Omnichannel Cybersecurity: Optimizing Security by Leveraging Asymmetric Motivation

Spencer Oriot
Booth School of Business
The University of Chicago
Chicago, Illinois, USA
spencero@chicagobooth.edu

Adam Williams
Booth School of Business
The University of Chicago
Chicago, Illinois, USA
adam.williams@chicagobooth.edu

Josiah Dykstra
Cybersecurity Collaboration Center
National Security Agency
Ft. George G. Meade, Maryland, USA
josiah.dykstra@cyber.nsa.gov

ABSTRACT

The current incentive structure of internet platforms hinders system security, imposes unnecessary costs on end users, and further entrenches the status quo. Establishing and maintaining effective cybersecurity without significant changes to system infrastructure could be much faster, cheaper, and easier if end users could contribute more to their own protection. However, evaluating the relative costs and benefits of investments in information security is not easy. Not only that, but current models of investment for firms do not suggest what to invest in, leaving firms to figure it out individually. In addition, the social costs of poor information security and potential benefits from coordination are often omitted from the analysis.

In this paper, we propose a new theoretical framework for omnichannel cybersecurity which draws on cyber supply chain risk management and analytic marketing principles. The framework recalibrates incentive structures, accurately allocates cost to risk, and determines the optimal set of security measures for each individual based on their perceived cost-benefit, threat profile, and real-time attack status within the context of the system as a whole. We expand the applicability and scale of the Gordon-Loeb model to continuous time using the Kelly Criterion. Our proposals are informed by results from 40 interviews with stakeholders, and we describe the benefits of the framework to their needs. The result is a Cybersecurity Utility Platform to manage supply chain risk with indirect coordination for a more secure and more resilient ecosystem for all participants in the ecosystem.

CCS CONCEPTS

• **Security and privacy** → **Network security**; • **Applied computing** → **Economics**.

KEYWORDS

cybersecurity, economics, incentives, supply chain

ACM Reference Format:
Spencer Oriot, Adam Williams, and Josiah Dykstra. 2021. Omnichannel Cybersecurity: Optimizing Security by Leveraging Asymmetric Motivation. In *WEIS '21, June 28–29, 2021, Virtual.* , 15 pages.

WEIS '21, June 28–29, 2021, Virtual
2021.

1 INTRODUCTION

Evaluating the relative costs and benefits of investments in information security is a difficult task [17]. Moreover, security is a complex problem and current models of investment for firms do not suggest what to invest in, leaving firms to figure it out individually. In addition, the social costs of poor information security and potential benefits from coordination are often omitted from the investment analysis. This shortcoming exist in part because it is difficult to systematically account for third parties that are not always readily apparent in the environment. In this paper, we address each of these issues by describing a theoretical framework to improve on existing models and drive further research.

From small firms to internet scale, the openness of the information and communication technology (ICT) infrastructure has two significant security implications. First, any threat or vulnerability will continue to be a threat or vulnerability for affected users until fixed, even if risk-reduction or mitigations emerge. Second, new threats and vulnerabilities will continually be discovered. Thus, information security investment should be thought of as consisting of two distinct components: one part protecting against known threats, one part detecting, communicating, and defending against new and emerging threats.

Much of the existing literature on economics and information security is oriented toward establishing the optimal cost-benefit investment strategy as a static value at one point in time. This value is the cost of protecting against known threats. On the other hand, cyber supply chain risk management (CSCRM) is primarily concerned with the ongoing processes of managing the risk of cyberattack throughout a supply chain. This is the cost of detecting, communicating, and defending against new threats. Combining these approaches would enrich and improve the cost-benefit analysis and provide insight to firms into the best way to allocate finite information security resources against ever-changing threats.

To aid this goal, we first present an investment model that can be used for cost-benefit analysis in a broad variety of situations and at any scale. The widely-accepted Gordon-Loeb (GL) model provides a solid foundation and useful framework for analyzing information security investment. We extend it to a broader array of circumstances and risk tolerances, most importantly information with extreme importance and risk characteristics, and where the potential loss would be catastrophic. To do so, we use a model developed as an interpretation of the rate of receiving information. The result, known as the Kelly Criterion, can be interpreted as the optimal proportion to invest in information security. The GL framework can be applied in the Kelly model to determine the optimal investment under changing circumstances. Furthermore, allowing

investments in earlier periods to impact the GL model's security breach function allows investment into more effective information security technology to be represented in the model by increasing the marginal benefit of investment. This improvement provides a better basis for evaluating the cost-benefit of investment into information security. We will show that investment in technology that facilitates coordination for the purpose of increasing system-wide resilience would have the highest cost-benefit.

Second, we consider and incorporate social costs, the benefits of coordination, and the connection between cybersecurity and cyber supply chain risk management. By using two principles from the study of law and economics—the single owner principle and least cost avoider—we describe how our framework incorporates social costs and identifies opportunities for coordination. We then apply these principles more directly to the field of cybersecurity and CSCRM. We describe how investments into CSCRM could provide particularly high returns because of its effect in the combined GL/Kelly model. The effect is to significantly increase the information the firm has available before placing a security "bet," resulting in lower lifetime losses. This improvement identifies an opportunity for high-return investment and a way to evaluate its performance.

Third, we address the heterogeneous preferences of different users, and how to deal with an unmotivated majority. The level of interconnectivity in the ICT infrastructure inherently creates means of compensation. Current practice, especially arrangements such as chargebacks, are especially inefficient forms of achieving the desired level of protection. We claim that the single owner and least cost avoider analyses provide an additional benefit of an alternative means of distributing benefits; namely preference satisfaction. The end result is that the need for direct monetary compensation is largely avoided. However, in some cases we may need the cooperation of an unmotivated majority that may be impractical to directly incentivize. To address these situations, we propose that an alternative is to change the rules of the game and either eliminate the need to coordinate entirely or change which parties need to coordinate. Some commercially successful examples are described, along with the interpretation and implications for cybersecurity.

Finally, we put forward one solution that could achieve the benefits of these principles and models. We show that a utility which facilitates many of the transactions proposed throughout could be an effective solution. There are a variety of reasons these transactions do not happen currently, and we discuss how a cybersecurity utility helps alleviate these obstacles to coordination.

We conclude by commenting on some of the technical capabilities that such a utility may need in order to achieve these benefits. These capabilities are needed in order to address concerns end users may have, as well as to bridge the capability gap between what average end users are capable of doing and what technological tools can do.

2 BACKGROUND

In this section, we briefly describe the relevant concepts from the law and economics literature and highlight the relevant aspects of key economic phenomena.

2.1 Economic Principles

Externalities, or spillover effects, are the costs (or benefits) imposed (or conferred) on others as a result of an individual's actions. These can be negative, such as the pollution from a factory deteriorating air quality for nearby residences, or positive, such as the additional foot traffic to a store as a result of being located next to a popular restaurant. The costs and benefits associated with externalities can be difficult to capture and allocate, particularly over distributed groups. In this way, externalities transfer value from one group to another indirectly. Herley famously explored externalities in cybersecurity concluding that users are rational in rejecting security advice that offers a poor cost-benefit tradeoff [18].

Network effects occur when the value of a network-based product depends on the number and type of users of the network. The telephone is a canonical example: its utility to Alexander Graham Bell was relatively low on the day he invented it, before there were many other people to call. As more people got phones, the number of people reachable by telephone increased, and the value to each existing user increased.

Network externalities, commonly used interchangeably with network effects, are used more precisely in this paper to describe externalities that are conveyed through networks. The value is not derived by the presence or absence of another user or type of user; rather, it is a cost or benefit conveyed indirectly via the interconnections of a network. The 2013 Target data breach is one example. In that case, credit card information, including names and card numbers, for about 40 million customers was compromised, allowing the hackers to generate duplicate cards. The cost to consumers resulting from use of the compromised card data is an externality arising from the fact that Target has access to this information in the course of doing business, but does not bear the full cost when the data is compromised.

Information asymmetry, where one side has more information than the other, commonly exists between two actors. However, when one side has virtually no information on the implications of a set of options, the decisions they make necessarily can not fully reflect the true costs and benefits. Moore pointed out in 2010 that "Ill-informed consumers and businesses are prone to invest in snake-oil solutions if they do not possess an accurate understanding of threats and defenses" [22].

In the cybersecurity context, significant negative network externalities arise from information asymmetry and misaligned incentives. This occurs in situations where one party (or group) creates high costs on another party (or group), often without knowing it. In the extreme, *liability dumping* effectively allows two parties to collude to push risk onto an unwitting third party, or even better, on "no one." Consider bulletproof hosting, where a provider allows customers to host online content in another country where the rule of law is not strong and would otherwise violate laws or terms of service of the customer's home country. Liability is transferred from the customer to the hosting provider, and the customer creates a high cost of takedown on their home country to pursue extraterritorial jurisdiction. The analysis we propose attempts to capture the true social cost of these situations.

The *single-owner principle* is used in the economic analysis of law to assign liability where one party has been harmed by another [15,

29]. The general idea is to consider how someone who owns both the source of a threat as well as the property being threatened would behave. One well-known formulation of the principle begins by asserting that trains run for the public benefit, but only if they pay their expenses. Now suppose one expense of running these trains is that a forest is burned down. If the railway owners also owned the forest, and would choose not to run the train to avoid burning down their forest, then we conclude it is not in the public benefit for them to run the train and burn the forest when they do not own it. If the railway owners would consider themselves net better off running the trains, despite burning down their own forest, then we conclude the benefits outweigh the costs, and the train owners should compensate the true owners of the forest when their trains cause it to burn down.

In many cases, the damage is not so localized as a burned forest with an identifiable owner. The difficulty of identifying the owner can be seen as a *transaction cost* that prevents the transaction from happening [12]. In the pollution example above, the cost of coordination among the residents near the factory is also a transaction cost. Suppose the coordination cost is greater than the benefit they would collectively obtain from the factory not polluting; it then becomes rational *not* to approach the factory owners, despite their imposition of a negative externality, simply because the cost of doing something about it is too high to offset the harm it would prevent. Later in this paper, the solution we propose relies on reducing transaction costs to better allocate the costs of risk creation, thereby increasing the efficiency of investment in information security.

The *least cost avoider* is another concept commonly used in law and economics to analyze the liabilities stemming from negative externalities and determine what decision rule for allocating costs leads to socially optimal behavior. In one canonical example, a patron sues an amusement park for failing to put up a handrail after falling off a set of stairs. The least cost avoider analysis suggests that if the amusement park does not have a railing and a patron is injured, it should be liable because it could have prevented the incident for the lowest cost, by installing a railing. We use the least cost avoider principle to identify potential avenues for cost reduction. Specifically, those with no obligation but who are the least cost avoider could potentially be incentivized with up to the amount of savings, resulting in a net social gain. Rosenzweig previously explored least cost avoider principles with cybersecurity as they applied specifically to software liability [25].

Multi-sided markets are economic platforms with two or more distinct customer groups, where at least one group has a preference about the number of users in each group. The internet is one example, with developers and end users who each want the other group to be large so they have more potential transaction partners. Initially, the infrastructure underlying the internet was designed with developer priorities in mind in order to incentivize development of applications and services for the network, thereby attracting additional users [2]. As the internet expanded and was adopted by more users with less technical expertise, the prioritization of developer convenience over user-friendliness persisted. Accommodating developers is in large part responsible for the incredible variety of functionality that maintains ICT's universal appeal. However, it is also largely responsible for the difficulty in establishing and maintaining effective cybersecurity.

2.2 Economics of Cybersecurity

Economics applies to many aspects of cybersecurity. To better understand the allocation of cost and risk, this section briefly describes the tension between offense and defense.

2.2.1 Interdependence and Co-evolution of Information Security and Cybercrime. Those who study and work in cybersecurity see first-hand the "cat and mouse" game between attackers and defenders. With new technology comes an opportunity and incentive for hackers to exploit it, which begets further defenses and new technology. Because technology offers and protects assets of value, criminals have little incentive to stop trying to break in. Given this persistence, system owners need resilience to recover quickly from attacks.

The interdependence and co-evolution of the markets for cybercrime and information security discussed by Bauer and Van Eeten suggests economic effects can and should be leveraged in a centralized manner to increase the effectiveness of security investment [4]. As they illustrated in Figure 1, there is a web of interconnection and interdependence between criminal activity and cybersecurity. Opportunities to short circuit such connections would increase security efficiency and effectiveness. For the most part, firms have been unable to opt out of having attackers on networks, which could be explained by the fact that offense is seen as the economically favored position [2].

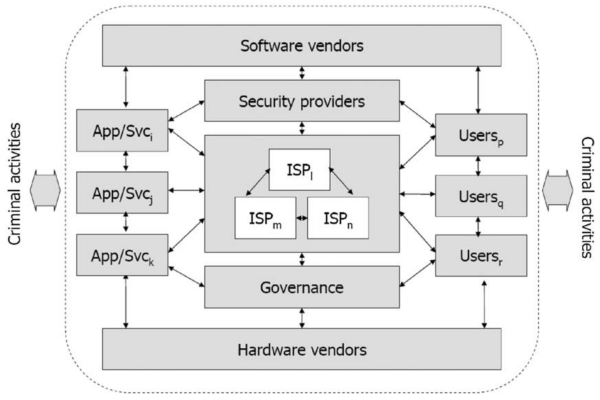

Figure 1: Important interdependencies in the ICT ecosystem [4]

In short, the probability is low that two parties will find the same vulnerability at the same time, assuming that vulnerabilities can arise in any portion of code, that the capabilities of offense and defense are similar [27], and that the vulnerabilities are equally likely to be discovered by any capable party. As a result, the cost of defense with a search and patch strategy far exceeds the cost of offense. Factoring in the low probability that a given user has installed patches for discovered vulnerabilities before an attacker is able to exploit them tips the scales even further in attackers' favor.

2.2.2 Murphy's Law of Cybersecurity. One of the difficulties pointed out in existing work is dealing with Murphy's law: that anything that can go wrong, will. The significance of this for cybersecurity

is that a bug could occur anywhere in all the lines of code and create problems elsewhere. Searching through code has limited value because there is no guarantee defenders will find the same vulnerabilities as attackers [8]. Thus, there is a relative advantage for offense due to the size of the code space and the need to find bugs before adversaries do [2].

Alternatively, if there were a way to submit all suspected or potential threats for evaluation by a utility (as discussed in Section 4), and a fix pushed out to all devices quickly, the benefit to adversaries of finding new threats would be limited to the number of devices that could be compromised before the fix rolls out. Currently, these patches and updates are split over all of the products that need to be patched, making patches a vulnerability in and of themselves. Again, knowing this quickly and telling people *not* to apply a patch is equally valuable.

2.3 The Gordon-Loeb Model

The Gordon-Loeb model (GL) is a widely-recognized framework, first introduced in 2002, for evaluating how much to invest in information security in a one-shot game based on the value of the information at risk [16]. It is concerned with three factors: the expected potential loss, probability of breach, and relationship between investment and probability of breach. In this paper, we focus on the investment and its relationship to the probability of breach, although reducing the potential loss is another lever to consider.

While the GL model offers a maximized one-time investment, it is not a cyber resilience investment strategy. A strategy for resilience should be dynamic and continually updated [10]. These features account for change and recovery over time in the face of ongoing cyber exploitation.

The main objective of our investment model is to expand the applicability of the GL model to continuous time. By doing so, we can evaluate the effects of selective investment and a time-varying breach probability function, giving a more precise cost-benefit analysis. In Section 3.1 we will discuss how to incorporate social value into the analysis.

We extend the GL model in two ways not included in the original model. First, we allow for investment in one period to affect the relationship between investment and the security breach function in later periods. In particular, this allows the effect of improving technology to increase the marginal benefit of investment, as would be the case in practice. Krutilla et al. recently proposed a different dynamic extension of GL, but influenced by the rate at which cybersecurity assets depreciate and the rate of return on investment [21]. Second, we allow for non-risk-neutral parties and catastrophic loss. While these assumptions made sense in the original model, we need an approach that applies to a broader range of risk preferences.

To achieve these goals, we let the GL model be the investment strategy in the Kelly Criterion, presented in the next section. We also need a way to measure value before and after a bet, such as the value of information before and after a breach. Total value can include social community value, which we discuss in Section 3.1, and could also include individual value.

2.4 The Kelly Criterion

The Kelly Criterion, also known as the scientific betting method, is the optimal amount to bet in a game with known odds in order to maximize the long term growth of a gambler's wealth [20]. In cybersecurity, the person making cybersecurity decisions can be considered the gambler who needs to protect her firm's data and other assets. We apply Kelly's interpretation of information rate to extend the Gordon-Loeb Model from fixed time to continuous time.

A cyber defender, like a gambler, benefits from having the best information for making an optimal decision at any instant in time as the situation changes. The complete knowledge of threats and vulnerabilities is constantly changing. The sooner a defender knows about new threat information, the better the probability of defending against new attacks. In the limit, the rate of growth is shown to equal the rate of transmission of information through the channel, and the gambler's best strategy for given odds is to place a bet for the same amount each time, as you would expect for bets with identical odds.

Most importantly, the gambler analyzed in Kelly's paper and the investor in information security in the GL model are making analogous calculations: what proportion of total wealth to bet on a game with known odds is the same as determining how much to spend to secure the value of all your data. Therefore, we can use the Kelly analysis to determine the optimal investment in information security given arbitrary odds (probability of breach). Updating the GL model to account for changes in risk or potential loss is equivalent to the gambler making the optimal bet based on the information provided by the wire.

If one can get the stream of information necessary for Kelly fast enough to enforce the risk probability distribution assumed in the GL model, the result is the optimal rate of investment in security for the firm considered in the GL model.

The Kelly model looks at returns over time, so in order to use it, we need a way to measure value. Any measure can be used, but we propose that the two sources of social value in the next section allow us to evaluate social costs and the benefits of coordination in the context of cyber supply chain risk management.

2.5 Implications of ICT as an Economic Platform

We treat information and communications technology (ICT) broadly as a multi-sided market with two customer groups, end users and cybersecurity experts, who want to collude to keep a third "customer" group, adversaries, out of the market. The network analog of the single owner principle is the vertically integrated, distributed, unitary firm consisting of all non-malicious users. The unitary firm would seek to minimize the harm adversaries can inflict at the lowest possible cost (i.e., the unitary firm is cost-benefit optimizing). In order to minimize the harm sustained, the firm could hire a team of experts to comb through code. However, this would seem to be a wasteful strategy, especially for a single entity. Alternatively, it could engage in a variety of activities to detect threats as they occur and predict threats before they occur, aggregate this information efficiently over all nodes, analyze the results, identify patterns, and redistribute the results to all nodes. To be clear, we adopt a broad definition of "threats," including internal, external, active, passive,

intentional, accidental, and any other category of circumstance. The important point is that an actual or potential expense can be detected and avoided, perhaps at some cost, which may also be determinable. Under certain conditions, this strategy can be more cost effective than current information security investment models. This argument relies on two primary assumptions.

The first assumption is that, all else being equal, non-malicious actors unanimously prefer to have as few malicious actors on their network as possible. We treat losses due to cyberattacks as externalities that reduce net social wealth of non-malicious users. Since not all users suffer losses, this creates an opportunity to identify a least cost avoider (LCA). Risk-neutral users and firms who in fact suffer losses should be willing to pay *ex ante* as much as the amount of the *ex post* loss in order to avoid it [19]. This amount, less the cost of identifying the LCA, is the most the injured party would have paid in advance to avoid it. Since we are evaluating a potential loss, in our analysis we must take the expected value of this amount to get the actual amount a rational actor should be willing to pay to avoid a loss. Thus, minimizing the cost of identifying the LCA is a critical determinant of how few malicious actors it is economically justified to leave in the system, and therefore how close we can get to the single-owner ideal.

The second assumption is that in the short term, when faced with obstacles, adversaries will switch to alternative means of accomplishing their objectives because there are sufficient substitutes. For example, they may increase social engineering attacks in response to increased defenses such as stronger encryption or password policies.

These assumptions are consistent with findings that adversaries can change tactics quickly and for low cost at any time, and do so in response to defensive measures [24], which support findings that technological means alone are insufficient [7]. This makes the protection offered by any particular measure perishable, because adversaries are already searching for a workarounds to existing measures while the continuous innovation of new software, hardware, etc. create new spaces to find vulnerabilities. Therefore, solutions that facilitate a dynamic set of continuously-improving processes are most desirable [23].

3 GROUP DYNAMICS AND OPTIMAL RATE OF INVESTMENT

Unlike most other approaches for calculating cybersecurity investment, we propose that social value is a keystone component. In this section, we describe how group dynamics and coordination can be successfully incorporated to recalibrate incentives for cybersecurity.

3.1 Defining Social Value

In economics, social value and social cost are defined as the total value or cost to a society or community of a given transaction. It includes both private costs (direct costs to the producer for producing) plus externalities (imposed on a third party who did not agree to incur that cost). The implication of these externalities is that solutions which are optimal for transacting parties (on the basis of only private costs) are not necessarily optimal for the ecosystem as a whole. In the context of cybersecurity, social costs and benefits of individual information security decisions are critically important but difficult to measure and attribute. This creates two analytical difficulties. The first is that, in practice, the externality value of decisions is ignored frequently enough that the predictions of models accounting for social value may only hold where users *are* considering social value. The second is that much of the literature also ignores social costs, both because they are difficult to measure and because users frequently ignore them.

However, considering social value and cost can unlock tremendous value for any ecosystem. In his 1960 paper "The Problem of Social Cost," Robert Coase establishes that, in the absence of transaction costs, affected parties can coordinate and use free-market mechanisms to ensure that all resources go to their highest and best uses [12]. *Coasean Transaction Costs* are now understood to refer not to financial fees but to the intangible costs that have a tendency to prevent otherwise socially value-creating transactions from taking place. In Coase's hypothetical, harmed third parties transact either directly or through an intermediary to ensure that the socially optimal outcome is achieved. We posit that a more deliberate consideration of social value and Coase theorem in information security can uncover novel opportunities to improve overall security infrastructure. For our purposes, we consider only network externalities that pertain to system security, and are primarily concerned with the discovery and capture of positive network externalities rather than negative consequences. The rationale for this decision is rooted in the least cost avoider principle, which assigns liability to the party with the least cost to mitigate the externality. Combining this assignment of liability with the market approach of Coase's theorem results in a structure wherein the party best situated to act is compensated some amount by a third party who would otherwise not achieve the externality benefit. For example, the vast majority of non-malicious users lack the resources, knowledge, or incentive to actively participate in the removal of malicious actors in their network. When malicious actors are removed from a network all non-malicious users receive the benefit of increased security. It stands to reason that all non-malicious users would be better off if they could collectively compensate a third-party actor to take initiative and remove those malicious actors. This simple principal forms the basis of the Cybersecurity Utility Platform outlined in Section 4.

3.2 Social Value and Transaction Cost Optimization Provide a Basis for Incentive Recalibration

The real world is not as simple as the one constructed by Coase. Namely, the assumption of zero transaction costs does not hold in practice. It does provide a structural framework from which one can analyze the relationship between social value and transaction costs. In fact, Coase stated that the intention behind the Coase Theorem was "to make clear the role which transaction costs do, and should, play in the fashioning of the institutions which make up the economic system" [12].

Specifically, Coase was concerned with "misallocations," or situations improvable by bargaining. Social value is not optimized for situations that a single owner would not accept, and should be no more accepted by society at large. When social value is not

optimized, this is a misallocation. Coase concluded that in an environment with rational actors, costless bargaining, and no legal barriers to bargaining, the market would correct misallocations through transactions until the optimal result is attained [9]. We are most concerned with the rationality and costless bargaining prerequisites of Coase's conclusion.

In the current system, social value is not captured effectively. Therefore, we can conclude that current transaction costs between all of the involved parties must be prohibitively expensive. It takes time and resources to coordinate between multiple parties, especially vast number of network end users that would need to be engaged. Additionally, there is an information barrier preventing effective communication and transaction. Cybersecurity benefits are somewhat unique in that they are probability-based. Security does not prevent a fixed number of incidents, it lowers the overall probability of a possible incident. This level of separation between actions and consequences means that many users are underinformed or uninterested in taking an active role in their own security, so they have self-selected out of the coordination process. The current transaction costs of cybersecurity collaboration at the necessary scale are greater than the expected social value gained through collaboration. To enable coordination, and ultimately achieve the socially optimal behaviors, we must find a system in which social value is greater than the sum total of all transaction costs. Fortunately, transaction costs are not inherent characteristics of any market. Transaction costs can be reduced by effective institutional structure, technological change, or simply new ways of working. In the following sections we explore potential changes to existing cybersecurity sector structure and their effectiveness at reducing transaction costs.

3.2.1 Barriers to group decision making. Coasean transaction costs are different from financial transaction costs but nonetheless can prevent otherwise desirable transactions from occurring. Coordination of a large and potentially ill-defined group and excluding free riders are canonical examples, each reducing the potential value of organizing a transaction. Additionally, even if coordination is possible, there is the risk of holdouts, or group members who threaten to withhold their assent to the deal and prevent the rest of the group from benefiting unless they receive a greater amount of compensation. One commonality among these examples of Coasean transaction costs is that they envision a hypothetical negotiation between parties, individual or collective. We loosen this requirement to better correspond to the situation we want to model, and discuss the consequences in Section 3.3. In short, our assumption of a single, unanimous, consistent objective obviates the need for individual negotiations and makes distributed decision-making possible.

Collective decision-making among a large group presents difficulties beyond Coasean transaction costs. There are three traditional avenues for achieving goals as a group: rules or laws, incentives, and voluntary actions. These are not mutually exclusive, and combining the effects of each can be advantageous, as we will discuss in Section 3.2.2 below. However, many of the aspects of the ICT ecosystem we have already discussed make each individually inadequate to improve security. Our proposal will therefore involve leveraging the positive aspects of each to minimize the weaknesses in others.

In particular, the set of motivations and objectives among all ICT users is incredibly diverse. Bindewald aptly notes that "[i]f one attempts to motivate a large group to achieve a certain goal, this diversity of priorities among group members poses the non-trivial problem of how to convince a sufficiently large fraction of group members to contribute to the goal[,]" and that the "most promising strategy depends on the problem at hand," especially "the certainty and ability to implement prescriptive or incentive actions" [5].

With respect to the ICT ecosystem, the ability to implement prescriptive and incentive actions is limited. Thus, we propose that voluntary action not only *can* be used, but must be. Bindewald and Atallah provide a game-theoretical analysis of multiple-goal achievement in groups through voluntary efforts under a variety of distributions of individual motivations within the group [6]. Their findings suggest that in large groups, a diverse set of objectives actually increased the robustness of the outcome to "black sheep," or agitators with the intention of preventing the beneficial outcome. The results are particularly applicable because they pertain to situations "where only a minority is motivated to achieve certain goals and enforcing mechanisms like coercion are not reliably available," which is consistent with the problem of information security [6].

3.2.2 Dealing with an uncooperative majority. In their daily lives, individuals have to choose how they spend their time and attention. For many, cybersecurity ranks relatively low on the list of priorities [28]. As a result, a system of network security that relies on significant individual effort to be successful is strategically flawed. However, as long as the goal is desirable to all involved, a sufficiently capable and coordinated minority can successfully bring about the outcome for the entire group. It should not matter if a large proportion of the population has zero willingness to pay or contribute to accomplishing the goal.

In the context of cybersecurity, it is especially unlikely that a "majority strategy" will be effective. The technical complexity of many of the vulnerabilities is beyond most peoples' ability to anticipate. This is compounded by the persistence of attackers, making it nearly impossible to stay fully informed of new vulnerabilities as they arise. Assuming that non-malicious users unanimously and rationally prefer fewer attackers, this complexity and the individual burden of staying up to date represent significant obstacles to effective participation in addition to the barriers to group decision making just discussed. Most significantly, these burdens are great enough that the average user is *rationally inattentive*. As a consequence, they cannot behave as the rational actors Coase envisioned because they are not sufficiently informed of all relevant details that help to identify misallocations. Therefore, rationally inattentive users in general cannot engage in misallocation-correcting transactions. The "majority strategy" is unlikely to be successful because most users are rationally inattentive with respect to cybersecurity and are therefore unlikely to independently make misallocation-reducing decisions.

Instead, the fact that reducing the number of attackers and losses to cybercrime is universally acceptable among all non-malicious users can be leveraged by making implementation effortless for most people [14]. By doing so, the uncooperative majority are not made to decide whether to allocate some of their scarce resources to system security, and can instead passively contribute. At the

same time, the reduced number of decision-makers increases the likelihood of success of the minority preference.

Recalling the three traditional means of achieving goals as a group discussed in Section 3.2.1, incentives and voluntary action are most relevant here. One part of the difficulty with cybersecurity is that incentives between any two individuals or the groups individuals regularly transact directly with are not strong enough to assure optimal precautions in general. The other part of the problem is that the transaction costs to correcting the imperfect starting incentives are prohibitive. However, the resulting social loss described above provides an ample basis against which to offset the expense of relieving transaction costs. In other words, since the losses from ineffective cybersecurity are so large, the potential value of reducing transaction costs are correspondingly large. This is so because, as discussed above, those who suffer avoidable losses should be willing to pay some amount to avoid the loss. We expect this willingness to pay to be function of the size of loss, probability of loss occurring, and an individual risk-aversion factor. In practice, transaction costs should also be included in this calculation, so the value of reducing transaction costs is therefore not the reduction in cost itself, but the increase in social value generated by the incremental transaction. Therefore, the voluntary action taken by the motivated minority need not be charity but could instead be strongly motivated by incentives.

There are numerous examples of the efforts of a motivated minority effectively reducing transaction costs and increasing coordination as a result. Reducing transaction costs to allow a transaction to occur is analogous to establishing a price signal and creating a marketplace. For example, high frequency trading firms use fiber optic cable and running it as close to the central exchange as possible. This allows them to place and cancel a large number of transactions in rapid succession, probing the market for the bids of other market participants. The difference in speed is meaningful enough for the trading firm to adjust its order before the slower bids close. For stocks that have enough trading volume, there can be enough signal to meaningfully have an advantage from this signal. A second example is the Robinhood investment and trading app. The app allowed retail investors to buy and sell stocks with no trading fees, and the company sold retail traders' aggregate trading data to hedge funds, providing them a better pricing signal. Another example is Google. Before Google, advertising firms' competitive advantage was in knowing the market with better consumer testing data than their peers. In other words, firms compete on their ability to gather market data. Once Google became the dominant search engine, it had access to the freshest market data around: individual real time search data. Consumers were not consulted because there was no need. They wanted search results and Google provided them. This data allowed vendors to no longer need expensive to produce market research and could instead rely on the greater price discrimination offered by Google's more targeted advertising to make up for the loss of proprietary (and stale) information.

Bindewald notes that where a minority objective depends on the majority for success, "one strategy to cope with the difficulty of coordination games may be to "change the game" such that the need for coordination is minimized" [6]. Each of these examples changes the game in a significant way and supports our claim that this extends to large groups as well as the small groups in the particular study.

3.3 Value of Indirect Coordination Solutions

The opportunity to increase social wealth presented by reducing transaction costs is particularly promising because, as noted above, the existing ICT infrastructure is effectively a multi-sided platform and much of the value derivable from ICT comes from network effects.

There are two significant consequences from these facts. First, the platform orientation and high level of interconnection makes it virtually trivial to establish a connection between any two nodes, so the cost of making contact is negligible. Second, since the value of the network changes based on how well users' preferences are being met, both direct compensation as well as indirect compensation through preference satisfaction are available options for incentivizing the LCA to take optimal precautions. At the same time, these preferences could be (and often are) with respect to other groups of users, so the value delivered doesn't have to be perceived by the recipient, but could be anyone. This means firms that place a greater value on security can pay for others to be better protected without their knowledge. Taken together, this means the cost to establish the connection necessary for a transaction is low and value can be delivered in a variety of ways, most importantly indirectly by preference satisfaction. Thus, two of the three Coasean transaction costs can largely be avoided: coordination of large groups and the threat of holdouts that arises in direct negotiations. The third Coasean transaction cost, excluding free riders, can be ignored here because, by assumption, the objective is to exclude malicious users from the public network. There is no free-rider because a non-paying user is just another opportunity for the motivated minority to meet some third party's preference for the non-paying user's security.

Although in some cases it suffices to create a pricing signal and establish a marketplace to reduce transaction costs, this is not so in cybersecurity. The highly inconsistent ability to determine the value of security measures among the general population prevents the pricing signal and marketplace from being meaningful to the average user, which is precisely the current situation. Relying on indirect transactions creates the circumstances necessary for a motivated minority to successfully accomplish an objective that depends on the majority's participation but that is only strongly held by the minority. Namely, the goal of minimizing the number of malicious users is acceptable to all non-malicious users, and indirect value creation through preference satisfaction makes it possible for the majority of users to contribute nothing personally.

Indirect coordination offers another advantage as well: information on threats and vulnerabilities is generated and can be aggregated by the motivated minority. The mechanics will be discussed in Section 4, but for now we take this as given and explore the potential uses of threat data. Information on the vulnerabilities faced by actual end users under real-life conditions is currently incomplete at best, and real-time insight into current vulnerabilities could be useful in a couple of ways.

First, the "Murphy's Law" problem and the inherent shortcomings of software testing create a dynamic where offense can be favored if finding vulnerabilities is purely probabilistic, even with

several orders of magnitude fewer resources [8]. By collecting threat and vulnerability data from actual users, the motivated minority effectively make real-world use the testing regime. Not only are these the most relevant conditions for testing software, but a dataset of all known vulnerabilities is more useful the more rapidly it includes new threats.

Second, this information could be used both proactively by developers to anticipate and avoid issues, as well as reactively to enlist these developers for help with resolving vulnerabilities discovered in their programs. As a result, developers could have more thorough testing protocols before release as well as more visibility into cross-compatibility issues after release.

4 A CYBERSECURITY UTILITY PLATFORM

Our single-owner social value analytical treatment has four implications for effective cybersecurity. The first is that the marginal costs and benefits of security investment have both private (subjective) and public (objective) components. The second implication is that the incentive problems widely recognized throughout the literature as an obstacle to information security can be quantified against a theoretical optimal outcome, the single-owner outcome [4]. The third implication is that since the objective is to maximize collective social value with respect to losses from cybersecurity, direct transactions between affected parties are unnecessary and we can instead adopt a "net costs in, net benefits delivered" black box approach. Lastly, the connection between the single owner principal analysis in theory and practical implementation lies in synthetically recreating the single owner outcome for independent firms. This is the job of Bindewald's "motivated minority."

In the cybersecurity context, this means that rather than each firm searching for the least cost avoiders, the motivated minority's job is to make the necessary matches happen. As we will discuss below, the necessary transactions consist primarily of communication and coordination of efforts to discover and broadcast vulnerabilities as soon as they are discovered. Centralizing costs in the motivated minority and those who most value security is the efficient allocation, which provides better incentives to the other players in the system.

We now propose a novel Cybersecurity Utility Platform that serves this function, describe its implementation, the effects on stakeholder incentives, and outcomes. The Utility to which we refer throughout this section need not necessarily be a government-run or regulated entity, nor even a single entity. It consists of, at minimum, a communication channel or channels that enable the identification of threats and vulnerabilities precisely where users are finding them "in the wild," each of which can then be allocated to the most efficient party to resolve. The Utility's function is primarily communication and coordination, and therefore maintaining the channels and infrastructure necessary for maximum availability is critical. In practice, this is just like many other communication networks, such as the internet, cell phone service, or cable TV, where the communication channel and the content that draws users in may be provided by different firms. Here, the Utility serves a similar market-making function by matching valuable information with those who are most able to act on it, letting everyone get the benefits of the actions of a few.

4.1 Cyber Supply Chain Risk Management

The Utility's role as the "motivated minority" is effectively equivalent to implementing cyber supply chain risk management (CSCRM) processes over all ICT users.

The core concept of CSCRM is not novel, and the need to manage cyber risk within supply chains is well recognized [7, 23, 26]. However, this is currently regarded as an organization-specific endeavor [17]. We instead refer to the collective cybersecurity concerns shared across the supply chain in CSCRM. At the organizational level, CSCRM requires management of both physical and virtual processes to be effective. At the collective group level, we focus our attention solely on virtual processes, leaving the physical processes within the individual organization.

An effective Cybersecurity Utility Platform must accomplish a diverse set of goals simultaneously. First, it must be consistent with the core principles discussed in Section 2 of this paper. Second, it must effectively reduce the transaction costs in the industry significantly enough to unlock the maximum social value. Third, it must be self-sustaining and able to operate in the absence of any regulatory intervention. Creation of a centralized platform which operates collaboratively with, but fully independently of, any individual entity offers the greatest potential to satisfy these goal.

In the fall of 2020, we conducted semi-structured exploratory interviews with 40 stakeholders across the Defense Industrial Base (DIB) to gain insights about their needs in cybersecurity. Over the course of these interviews, we found that many average users do not find cybersecurity to be a serious threat and are unlikely to take it upon themselves to learn more about their posture or how to improve it. Similarly, many companies view security expenditures as purely a loss item because they do not perceive any return on the investment. Interestingly, a common complaint among security experts was that when they arrived on scene, they have little to work with and their ability to diagnose and/or revert back to an earlier state is limited. This could have a subtly negatively reinforcing effect on average users taking proactive measures because they experience the loss coupled with the expert's inability to do anything. On the basis of information collected in these interviews we have concluded that the platform will need to address four specific points of failure that arise from trying to manage CSCRM within individual organizations. Namely, it must address scanning, information, learning, and incentive failures. Effectively resolving these four issues will have the required impact on transaction costs. We also posit that resolving incentive failures will ensure that the platform is truly self-sustaining, thus fulfilling our primary goals for the platform.

Scanning failures occur when a threat or event is not detected due to lack of attention or insufficient resources. Cybersecurity threat intelligence can provide information necessary to reduce the probability of breach for participants throughout the system. More rapid detection and awareness of the existence of vulnerabilities, therefore, could create value by preventing the loss that would otherwise follow from the avoided breaches. This value is measurable by the number of devices that face the reported threat and are not compromised as a result of already having implemented protective measures. When we asked cybersecurity experts what would be

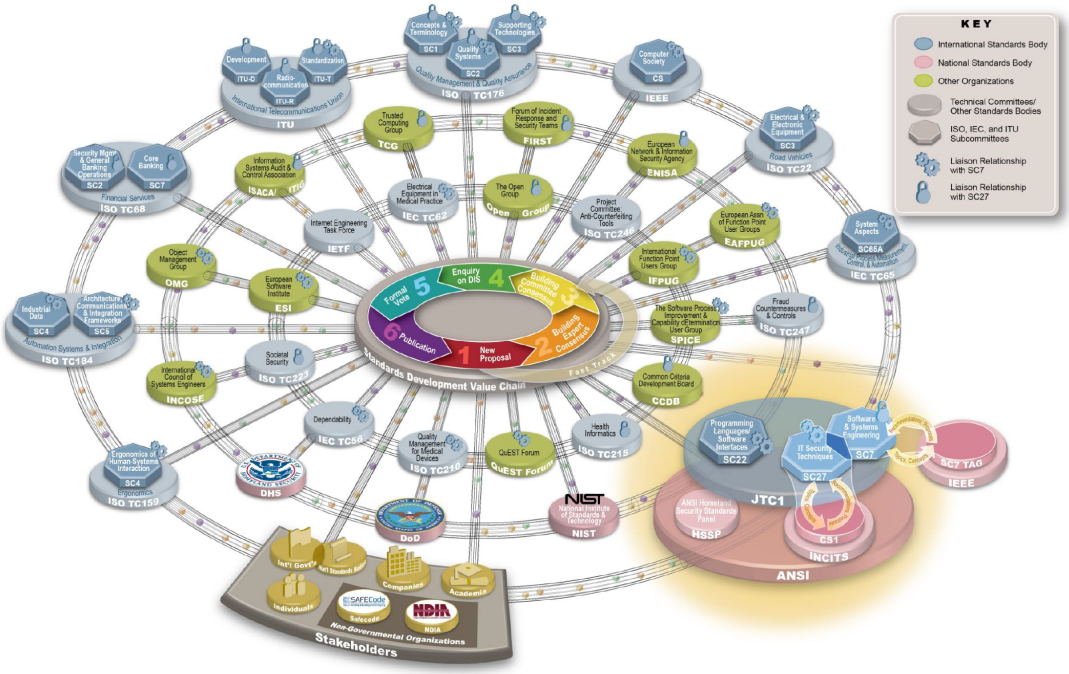

Figure 2: Cyber supply chain standards landscape in 2009 [3]

most helpful to better protect the internet, many reported a lack of visibility into the threats faced by users at endpoints. This includes critical data such as real-time information on adversary attacks, specific methods used, and any early indicators of trouble, as well as less urgent information such as potential system vulnerabilities, bug reports, and endpoint status. On the other hand, people working at the firms we interviewed reported that cybersecurity was thought of as a necessary nuisance, reflecting a low willingness to spend additional time learning more about effective security. For smaller firms, lack of resources is the biggest contributor. Whether time, money, or expertise, many small firms lack the resources to detect threats. We found these firms often took no steps to attempt to detect threats. Larger firms tend to devote more resources to information security, up to and including a dedicated internal security operations center (SOC). However, management responsible for security reported difficulties in convincing the rest of the executive management that additional investment in security is worthwhile, in large part because it is not viewed as a revenue generator. This is offset to some extent by risk management and legal departments increasingly advocating for better protection, but the decision is still predominantly viewed with a return on investment (ROI) mindset. Given the mismatch between what experts need and the willingness of most users to provide it, it is clear that cybersecurity must be made as simple as possible to generate and convey this information.

Information failures occur when the information generated from scanning activities is not effectively aggregated, analyzed, and distributed to those who need it within a system. We found this to be both a top-down and bottom-up issue in our interviews. In particular, the small and midsized businesses (SMBs) failure to take action to detect threats could be attributed to lack of resources, but at the same time, we found that these firms were also largely unaware of being potential targets for foreign adversaries and a few affirmatively (though mistakenly) stated that they were not targets. This reflects in part a failure of the larger and more capable firms getting information on the relevant indicators to these smaller firms. On the other hand, Security Analysts at both private sector firms and government agencies report a widespread lack of visibility into the threats users face at endpoints, limiting their ability to proactively protect users by providing information on indicators of attack or compromise. In other words, the average end user is rationally inattentive as to their own cybersecurity posture and the impact it may have on the system broadly. Even those who have been hacked before do not spend much time or effort researching how to improve their protection. This behavior can be explained in part by the lack of interpret-able information on what various measures or services actually do. Several firms reported an inability to determine quality of protection when selecting between alternatives, and ultimately deciding based on price.

Learning failures occur when lessons from earlier events are not sufficiently documented, synthesized, and distributed to the relevant people. In the context of cybersecurity, this includes threat reporting and successful defensive measures. In an ideal state, data from end users could be informative for the entire market, including identifying risks as they develop and detecting potential compromises earlier. Instead, threat information generally becomes available only, if at all, long after the attack has happened and then has been selectively filtered to protect confidential information. Creation of an anonymous pool of end user threat data would enable

security analysts to learn from the true population of incidents rather than an extremely small sample.

Incentive problems arise when an organization fails to act on available information due to a lack of incentive. The scanning, information, and learning failures described above each contribute to, and are exacerbated by, the current set of incentive structures, leading to worsening circumstances and increasing vulnerability across the system. As described above, many users do not care about cybersecurity at all and do not consider it a relevant concern, while a small number of users place a high value on cybersecurity. When the costs of cybersecurity fall disproportionately on the group that puts a low value on cybersecurity, the measures taken will be primarily motivated by minimizing this cost rather than achieving some specified level of security. Costs are not allocated to the optimal parties and we currently have no effective mechanism to facilitate side payments that would enable the Coase Theorem to take effect. Without an efficient means of incentivizing ambivalent users to actively contribute to overall security (i.e., coordinate as part of the whole), coordination costs will be prohibitive and increase rapidly with the complexity of the organization, and overall protection will be sub-optimal.

4.2 Building a Cybersecurity Utility Platform

We propose the creation of an independent service provider that will fill the role of the motivated minority in the cybersecurity market. This independent service provider will bear the responsibility of coordination across all actors in the market and actively manage the overall security of the network. By becoming both a market intermediary and new least cost avoider, this entity can execute socially optimal behaviors without incurring the high transaction costs that hinder current players. The majority of players will no longer need to put in any effort beyond what is optimal for their personal interests and yet the socially optimal outcome will still be reached. The service provider will take on the costs of coordination with end users as a cost of doing business and, in exchange, could turn around and sell improved overall network security to the small group of users with a high value of cybersecurity. Essentially, this will be using the same ICT infrastructure connections that contribute to system vulnerability and disaggregated risk creation in the opposite direction to collectively reduce the impact of vulnerabilities.

Extending Bindewald and Atallah [6], we propose that the general ambivalence toward social cost when it comes to information security investment coupled with an (assumed) near-universal preference for fewer attackers is equivalent to the scenario in which the group generally has consensus but is asymmetrically motivated to invest, whether time or effort, in achieving the desired outcome. Thus, the Unique Individual Motivation game criteria are met. The Cybersecurity Utility Platform creates a multi-sided market which enables actors to actually make asymmetric levels of investment that match their own best interests. Simply by conducting security operations through this entity, we are enabling the side-payment system described in the Coase Theorem. Actors with a high value of security pay according to their subjective valuation, creating a cross-subsidy and ensuring the cooperation of actors with a low value of security.

At a high level, the Cybersecurity Utility Platform will perform three distinct functions to fulfill it's role as a market intermediary.

The first function is recalibrating incentives in the marketplace. For the majority of users who place a low value on security, it will offer free or deeply discounted security services. These services will come in the form of user security education and basic threat monitoring. In exchange, these users will provide the Cybersecurity Utility Platform with broad access to their threat data. Any potential vulnerabilities or weaknesses in a given piece of software will be reported back to the central platform. This data can then be aggregated and used by security experts (within the platform or externally) to resolve the discovered vulnerabilities quickly and efficiently.

The second function is enabling cyber resilience across the whole network by coordinating the efforts of experts and other actors in the ICT ecosystem. For the niche set of users who place a high value on security, the Cybersecurity Utility Platform will provide assurance that the entirety of their supply chain is secure by ensuring that all individuals and entities are capable of defending themselves, threats are resolved as soon as discovered, and that evolving threats are identified as quickly as possible. In exchange, these high-value users will directly compensate the Cybersecurity Utility Platform and serve as their primary source of revenue.

Lastly, the platform will also have the ability to advance overall industry knowledge. By aggregating all available end user threat data in a central location, it can facilitate proper documentation, synthesis, and dissemination of key knowledge. In an ideal scenario, this database could become the basis for an industry-wide knowledge platform that could further subsidize the cost of operating the overall platform. More importantly, by serving as a source of customers for security products, developers have a strong incentive to interface with the Utility, allowing for greater coordination among independent organizations than would otherwise occur without some formal joint venture or co-investment agreement.

4.3 Implementing the Cybersecurity Utility Platform

Boyson once said: "In highly volatile operating environments, both in IT systems and in supply chains, the very structure of the organization and how it is configured determines adaptability and performance, with higher degrees of integration leading to better enterprise performance" [7]. Security in general is a highly volatile operating environment, which explains the growing emphasis on cyber resilience as an objective. Therefore, the ability to adapt to changing circumstances and the overall level of network security depends on how well the various tools available to establish security are integrated. This insight applies whether at the individual or systemic level.

The potential value we envision is measured as a function of (network aggregate) preference satisfaction, effectively allowing for indirect rather than direct compensation. The structure of the solution itself is therefore at least in part responsible for recalibrating incentives, alleviating bottlenecks, and generating presently unavailable information. By minimizing the role of the least capable, incentivized, and avoidable ICT users in maintaining system security, the remaining participants in the marketplace are informed,

capable, can be properly incentivized, and are best situated to act. By providing a clear pricing signal on a dynamic basis, these market participants can make optimal decisions, leading to tighter integration. As discussed above in Section 3.3, indirect value distribution avoids many of the potential difficulties involved with direct compensation, most significantly the Coasean transaction costs that currently prevent the free flow of system threat information and effective coordination.

A Cybersecurity Utility would share several characteristics with other types of utilities, such as a high level of upfront investment and a large number of potential users to secure economies of scale. However, unlike traditional utilities, in this case there are many alternative options, including the choice to do nothing. As a result, in order to be successful this Utility will have to successfully compete with these alternatives. To do this, the Utility can focus on maximizing the economic value perceived by end users, which, in the majority of cases, will be average users who are provided these benefits in exchange for information on threats faced or bugs.

Due to the large reliance on indirect network effects to create and distribute value, it will be critical to quickly acquire both the large companies on which it relies for revenue and the smaller companies on which it relies for data. This is because of the need to get buy in from investors and decision-makers; both of whom will demand a high probability of success. In addition, the expectations of end users are also vitally important. To achieve the scale at which efficiency begins to rapidly increase, the Utility will need to reach a tipping point where it has enough users that prospective users begin to increase their estimates of the probability of the Utility remaining in existence several years later, increasing the expected value of the Utility [13]. One explanation for this behavior is that since the value is derived from indirect network effects, early adopters will get less value at adoption than later adopters. Therefore, consumers expect greater value from products and services they believe will be around long enough to have later adopters. The expectation of greater adoption in the future increases the perceived value now, affecting how much initial adoption occurs. This applies equally to decision-making by both SMBs and large enterprise users, although the two groups likely differ in their ability to accurately estimate probability of adoption. This is in part because some large enterprises may be able to endogenously affect the probability of adoption whereas it is unlikely any individual small and midsize firm can do the same. Investors in a potential Utility will likewise want to see either or both parties placing a high value on the Utility, potentially shown by promising numbers of early users, which could reflect high expectations of its success by the users themselves.

Another feature distinguishing a Cybersecurity Utility from a traditional utility is that there are not just alternatives but vigorous competition and a relatively high level of information asymmetry. For most traditional utility industries, the quality of the service received is readily observable or at least determinable by the end user. People know when the lights are out; water quality can be tested. But the average internet user can not tell how secure they are, which makes it difficult to judge the value of a particular security product as a function of vulnerability before and after investment. This parallels the well-known Market for Lemons problem [1], and in the case of cybersecurity the analysis suggests that there is a thick market populated with products and services of variable quality in part because consumers have a broad spectrum of capability of evaluating quality or cost/benefit.

One possibility for successfully competing in the crowded and fragmented security market is to integrate the functionality of various specialists, compounding economies of scale and scope to boost effectiveness by complementing rather than competing. Incidentally, this is precisely what has developed naturally among adversaries as a result of the co-evolution discussed in Section 2.2.1. Offensive capability is virtually costless to reproduce once in the wild, and adversaries benefit from and build off of the compromises achieved by each specialist. Similarly, defensive capabilities and processes can be performed relatively costlessly, but many industry standard pricing models rely on the premise of individual investment, which leads the optimal strategy for the security specialists to charge in units that correspond to the end user's perception of value. To these users, the perception of a scan that shows no vulnerabilities is an expense for nothing in return. The perception is likely especially negative when these scans are performed as part of a compliance regime because the end user presumably does not see enough value in the scan to perform it absent a regulatory requirement. These are precisely the wrong incentives for system security, and is an example of the types of transaction costs that arise under a direct compensation scheme. A Utility that integrates the functionality of various tools on the market could presumably use the combined signals to greater effect than each individually. To see this, suppose a new vulnerability has arisen where an additional piece of hardware is present on a particular network. It may be difficult to detect this piece of hardware via software means, but sensors that detect additional current could reveal its presence. Many electronic devices have signatures that make them identifiable with this type of analysis, which could supplement software detection as described. Extending this logic, a Utility would sit in a position to achieve economies of scope by integrating the signals from various tools. With the additional economies of scale enjoyed by Utilities by assumption, the potential value deliverable by a Utility is greater than the individual capability providers.

This does not mean that a Utility would replace individual vendors, however. On the contrary, the innovation fostered by competition is critical to keeping up with the development of new offensive techniques in addition to unintended but unavoidable bugs. The Utility is responsible for making sure the communications channels are working properly, and coordinating with other stakeholders when problems are identified.

The actual implementation of the Utility Platform can be broken down into the product, services, and business activities performed.

Starting with the product, the average user should only have to make as few security-related decisions as is practicable, but should neither be expected to learn much beyond the intuitive nor entirely relieved of considering security. Taken together, a simple app with basic cyber hygiene tools is sufficient for the base level of the Utility's products. The user interface will be important to meeting the constraints simultaneously and should combine human-centered design with information design principles in an iterative design process with input from end users.

The Utility's services consist largely of ensuring proper information flow for maintaining a secure posture. Assuming the "single

owner" outcome is the objective, each of the four supply chain information failures described above must be addressed.

First, information needs to be generated about vulnerabilities of any type. In order to prevent security breaches, they must be known to begin with. In practice, each device should be able to assess its own status, or at least connects to a device that can assess its status, against all known vulnerabilities. This corresponds to the scanning failure described above.

Second, the vulnerability information needs to be aggregated in a way that meets the economic requirements of the universe of non-malicious users. In other words, in a way that is generally not objectionable to most if not all users. This information would consist of anonymized reports of threats experienced and vulnerabilities discovered so that Indicators of Compromise or Attack can be more quickly discovered, and preventative action can be taken more quickly. A private deployment of the system we propose can be used by a firm with heightened disclosure enabled in order to increase visibility into its own processes and risk exposure, as well as manage privately negotiated cybersecurity requirements. This mitigates the Information Failure described above.

Third, the aggregated information will be analyzed and distributed among select trusted capability providers. By sharing this information with partners, each has the best possible basis on which to make design decisions and improve their own products. At the same time, the individual providers do not need to worry about revealing their intellectual property through joint ventures or maintaining relationships with hundreds of other companies. In addition, much of the information shared pertains to developing threats and present status rather than output that could be reverse-engineered. This mitigates the incentive failure described in Section 4.1.

Lastly, redistributing the most up-to-date remediation and defensive configuration information to all endpoints completes the cycle, resolving the Learning Failure described above. This connection is less obvious because of the indirect value creation employed, but the outcome is the same as if efficient bargaining on the market were possible. In other words, the Utility synthetically recreates the same outcome as the market would if there were no transaction costs. However, in this paper we have described how the Coasean Transaction Costs involved in direct bargaining are substantial and in fact prevent information transactions and described how indirect, distributed compensation synthetically recreates the outcome of efficient direct bargaining through the work of a motivated minority. In practice, it is not important who paid whom in any particular transaction because the socially optimal outcome maximizes the sum of individual outcomes, not any individual outcome. Since we are only considering security-based misallocations and transactions, the Utility effectively removes security issues from the negotiating table of most deals because it is handled separately. By generating information from individual endpoints, aggregating over the network, analyzing to identify new vulnerabilities, and distributing the most complete and current set of instructions for protection, the Learning Failure described above is mitigated.

Having implemented a process to resolve the Scanning, Information, Incentive, and Learning failures, the Utility is then in a position to offer valuable services. For instance, the app given to individuals can be backfilled with whatever defensive capability is needed without requiring user interaction. Firms with higher security requirements likely have internal security operations, but visibility does not extend far beyond the corporate border. This can be problematic, especially for firms with regulatory obligations to enforce compliance throughout their supply networks. Suppliers are unlikely to reveal their own suppliers for fear of getting cut out, among other concerns, so even the identities of these firms may not be known, let alone their security configurations. Typically, the obligation to enforce compliance is required to be included as a condition to each subsequent subcontractor, but this solution is imperfect. The Utility could, however, include in its app instructions for scanning a local network for compliance with some set of rules, and anonymously post the results to a distributed ledger that the Utility maintains. A shared key is generated for each transaction, allowing only the two participants in the transaction to see the results of the compliance scan and who the results correspond to. When any node in the network goes out of compliance, the rest of the nodes are aware but can not see who. Either the Utility or the firms themselves could then check to see if they are the cause and remediate as necessary. For the security-sensitive firms who have ultimate responsibility and liability for compliance throughout, the ability to get reliable assurance of compliance, even without identifiability, is a big improvement over contractual requirements. This service would be particularly valuable because not only are successive subcontractors presumably less capable of monitoring the next, but the firms are also more likely to be judgment-proof further down the chain.

Lastly, the Utility's business activities revolve around coordinating amongst the motivated minority and interfacing with the passive majority. Coordination should include integrating vendors' functionality into the user-facing app, optimizing the user experience to maximize usability for the user, and contacting the best-situated to respond to emerging threats.

4.4 Omnichannel Cybersecurity

To illustrate the potential for indirect coordination solution in the information security space, we propose that the information that could be generated by a Utility would allow methods adapted from omnichannel marketing to help accommodate passive participants. Omnichannel marketing is a strategy where firms optimize their use of various content channels to deliver messages to, and learn about the preferences of, potential customers. For example, when a user searches a product and subsequently sees ads for that product on other websites, that was likely omnichannel marketing. The vendor maximizes the amount of profit from their customer base by leveraging all the information available to them about their customers to determine the highest-value offer to make at any given time.

Omnichannel cybersecurity applies the same strategy over all ICT channels as a platform in order to facilitate collusion between end users and security experts and to impose as high of costs as possible on adversaries to participate in the market. In practice, the "collusion" takes the form of better coordination between the other two groups and the "cost" imposed would take the form of shrinking losses due to hacking incidents.

A Cybersecurity Utility could, for example, better understand its customers' technical capabilities and needs using churn management techniques otherwise employed to anticipate when customers are about to cancel a subscription by looking at engagement with the product. This would help design products that are within customers' capability to deploy fully and correctly; something that is far from the norm currently. Similarly, by gamifying the experience and letting users track their progress, the interface could be made at least somewhat informative. On the other hand, trends in threat development could indicate adversary capabilities and objectives, allowing more targeted defense against active attacks and better proactive measures. Risk profiles could also be created at each node for each connection based on what behavior is observed and knowledge of the complete set of vulnerabilities, plus any newly discovered or suspected threats.

Another option is for the Utility to develop "invisible" security features which reduce some risk of loss and are hands-off for users [14]. For example, redundant encrypted backups or microsegmentation could limit the scale of loss or compromise. Similarly, running a script that checks all of a user's credentials against known compromised credentials, and automatically updating any found to have been posted, could prevent identity theft entirely unbeknownst to the user.

In these ways, the Utility could evaluate the level of effort required of the motivated minority based on the majority's expected contribution as well as adversaries' activity and intensity. In effect this lets the market determine what level of social loss is acceptable because the most incentivized and informed (i.e., the minority) are the only participants and transaction costs are nearly zero [12].

4.5 Evaluation for Stakeholders

A Cybersecurity Utility as described cannot exist by fiat because of the inability to compel volunteer users to use the service. As a result, it must produce economically significant value in order to exist. As mentioned above, the Utility will have to compete with existing tools and services as well as the do-nothing alternative. One possibility is the system we have proposed, which leverages economies of scope in addition to economies of scale. The interconnected nature of the ICT ecosystem means the applications on the network are inextricably interconnected, making a holistic perspective, coupled with specialized capability providers, an especially valuable tool in establishing and maintaining network security.

In the system we have described, coordination with the Utility occurs because the interaction satisfies the *economic* constraints on information flow imposed by individual end users. Stakeholders benefit indirectly, but significantly, from the reduction in successful attacks and losses at the system level that result. One of the Utility's responsibilities is therefore to ensure that the chosen implementation continues to meet the constraints implicitly demanded by end users. It should be noted that these constraints are distinct from the perceived economic value in that they are better represented as binary (e.g. information is or is not shared based on whether it can be shared anonymously) rather than a continuous scale (-10 to +$10 based on features). Constraints then can be thought of as bright red lines and perceived economic value as a dependent variable that determines users' decision rules. The Utility is in effect most concerned with keeping subjective perceived economic value above zero for all users. This is in effect a restatement of the Unique Individual Objective game's requirement that the uncooperative majority not incur any cost in order for it to be possible for a coordinated, motivated minority to change the group's behavior. We have now arrived at this conclusion via the Coasean Single-Owner analysis as well as extending Bindewald's game theoretic result for coordinating group behavior, supporting our proposal's feasibility.

Consider also the sources of perceived value to each stakeholder. We segment the stakeholders by their present use to determine present sources of value because this will form their reference value.

Users with average risk (i.e., those having no specific reason to be targeted) gain far more from connectivity to the network than information security alone. Thus, these individuals consider relatively little of the social cost or benefit of their behaviors, and must be appealed to with private benefits. Typical users in this category include most individuals and many SMBs. One possibility that arises when a Utility is responsible for aggregating and delivering capabilities is to deliver and orchestrate more than an individual user pays for in order to reduce its own expense of responding to incidents in the future. Features such as automatic encrypted backups or micro-segmentation fall into this category of proactive harm mitigation and reduce the cost of unavoidable compromises. These features are simultaneously lowest cost to deliver due to their relative simplicity and ubiquity, and highest perceived value add for this particular user group. Existing solutions charge extra for these "premium" features, likely due to the fact that the paradigmatic business model is one that caters to individuals seeking to make private investment in security for personal benefit. By instead delivering these features either free or heavily subsidized, these users receive the greatest increase in perceived value possible, giving the best chance of enticing these users to continue using the tools offered by the Utility. Another option is to change the rules of the game as Bindewald suggests. This could be seen as removing options, but a more accurate characterization would be the establishment of a value-optimizing default rule such that the effect of those who make no changes is optimal. Law and economics frequently use this type of *ex ante* hypothetical bargaining analysis to evaluate the merits of various legal rules, and one interpretation of contract law views the establishment of a baseline set of optimal default rules as a function of the law, confirming that the interpretation of limiting optionality as a matter of efficient default rule selection is neither uncommon nor untoward.

Users with above average security risk value protection in addition to connectivity. This may include high-profile individuals, journalists, activists, firms in regulated industries, or firms with significant value in digital assets, such as intellectual property. For these users, the risk of loss is pointed enough that tangible reductions in risk are more easily attributable to non-probabilistic causes. In other words, the risk is more readily identifiable and distinct from average or background risk. Firms in this category currently face some of the most difficult incentive issues due to the varying motivations for valuing security. In particular, smaller firms subcontracting in regulated industries may primarily be motivated by meeting the requirements of the prime contractor at lowest cost, while the prime may be motivated by protecting its intellectual

property. This is an example of *ex ante* moral hazard, where the damages the small firm may have to pay in the event it is found responsible for a breach are capped at the lesser of its enterprise value and the amount of damage determined to be result from its breach of duty. In corporate finance, this is sometimes referred to as a "debt overhang" situation (or colloquially "heads I win, tails you lose"), and the typical resolution is to renegotiate the terms of their agreement such that management has some "skin in the game." The cybersecurity Utility effectively fulfills this reallocative role by creating a distributed market where these direct risk relationships are considered in aggregate, and the party with the higher value is made to reveal this preference.

The market for security products would be positively affected by a Cybersecurity Utility. It is empirically a market for lemons, in large part due to the information asymmetries between providers and consumers. In our interviews, consumers were generally uninterested in becoming more informed. Since network security benefits from a greater share of "cooperating" network participants, the theoretically maximum level of security achievable with indifferent participants is attained where the contribution required is as close to zero as possible [6]. This has two implications. First, any information security scheme that envisions indifferent participants making independent investment decisions will most likely not have a high level of network security. Second, a Utility resolves this by recalibrating incentives as discussed above. In particular, because the Utility occupies a position where its profit is greatest when security is highest, and because the Utility can transcend the boundaries of the relationships shown in Figure 1, it has both the ability and incentive to seek out misallocations in information security expenditure and correct them simply by delivering defensive capability. The end result is that the only market participants are informed and properly incentivized, driving lower-quality firms out of the market. The Utility's position at the center gives it visibility into a greater amount of data at once than is currently visible by most security providers. The Utility can then steer investment selectively into the products and developing technologies that offer the greatest promise based on developing threats. Similarly, the Utility's *ex ante* access to would-be victims gives it the opportunity to minimize its later response costs by implementing proactive measures such as microsegmentation and automated, redundant, encrypted backups. It is no coincidence that these are the same features the most indifferent users value the highest, because they minimize the severity of the situations that experts are later called in to clean up.

The Utility is incentivized to produce the highest level of protection achievable per dollar. In other words, because the Utility's compensation is reduced by any expense analagous to a transaction cost, they are best served by seeking out and achieving the highest marginal reduction in breach probability. Since vulnerabilities can arise in any section of code, not just the predictably security-sensitive portions, the Utility will have to coordinate with developers of other software products as well. In fact, it may ultimately be a more efficient arrangement for the Utility to do bug testing than for each developer to test their own products.

4.6 Measures of Success

The Cybersecurity Utility is primarily intended to reduce losses to adversaries. As a result, it derives much of its value from the reduction in uncertainty and severity of cyberattacks. The metrics for success correspond to the Utility's performance in carrying out the scanning, information, incentive, and learning functions described above. Alternatively, performance could be measured by proxy metrics, such as the number of supported systems or ability to deter attacks. The ability to measure deterrence is desirable but difficult to calculate a theoretical maximum.

5 CONCLUSION AND FUTURE WORK

In this paper, we presented a framework and Cybersecurity Utility Platform for omnichannel cybersecurity. The solution improves the cost-benefit investments of individual firms and increases their resilience in continuous time. It also leads to a platform whereby even an uncooperative majority can benefit from the cybersecurity of the network with little cost to end users.

While the solution we describe is not yet implemented, we have confidence that it will demonstrate improved cybersecurity. Many deployment use cases are possible, including industry sectors such as the energy sector and defense industrial base. We see potential for benefit to the Cybersecurity Maturity Model Certification (CMMC) framework developed by the U.S. Department of Defense [11]. Within five levels of maturity are processes and practices to protect information of varying sensitivity and threats. Establishment of a Cybersecurity Utility may aid and support companies in achievement and compliance with CMMC.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful comments. The views and conclusions expressed in this paper are those of the authors, and do not necessarily represent those of the Department of Defense or U.S. Federal Government.

REFERENCES

[1] George A Akerlof. 1978. The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in economics*. Elsevier, 235–251.
[2] Ross Anderson. 2001. Why information security is hard-an economic perspective. In *Seventeenth Annual Computer Security Applications Conference*. IEEE, 358–365.
[3] Nadya Bartol. 2012. ISO Cyber Security and ICT SCRM Standards. https://www.acsac.org/2010/program/case/wed-1330-Bartol.pdf
[4] Johannes M Bauer and Michel JG Van Eeten. 2009. Cybersecurity: Stakeholder incentives, externalities, and policy options. *Telecommunications Policy* 33, 10-11 (2009), 706–719.
[5] Eckart Bindewald. 2017. A survey suggests individual priorities are virtually unique: Implications for group dynamics, goal achievement and ecology. *Ecological Modelling* 362, C (2017), 69–79. https://doi.org/10.1016/j.ecolmodel.2017.
[6] Eckart Bindewald and Shady S. Atallah. 2017. Achieving multiple goals via voluntary efforts and motivation asymmetry. *Ecological Modelling* 354 (2017), 37 – 48. https://doi.org/10.1016/j.ecolmodel.2017.03.010
[7] Sandor Boyson. 2014. Cyber supply chain risk management: Revolutionizing the strategic control of critical IT systems. *Technovation* 34, 7 (2014), 342 – 353. https://doi.org/10.1016/j.technovation.2014.02.001 Special Issue on Security in the Cyber Supply Chain.
[8] Robert M Brady, Ross J Anderson, and Robin C Ball. 1999. *Murphy's law, the fitness of evolving species, and the limits of software reliability*. Technical Report. University of Cambridge, Computer Laboratory.
[9] Guido Calabresi. 1968. Transaction Costs, Resource Allocation and Liability Rules–A Comment. *The Journal of Law and Economics* 11, 1 (1968), 67–73.
[10] Juan Francisco Carías, Leire Labaka, José María Sarriegi, and Josune Hernantes. 2019. Defining a cyber resilience investment strategy in an industrial internet of things context. *Sensors* 19, 1 (2019), 138.

[11] Carnegie Mellon University and The Johns Hopkins University Applied Physics Laboratory LLC. 2020. Cybersecurity Maturity Model Certification (CMMC), Version 1.02. https://www.acq.osd.mil/cmmc/docs/CMMC_ModelMain_V1.02_20200318.pdf

[12] R. H. Coase. 1960. The Problem of Social Cost. *The Journal of Law & Economics* 3 (1960), 1–44. http://www.jstor.org/stable/724810

[13] Jean-Pierre H Dubé, Günter J Hitsch, and Pradeep K Chintagunta. 2010. Tipping and concentration in markets with indirect network effects. *Marketing Science* 29, 2 (2010), 216–249.

[14] Josiah Dykstra. 2020. Invisible Security: Protecting Users with No Time to Spare. In *2020 IEEE 6th International Conference on Collaboration and Internet Computing (CIC)*. IEEE, 184–190.

[15] Richard A. Epstein. 1973. A Theory of Strict Liability. *The Journal of Legal Studies* 2, 1 (1973), 151–204. http://www.jstor.org/stable/724030

[16] Lawrence A. Gordon and Martin P. Loeb. 2002. The Economics of Information Security Investment. *ACM Trans. Inf. Syst. Secur.* 5, 4 (Nov. 2002), 438–457. https://doi.org/10.1145/581271.581274

[17] Lawrence A Gordon, Martin P Loeb, and Lei Zhou. 2020. Integrating cost–benefit analysis into the NIST Cybersecurity Framework via the Gordon–Loeb Model. *Journal of Cybersecurity* 6, 1 (2020), tyaa005.

[18] Cormac Herley. 2009. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*. 133–144.

[19] C Derrick Huang, Qing Hu, and Ravi S Behara. 2008. An economic analysis of the optimal information security investment in the case of a risk-averse firm. *International journal of production economics* 114, 2 (2008), 793–804.

[20] J. L. Kelly Jr. 1956. A New Interpretation of Information Rate. *Bell System Technical Journal* 35, 4 (1956), 917–926. https://doi.org/10.1002/j.1538-7305.1956.tb03809.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.1538-7305.1956.tb03809.x

[21] Kerry Krutilla, Alexander Alexeev, Eric Jardine, and David Good. 2021. The Benefits and Costs of Cybersecurity Risk Reduction: A Dynamic Extension of the Gordon and Loeb Model. *Risk Analysis* (2021).

[22] Tyler Moore. 2010. The economics of cybersecurity: Principles and policy options. *International Journal of Critical Infrastructure Protection* 3, 3-4 (2010), 103–117.

[23] PCAST. 2013. Immediate Opportunities for Strengthening the Nation's Cybersecurity. https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/PCAST/pcast_cybersecurity_nov-2013.pdf

[24] Anirudh Ramachandran and Nick Feamster. 2006. Understanding the network-level behavior of spammers. In *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*. 291–302.

[25] Paul Rosenzweig. 2013. Cybersecurity and the Least Cost Avoider. https://www.lawfareblog.com/cybersecurity-and-least-cost-avoider

[26] SCRLC. 2013. Supply Chain Risk Management Maturity Model. Issue May. http://www.scrlc.com Supply Chain Risk Leadership Council.

[27] Rebecca Slayton. 2017. What is the cyber offense-defense balance? Conceptions, causes, and assessment. *International Security* 41, 3 (2017), 72–109.

[28] Eric Stegman, Jamie Guevara, Nick Michelogiannakis, Shreya Futela, Shaivya Kaushal, and Sneha Sharma. 2020. *IT Key Metrics Data 2021: Overview*. Gartner.

[29] Bamford v. Turnley, 122 ER 25. 1862. .