winning.

In most auctions with a winner's curse, the item is of approximately equal value to all players, while in the risk market the players are only indirectly measuring the value of risk. Rather they are measuring the value of the resource they may trade it for. The incentives described in Section 5.4.2 combined with experience should discourage players from overbidding and succumbing to the winner's curse.

7. 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 a "Yes/No" access control system that represents a binary decision model based on risk thresholds. 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.

7.1 Simulation Design

We model a quantified risk-based access control system similar to the one described in [30], 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 [14] and the redistribution capital in [66].

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 where the cost in tokens is equal to the harm and hard upper bounds such as those discussed in [14] 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 with respect to the number of documents accessed. The intuition is that employees will access the most relevant and beneficial resources first.

7.1.1 Preallocation

We consider three possible preallocation methods where the risk tolerance is treated as a budget:

1. *Risk Level* - We use the inverse of an agent's risk value to obtain a trust value. A clearance level (similar to MLS) 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 of the organization's risk tolerance.

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 [42]. In real instantiations, risks may be traded by management and deposited into employee accounts, relieving the employees from the requirement of interacting on the market.

7.1.2 Market Strategy

The incentives from Section 5.4.2 encourage rational individuals to behave honestly when interacting with the market. For our simulation, agents place bids and asks as constrained zero intelligence (ZI-C) traders similar to [25], and discussed in Appendix A.2. 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 around the actual benefit for the resource and will bias an agent to bid higher rather than lower, a phenomenon often observed in experimental economics [7]. Employees with low competence may exhibit behavior similar to an employee that is irrational in random ways (doesn't always bid high or low), allowing us to test the limitations on our rationality assumption.

An agent's marginal benefit function is defined using one of two ways for each simulation[10]:

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.

[10] An simulation is run several times and its result is averaged over these runs.

7.1.3 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.

The optimal solutions are equivalent to the decision making process of a central entity (such as management) with perfect and complete information. VCG is a known way to obtain accurate utility functions from the agents who would otherwise have an incentive to be misleading, though it was previously noted that VCG is incapable of solving the risk allocation problem in practice [52].

7.1.4 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 tolerance, the marginal damage constant (Equation 3), and the average rate of drop in the benefit.

7.2 Monte Carlo Simulation Results

To test our hypothesis, we conducted multiple simulations. In each simulation multiple parameters were allowed to vary independently. In the following figures, each point on the graph represents the average value for multiple runs of the simulation for each possible set of parameters, or states. For example, random instances of an organization with n employees, accessing m resources with benefit and damage distributions ϕ and γ are averaged and plotted as a single point. Since values are chosen at random from a given range, each point illustrates a unique, random possible universe in which we may be in. The uncertainty regarding the distribution of parameters prevents us from selecting any given instance (point) as being the correct model for the ground truth.

7.2.1 Comparison with Yes/No Access Control

As one of our alternative access control models, we simulated an access control system that makes Yes/No binary decisions by comparing a subject's level of trust and an object's level of sensitivity. To accomplish this, we first assume that an agent's trust level and a document's damage (sensitivity) 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.

The idea is to capture the essence of traditional access control models with binary decision making abilities. This simple model is a multilevel model that resembles the Bell-LaPadula (BLP) model but without categories, though it should not be considered a direct comparison with the BLP model. We do believe our results may be directly applied to BLP. Our simplification captures the risk (and trust) aspects of BLP while leaving out the "need to know". It is our view that employees without a need to know are not trusted enough (the risk outweighs the trust) to access a resource. This concept is still captured. While categories do provide more control and security, trusted persons are often given enough access (Walker [20], Ames [61], Hanssen [59], etc.) and a lot of sensitive information may be stored in a single category.

By converting the parameters from the risk-based system into a multilevel system, we are able to make direct comparisons using identical resources and agents. Note that when simulating the multilevel system an agent's risk (clearance level) does not change, and the system has no concept of risk tokens, thus risk tokens are not used or considered when making access control decisions, just classification and clearance. To make a direct comparison to our risk-based system, we log how much risk the multilevel system expends. We also assume each resource is accessed only a single time.

Our results are shown in Figure 4. Figures 4(a-c) are for the multilevel system, and (d-f) are for our risk-based access control system using the risk market. In each figure, the horizontal axis is the mean potential damage for any transaction, presenting the impact if a resources is leaked or compromised. The vertical axis for the Figures 4(a,d) are the aggregate benefit obtained by the organization, (b,e) the aggregate harm from loss or compromise, and (c,f) the net gain (benefit minus harm). It should be noted that the vertical axis for Figures 4(a-c) are on much larger scales than those of Figures 4(d-f).

It is clear to see that the multilevel system yields greater benefit, but significantly greater harm than the risk-based systems, and it is more erratic (high volatility) and more likely to yield a net loss. The multilevel system also has a lower return on risk investment (ROI), which is not shown, but is implicit from Figures 4(c,f).

One of the main motivations behind risk-based access control systems, and the very motivation behind the commission of [30], is to allow information sharing securely to users who would otherwise be untrusted. From Figures 4(a,d) (and (c,f)) we can see that information sharing halts in the multilevel system when the damage exceeds a given level (classification exceeds clearance) yet by taking a quantified amount of risk, the risk-based access control systems may continue to produce desirable outcomes. Aggregate harm in the multilevel system, as shown in Figure 4(b), depends on the damage, which in turn depends on the number of allowed accesses, an illustration of the tragedy of the commons for information security. But aggregate harm in risk-based systems, as shown in Figure 4(e), does not depend on the damage, it is only highly correlated to the amount of risk.

7.2.2 Effect of Risk Allocation on Efficiency

Next, we would like to compare the various preallocation schemes with the market based distribution mechanisms. As stated, a market's efficiency is a measure of the extracted utility compared to an optimal distribution, which is determined using VCG.

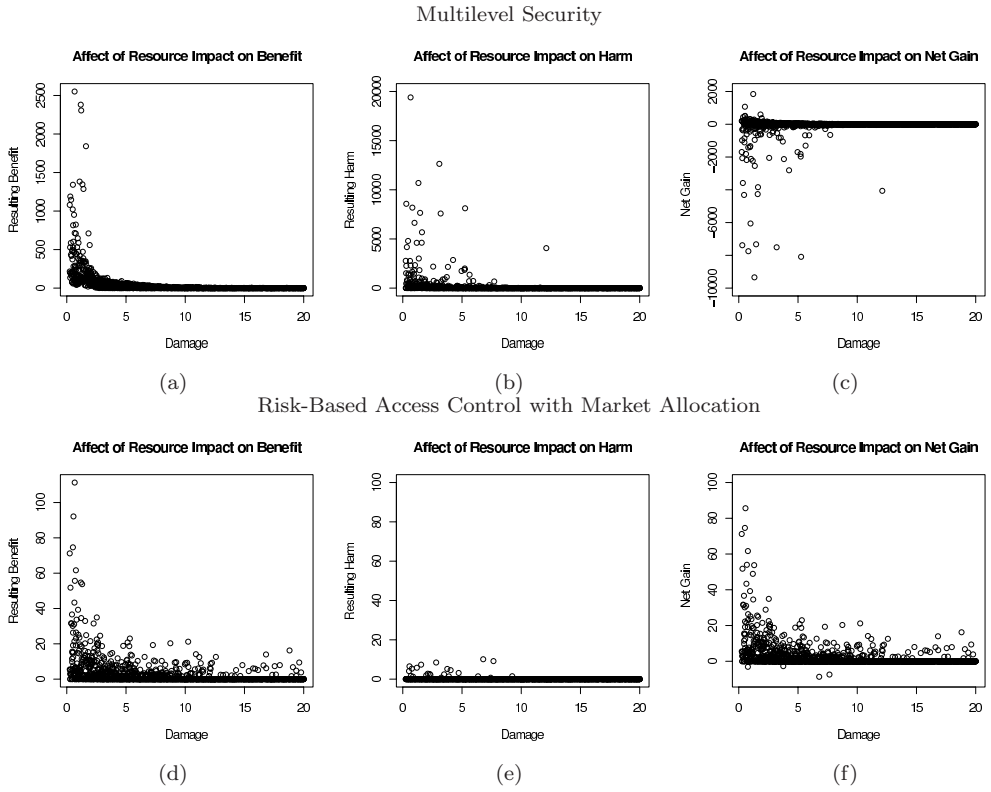

Figure 4: Multilevel Security (a-c) Versus a Risk-Based Access Control System (d-f) using a Risk Market

Since the market mechanisms, even given the zero-intelligence traders described in Appendix A.1, are known to lead to efficient allocations, the new results from our risk market shows how the agent's ability to estimate benefits affect the efficiency. Figure 5 show the effect of competence on the efficiency of the market. By performing a linear regression, we can see that efficiency reaches 90% when the mean competence is around 42% and standard deviation from 5-25%.

We also consider the effect of distribution methods on the amount of risk expended and efficiency of preallocation schemes. Preallocation methods performed extremely poorly in both categories (not taking risks and not performing well), and on average only used 15% of the risk tolerance R^*. The low utilization is due largely to a small surplus of risk tokens held by each agent, individually incapable of accessing additional resources. The inability to redistribute risk is unlikely to be as severe in practice, yet indicates that preallocation is a poor choice.

Conversely, the VCG mechanism consistently allocated 98-99% of R^* (the risk tolerance), while the market allocated 92-95%. Many risk tokens may be purchased by agents using the market and not used since they are not able to acquire the sufficient number to access the desired resources. The result is a single agent holding a surplus of risk tokens. In real scenarios not modeled here but discussed in [44], a rational agent may be able to increase their utility by trading those tokens at a loss if they may be used for some benefit by another employee. This is possible in scenarios where agents receive a cooperative wage.

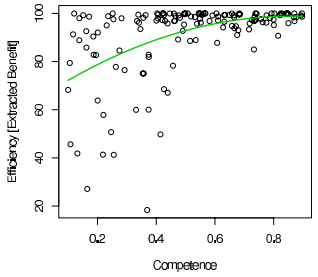

Figure 5: The Affect of Competence on Market Performance. A Linear Regression is Shown.

7.2.3 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 5.2.4 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^* (equilibrium quantity) 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 5.2.2 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 increas-

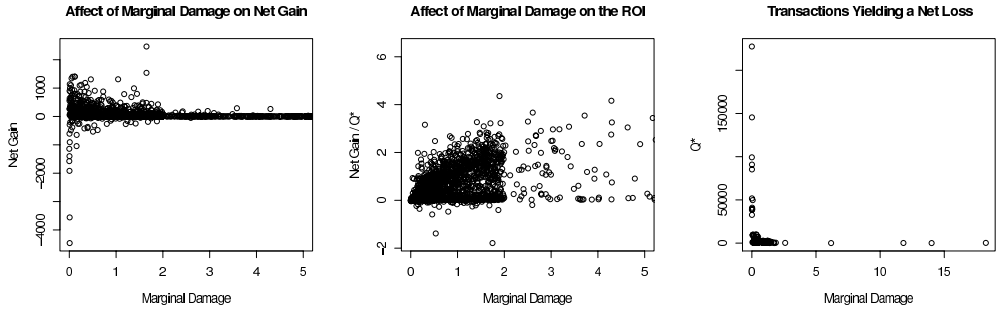

(a) (b) (c)

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.

ing the marginal damage too high, employees are starved and prevented from accessing valuable resources, and information sharing drops to zero. In Figure 6(b), we illustrate the return on investment (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.

7.2.4 *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[11]. We modeled an upper bound and classified runs of the simulator into two categories: net loss and net gain. For these experiments, all R^* risk tokens were released, corresponding to $MD = 0$ and employed a hard boundary.

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 tolerances and hard boundaries. While hard boundaries are insufficient for ensuring net gains, they are beneficial for restricting the long negative tail that is present when hard boundaries are not used. Figures 7(a) and 7(b) show the distribution of net gains without and with hard boundaries respectively[12].

7.2.5 *Effect and Identification of Malicious Agents*

It should be clear to see that the total aggregate amount of harm when using risk based access control systems is dependent only on the total risk expenditure Q^*, which the organization may limit to R^* by restricting risk token production. The proof relies on the assumption that risks and damages are known and quantified to a sufficient precision. We would like to see what impact a small number of malicious employees may have on the system if they violate our quantifiable risk assumption, or receive external incentive.

To model malicious agents, we assume they will receive an external benefit for the resources they leak that is a function of the resource's impact. Thus the more damaging the resource, the greater the incentive to leak it. We augmented our simulation to add malicious agents that receive such an external incentive. The incentive is a fraction $(0 < \zeta \leq 1)$ of the resources impact. In our simulation thus far, risk estimates are known to be within one standard deviation of their actual value (we allow for errors and uncertainty in risk estimations). In this section, we assume arbitrary errors in risk estimates for malicious agents only. That is, the actual probability is 100% for a malicious agent to cause damage, yet the estimation of the probability made by the risk based access control system is in the standard range $(0, 100\%]$.

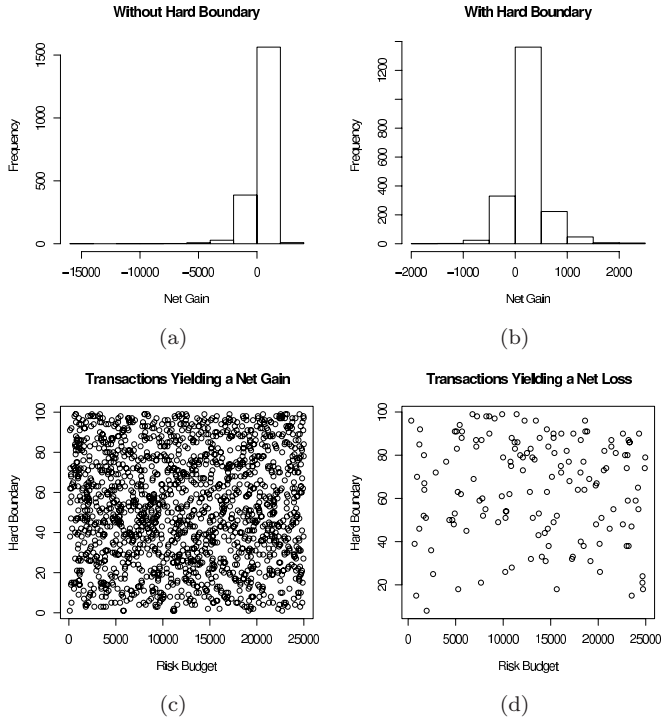

(a) (b) (c) (d)

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

[11] And noting this region may be unbounded.

[12] The scale on the X-axis of Figure 7(a) is much larger than that of Figure 7(b)

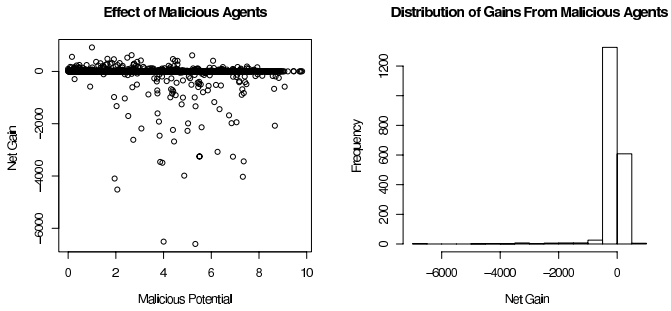

(a) (b)

Figure 8: Impact from malicious agents and a violation of the assumption of quantifiable risk.

Figure 8 shows the potential impact a small number of malicious agents may cause. In Figure 8(a), the horizontal axis represents the malicious potential: the percentage of employees that are malicious (from 0-1%), times the fraction ζ of the resources impact they will receive. In these experiments we restrict ζ from 5-10%. Hard boundaries and risk tolerances were chosen at random to maintain consistency with the Monte Carlo simulations. Figure 8(b) shows the distribution of net gains. We can see that a small number of malicious agents are still capable of causing a large amount of damage (albeit we must violate our assumption on quantifiable risks), though the distribution is centered strongly around zero net gain or loss.

The attacks in the above simulations are no different than those seen in practice that are executed in multilevel security systems by Hanssen and Ames. However, the risk market allows for rapid detection of the few malicious agents who performed such attacks. This is the same risk management philosophy used in many financial institutions that have shifted from prevention to detection and resolution of incidents [47].

Next, we illustrate how the market data may be used to identify malicious employees using the accurate prediction wage incentives. For rational agents, the benefit they receive from accessing resources must exceed the cost to purchase the total number of required risk tokens. When the agent will receive an external benefit, some transactions that were once irrational may now be rational; the external incentive offsets their internal operating loss. By observing the agent's market wage (π_i, or the benefit the resources provided to the organization minus the price paid for risk tokens), we may detect malicious agents. Since agents place random bids in our simulation, the market wage is estimated from the marginal cost to produce risk tokens and the number of risk tokens purchased, i.e., the agent must pay a minimum of the marginal cost times the number of tokens purchased. In these simulations the risk tolerance is determined by the market.

Figures 9(a,b) illustrate how the market wage of malicious agents is different from their honest (yet potentially incompetent) counterparts. In these plots, each point represents a single agent; the red triangles are malicious, and the black circles are honest. The horizontal axis represents the agent's market wage, while the vertical axis is the benefit the organization obtains from allowing the access. If the interests of the organization and its employees are perfectly aligned, there should be a strong correlation between these values (they should be clustered around and above the line $y = x$).

Figure 9(a) shows the long tail distribution (strong incentives to misbehave) present in the system for malicious employees. Figure 9(b) illustrates the distinction between malicious and honest agents in more detail. While there are many honest agents with negative market wages, they are likely incompetent, and only appear marginally irrational compared to malicious agents. Classifying agents based on market wage using fuzzy sets, neural networks, support vector machines, or other mechanisms should work well.

To test this theory, we classified the agents in Figure 9 using a support vector machine (SVM) into the two classes: malicious, and honest. To train the SVM, we selected a random set of 2000 agents. The SVM was provided the classification (malicious, honest), the market wage (π_i) and the organization's benefit. We then tested the data on the remaining agents. Agents that made no purchases were removed from the training and testing sets. Without tuning the SVM, we were able to classify 93.6% of the agents correctly.

In addition, the risk estimation may use the market wage and organizational benefit from previous time intervals to

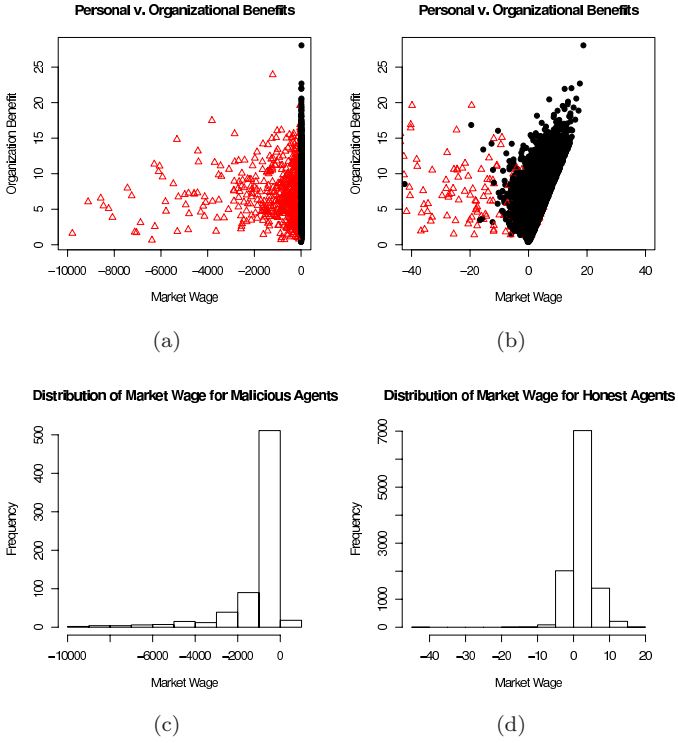

(a) (b)

(c) (d)

Figure 9: Malicious agents (red triangles) may be distinguished from honest agents (black circles) based on their market wage. The full distribution (a) illustrates the extent to which they have diverge in market behavior. Figure (b) shows the blurred line between incompetent and malicious employees. Figures (c,d) show the distribution of market wage for the group of malicious and honest employees.

update risk estimates. Employees with negative market wages obtain higher risk estimates for future transactions, and are penalized for their incompetence or alleged malice.

8. 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 intrusion detection system, 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 [30] 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.

9. ACKNOWLEDGMENTS

We would like to thank Grant M. Wagner and Angela S. Reninger for suggesting to us the potential usage of market mechanism in risk management, they also worked with us on our earlier research on risk management. Claudia Keser taught us the basics of market mechanism and helped us in modeling the risk market. We would also like to thank Angelos Keromytis, Steven Greenwald, and the other NSPW attendees for their useful comments and feedback.

10. REFERENCES

[1] FCC: Wireless telecommunications bureau, October 2007. `http://wireless.fcc.gov`.

[2] I. Balepin, S. Maltsev, J. Rowe, and K. Levitt. Using Specification–Based Intrusion Detection for Automated Response. In *Sixth International Symposium on Recent Advances in Intrusion Detection (RAID)*, 2003.

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

[4] J. Bingaman, A. Specter, T. Harkin, T. Stevens, L. Murkowski, and D. Akaka. Low Carbon Economy Act of 2007. Technical report, United States Congress, 2007. Proposed.

[5] R. Böhme and G. Kataria. Models and measures for correlation in cyber-insurance. *The Fifth Workshop on the Economics of Information Security (WEIS 2006)*, June 2006.

[6] S. Brands. Untraceable off-line cash in wallets with observers. *Advances in Cryptology — CRYPTO'93*, 1993.

[7] D. Brenner and J. Morgan. The Vickrey-Clarke-Groves versus the simultaneous ascending auction: An experimental approach. *A1.133 WP 188*, 1997.

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

[9] V. Cahill, E. Gray, J.-M. Seigneur, C. D. Jensen, B. Shand, N. Dimmock, A. Twigg, J. Bacon, C. English, W. Wagealla, S. Terzis, P. Nixon, G. di Marzo Serugendo, M. Barbone, K. Krukow, and M. Nielsen. Using trust for secure collaboration in uncertain environments. *Pervasive Computing*, 2003.

[10] C. Castelfranchi and R. Falcone. Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In *ICMAS '98: Proceedings of the 3rd International Conference on Multi Agent Systems*, page 72, Washington, DC, USA, 1998. IEEE Computer Society.

[11] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *CRYPTO '88: Proceedings on Advances in cryptology*, pages 319–327, New York, NY, USA, 1990. Springer-Verlag New York, Inc.

[12] P.-C. Cheng and P. A. Karger. *Risk Modulating Factors in Risk-Based Access Control for Information in a MANET*. IBM T.J. Watson Research Center, February 2008. IBM Research Report RC24494 (W0802-051).

[13] P.-C. Cheng and P. Rohatgi. *IT Security as Risk Management: A Research Perspective*. IBM T.J. Watson Research Center, April 2008. IBM Research Report RC24529 (W0804-015).

[14] 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.

[15] G. Christodoulou, E. Koutsoupias, and A. Kovács. Mechanism design for fractional scheduling on unrelated machines. In L. Arge, C. Cachin, T. Jurdzinski, and A. Tarlecki, editors, *ICALP*, volume 4596 of *Lecture Notes in Computer Science*, pages 40–52. Springer, 2007.

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

[17] B. Cohen. Incentives build robustness in bittorrent. *NA*, 2003.

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

[19] N. Dimmock, A. Belokosztolszki, D. Eyers, J. Bacon, and K. Moody. Using trust and risk in role-based access control policies. In *SACMAT '04: Proceedings of the ninth ACM symposium on Access control models and technologies*, pages 156–162, New York, NY, USA, 2004. ACM.

[20] P. Earley. *Family of Spies: Inside the John Walker Spy Ring*. Bantam Books, 1988.

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

[22] M. Evered and S. Bögeholz. A case study in access control requirements for a health information system. In J. M. Hogan, P. Montague, M. K. Purvis, and C. Steketee, editors, *ACSW Frontiers*, volume 32, pages 53–61. Australian Computer Society, 2004.

[23] R. Garratt and J. Wooders. Efficiency in second-price auctions: A new look at old data. Technical report, Department of Economics, UCSB, February 2004.

[24] K. Garson and C. Adams. Security and privacy system architecture for an e-hospital environment. In K. E. Seamons, N. McBurnett, and T. Polk, editors, *IDtrust*, volume 283, pages 122–130. ACM, 2008.

[25] 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.

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

[27] G. Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, December 13 1968.

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

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

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

[31] A. Jøsang. Trust-based decision making for electronic transactions. *Fourth Nordic Workshop on Secure Computer Systems (NORDSEC'99)*, 1999.

[32] A. Jøsang. A logic for uncertain probabilities. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 9(3), 2001.

[33] L. Kagal, J. Undercoffer, F. Perich, A. Joshi, and T. Finin. A security architecture based on trust management for pervasive computing systems. Technical report, Maryland University Department of Computer Science and Electrical Engineering, 2005.

[34] J. Kagel and D. Levin. Independent private value auctions: Bidder behavior in first-, second- and third-price auctions with varying numbers of bidders. *Economic Journal*, 103:868–879, 1993.

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

[36] Knowledge@Wharton. How we got into the subprime lending mess, September 19 2007. `http://knowledge.wharton.upenn.edu/article.cfm?articleid=1812#`.

[37] Knowledge@Wharton. Victimizing the borrowers: Predatory lending's role in the subprime mortgage crisis, February 20 2008. `http://knowledge.wharton.upenn.edu/article.cfm?articleid=1901&CFID=66192274&CFTOKEN=98696674&jsessionid=a83075bc94e55b3352f5`.

[38] B. W. Lampson. Protection. *Operating Systems Review*, 8(1):18–24, Jan. 1974. initially appeared in Proceedings of the Fifth Princeton Conference on Information Sciences and Systems, Princeton University, Princeton, NJ, USA, March 1971, pp. 437–443.

[39] R. Lavi and C. Swamy. Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. In J. K. MacKie-Mason, D. C. Parkes, and P. Resnick, editors, *ACM Conference on Electronic Commerce*, pages 252–261. ACM, 2007.

[40] N. Li and J. Mitchell. Rt: A role-based trust-management framework, 2003.

[41] R. T. B. Ma, S. C. M. Lee, J. C. S. Lui, and D. K. Y. Yau. A game theoretic approach to provide incentive and service differentiation in p2p networks. In E. G. C. Jr., Z. Liu, and A. Merchant, editors, *SIGMETRICS*, pages 189–198. ACM, 2004.

[42] L. Makowski and J. M. Ostroy. Vickrey-Clarke-Groves mechanisms and perfect competition. UCLA Economics Working Papers 333, UCLA Department of Economics, July 1984. Available at `http://ideas.repec.org/p/cla/uclawp/333.html`.

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

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

[45] D. McAdams and T. W. Malone. Internal markets for supply chain capacity allocation. Technical Report 4546-05, MIT Sloan School of Management, 2005. MIT Sloan School of Management Working Paper No. 4546-05 and MIT Center for Coordination Science Working Paper No. 224.

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

[47] G. McGraw. Silver bullet speaks with dan geer. *IEEE Security & Privacy*, 4(4):10–13, 2006.

[48] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *FOCS*, pages 94–103. IEEE Computer Society, 2007.

[49] S. Na and S. Cheon. Role delegation in role-based access control. In *ACM Workshop on Role-Based Access Control*, pages 39–44, 2000.

[50] N. Nissanke and E. J. Khayat. Risk based security analysis of permissions in rbac. In E. Fernández-Medina, J. C. H. Castro, and L. J. García-Villalba, editors, *Proceedings of the 2nd International Workshop on Security In Information Systems (WOSIS)*, pages 332–341. INSTICC Press, 2004.

[51] T. G. Papaioannou and G. D. Stamoulis. Reputation-based policies that provide the right incentives in peer-to-peer environments. *Computer Networks*, 50(4):563–578, 2006.

[52] M. H. Rothkopf. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research*, 55(2):191–197, March–April 2007.

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

[54] E. H. Spafford, R. A. DeMillo, A. Bernat, S. Crocker, D. Farber, V. Gligor, S. Goodman, A. Jones, S. Landau, P. G. Neumann, D. Patterson, F. Schneider, D. Tygar, and W. Wulf. Four grand

challenges in trustworthy computing. Technical report, Computing Research Association, November 16–19 2003.

[55] S. Sunder. *Experimental Asset Markets: A Survey*, chapter 6, pages 445–500. Princeton University Press, 1995.

[56] P. P. Tallon. Critical steps in storage management: How business requirements shape policy decisions. Technical report, GlassHouse Technologies, 2003. `http://www.dscon.ru/docs/wp_critical_steps_stor_mgmt_web.pdf`.

[57] C. Thorpe and D. C. Parkes. Cryptographic securities exchanges. *Financial Cryptography and Data Security (FC07)*, February 2007.

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

[59] D. A. Vise. *The bureau and the mole: the unmasking of Robert Philip Hanssen, the most dangerous double agent in FBI history*. Atlantic Monthly Press, 2002.

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

[61] R. J. Woolsey. The Aldrich H. Ames case: An assessment of CIA's role in identifying Ames as an intelligence penetration of the agency, October 21 1994. `http://www.loyola.edu/dept/politics/intel/hitzrept.html`.

[62] Y. Yemini, A. Dailianas, and D. Florissi. Marketnet: Using virtual currency to protect information systems. *ECDL '98: Proceedings of the Second European Conference on Research and Advanced Technology for Digital Libraries*, pages 891–902, 1998.

[63] Y. Yemini, A. Dailianas, D. Florissi, and G. Huberman. Marketnet: Market-based protextion of information systems. *Proceedings of ICE'98, First International Conference on Information and Computation Economics*, October 1998.

[64] L. Zhang, G.-J. Ahn, and B.-T. Chu. A rule-based framework for role based delegation. In *SACMAT '01: Proceedings of the sixth ACM symposium on Access control models and technologies*, pages 153–162, New York, NY, USA, 2001. ACM.

[65] L. Zhang, G.-J. Ahn, and B. tseng Chu. A role-based delegation framework for healthcare information systems. In *SACMAT*, pages 125–134, 2002.

[66] L. Zhang, A. Brodsky, and S. Jajodia. Toward Information Sharing: Benefit And Risk Access Control (BARAC). *Seventh IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'06)*, 0:45–53, 2006.

APPENDIX

A. AUCTION THEORY

Adam Smith's *invisible hand* is the idea that in a free market, an individual acting selfishly attempting to maximize their wealth will benefit the community as a whole [53]. Smith argues that the wealth of a society was the sum of its parts, thus by maximizing each individual, society's wealth as a whole is maximized. Modern economics rarely works in terms of currency, but in measures of *utility*, a unit-less measure of desirability. Utility allows for more accurate modeling of real world incentives, and is a solution to paradoxes in expected value theory such as the St. Petersburg paradox.

The first fundamental theorem of welfare economics can be seen as a formalization of Smith's invisible hand [46]. It states that under ideal conditions, a competitive market will converge towards an efficient allocation of resources, such as money or other assets. These ideal conditions are known as *perfect competition*, where the buying decisions of agents cannot affect the market price. The *efficiency* of a market is a measure of the extracted value, or utility, versus an optimal allocation. For example, imagine a system with two parties, a single producer, and a single consumer. The producer creates widgets at a cost of five dollars that have a value of ten dollars to the consumer. By selling the widget for seven dollars, the utility of the system has increased by five dollars; two for the producer, and three for the consumer.

Some efficient allocations are particularly attractive due to their stability. For example, in a *Pareto-optimal* allocation, any redistribution that is beneficial to one individual is detrimental to at least one other. In *Pareto-optimal* allocations, there is a disincentive to change the allocation, making it stable. The second fundamental theorem states that by performing a lump sum redistribution of the assets to the players and allowing the market to take over will result in a *Pareto-optimal* allocation [46].

Finally, we need to discuss what each player knows about the system, and how they make decisions. A game has *imperfect information* if each player does not know what actions every other player has taken in the past, and *incomplete information* if each player does not know the payoff (utility or value function) of every other player. A *dominant strategy* will maximize an individual's utility regardless of the strategy played by other players. The market may be considered a game in game theory where the seller chooses from a series of "ask" actions and the buyer plays "bid" actions.

While the theorems of welfare economics indicate that free markets are capable of solving the resource-distribution problem, [42] proves that perfectly competitive markets are the only Pareto-optimal, individually rational, dominant strategy solution given imperfect, incomplete information.

A.1 Market Mechanisms and Auction Theory

Much of the literature on dominant strategy mechanisms are on based on the works of Vickrey [58], Clarke [16], and Groves [26] and are termed Vickrey-Clarke-Groves (VCG) mechanisms. A Vickrey Auction [58] 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 [16] 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 individual's assigned output and actual output. Groves [26] 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 the winning agent of a bid pays the highest amount that would have been bid for the objects had the agent's bid 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 and an optimal pricing policy. While VCG mechanisms are extremely attractive in theory, they do not work well in practice. Rothkopf [52] 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.

A.2 Double Auction

A standard free market, such as a stock or commodities market, are known as double auctions. Sunder [55] 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 [55].

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 [25] 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 [25]. These results are encouraging. When combined with the appropriate incentives, such as those discussed in [26, 45], agents need only be rational for the market to perform well.

B. 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 [4] and the United Kingdom Climate Change Bill [21]. 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 [43] 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 [55]—which is discussed in Section A.1—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 [43, 45]. HP has performed similar experiments to forecast sales figures for their printer devision [43] and as efficient scheduling algorithms for utility data centers[13].

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[14].

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

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