

Security Circumvention: To Educate or To Enforce?

Debabrata Dey
Foster School of Business
University of Washington
ddey@uw.edu

Abhijeet Ghoshal
College of Business
University of Louisville
abhijeet.ghoshal@louisville.edu

Atanu Lahiri
Jindal School of Management
University of Texas at Dallas
atanu.lahiri@utdallas.edu

Abstract

Deliberate circumvention of information systems security is a common behavioral pattern among users. It not only defeats the purpose of having the security controls in place, but can also go far beyond in terms of the total damage it can cause. An organization grappling with circumvention can try to (i) train its users, or (ii) take on enforcement measures, or adopt a combination of the two. In this work, we look at the trade-off between these two very different approaches towards circumvention and try to gain some insights about how an organization might wish to tackle this menace.

Keywords: *IT Security, security control, circumvention, work-around, training, monitoring.*

1. Introduction

Over the last few years, as interconnected, networked information systems have become more and more prevalent, security and assurance of information technology (IT) have gained tremendous importance within all types of organizations. In order to reduce the risks of security breaches, organizations have often invested heavily in IT security. For example, the global IT security market “topped \$75 billion in 2015,” and is expected to hit a whopping \$170 billion by 2020 [24]. A large portion of this expenditure goes towards different types of security *controls*¹—controls that are supposed to reduce, or even eliminate at times, loopholes through which hackers can gain access to a system [8]. However, despite such heavy investments, security breaches are a common occurrence in the networked world of today.

¹ Every security control is deployed with its own policy specification. For example, access authentication through userid and password is a control, and its associated policy is essentially a definition of what makes certain strings of characters acceptable as userid and password. Similarly site blocking is a security control, and the actual list of blocked sites is its policy specification. In this paper, we use the shorthand “control” to mean both—a security control as well as its associated policy specification.

It turns out that a significant part of the problem lies not with the controls themselves, but with the human users who interact with these systems.² Prior research has consistently found that, in the face of a stricter security control, users often try to bypass or work around it, essentially diluting the ability of the control to effectively thwart security attacks [e.g., 2, 14, 17, 19, 20, 21]. *Security circumvention*—a situation where a user works around a security control, thereby defeating its purpose, at least partially—can take on many different forms. When a user, faced with a stringent requirement for a complex password, writes it down on a sticky note to attach it to the corner of his monitor—or to the back of the keyboard for that matter—it is a case of security circumvention [21]. Similarly, when a user, faced with a list of sites that are blocked from a company network, deliberately connects to a third-party virtual private network (VPN) to access those very blocked sites, it is also a case of circumvention. When a doctor, facing repeated timeouts after a period of inactivity, places a Styrofoam cup on a proximity sensor to fool the system to think that it is still in use, she is actually circumventing a security control [19]. And, when a nurse on duty walks away from his station, if only for a few minutes, without signing out of the system and, hence, leaving it vulnerable, it is certainly a case of circumvention as well. As Blythe et al. [3] put it, security circumvention occurs any time “users either fail to follow an intended protocol or workflow process, or actively take steps to defeat it.”

Why do users circumvent? While it is not possible to pinpoint one single reason, prior field work has identified several. First, the inconvenience caused

² We use the term “user” in a generic sense to mean anyone who interacts with the system. Therefore, our user could be an end user, a developer, a tester, a security expert, or a system administrator. Irrespective of the actual category, these human users can and do engage in activities that bypass the intended purpose of security controls [3].

by security controls or policies may often be the primary motivation—the trouble of remembering a long complex password (only with a permissible combination of different keyboard characters) or the frustration at having to repeatedly sign in to a system after periodic timeouts from inactivity are only two of many such examples [21]. Second, the urgency to engage in an activity that has been forbidden—the need, for example, to access certain sites that have been blocked—could also be very strong, which might prompt the user to circumvent by, say, connecting to a “dark” VPN [10]. Third, a user may also not be familiar with the repercussions of his own activities, that is, he may grossly underestimate the extent of damage posed by his actions. For example, an employee taking a toilet break for a couple of minutes may think that leaving the system signed in for that small time window is perhaps harmless, when, in reality, it could pose a significant security risk. Finally, a user may not be fully conversant with the security policies and their ramifications; such would be the case when an executive shares her password with her personal assistant without realizing that it is not only against her own company’s security policies, but now could also be a federal crime [5].

Intentional or not, circumvention can pose significant security risks to an organization [e.g., 3, 4, 14, 25]. Not only do such activities dilute the effectiveness of a control, they could also open doors to newer attacks that were not present before the control was put in place. Consider the case of circumventing blocked sites using third-party VPNs. Before blocking certain sites, an organization does face certain risks from those sites. However, if a user connects to a third-party VPN to reach those sites after they have been blocked from the company network, not only do they bear the risks posed by those blocked sites, but there is also an additional threat coming from the VPN provider, typically an illegitimate site posing additional risks that were not present earlier.

Given these realities, the issue of circumvention and how to prevent it has become a critical one for many organizations [e.g., 3, 14, 19]. Essentially, there are two approaches an organization can take [13, 18, 25, 28]: On one hand, it can invest in *enforcement measures*, that is, towards better monitoring (auditing) of user activities and penalizing violations when detected. On the other, it can also invest in providing sufficient *training* to its employees, making them aware of the current security controls and policies, as well as the ramifications of circumventing them. Both these approaches can cost an organization significant time and effort [18]. Naturally, the following research questions emerge:

- How effective are these two approaches, and should an organization prefer one over the other?
- If so, under what conditions should one approach be preferable to the other?
- Do these two approaches act as substitutes or complements of each other?

Clearly, these are important questions for any organization grappling with how to stop security threats posed by circumvention.

To answer these questions, we set up a simple modeling experiment using constructs borrowed from standard microeconomic models. We consider a user base that is heterogeneous in the benefits derived from, and costs incurred for, circumvention. We also consider organizational losses arising out of security loopholes as well as circumvention. A game is setup where the organization first chooses its levels of investment for training and enforcement. Based on that, the users choose whether or not to circumvent. We solve for the equilibrium of this sequential game and perform comparative statics on the cost parameters for training and enforcement to answer our research questions.

Our answers are interesting. We find that neither approach dominates the other one throughout. We also find that neither approach is sufficient on its own, and a combination is usually the best way forward. These two results, although somewhat intuitive, provide important insights about organizational policy on circumvention. We also find that, in a significant portion of the parameter space, these two approaches to prevent circumvention complement each other, and curiously, their levels either increase or decrease together as the parameters are changed. This is surprising. Given that both the approaches work towards the same end—prevention of circumvention—it is natural to expect that they would be substitutes, but we find that they could actually be complementary.

2. Literature

Our research overlaps with the extant literature on the economics of information security. This literature has grown substantially in recent years. Of particular relevance are papers that discuss issues pertaining to investments in IT security. For example, Gordon and Loeb [11] consider the decision to invest in security using an economic model that weighs the cost of security against the expected loss from attacks. In a subsequent empirical work, Gordon and Loeb [12] make the point that such cost-benefit analysis is quite common in practice. Cavusoglu et al. [9] argue that a game-theoretic

approach actually leads to a more effective security investment decision when compared to such decision-theoretic approaches. This is because the attackers often strategically respond to the level of investment, which makes their activity level endogenous to the decision to invest in information security. Herath and Herath [16] propose a real-options model to evaluate security investment decisions. Anderson and Moore [1] and Varian [26] look at the provision of security from the perspectives of underlying incentives, legal liability, and network externalities. Our main contribution to this literature is that we discuss how investments should be targeted—should they primarily target enforcement or should they emphasize training—in an environment in which employees engage in circumvention of security controls.

The literature on security controls is also important. Lee et al. [22] examine the role of security standards in a context where not all security controls are verifiable. Our focus is neither on verifiability nor on standards. We simply focus on the consequences of circumvention and the economic losses resulting therefrom. As mentioned already, our motivation is actually rooted in the long stream of research that highlights how users often bypass and work around security controls, in essence rendering them useless [2, 14, 17, 19, 20, 21]. According to Koppel [19], in many cases, these circumventions have become the norm, rather than the exception. As Heckle [15] notes, they have become the norm so much so that clinicians in some hospitals offer logged-in sessions to one another as a matter of professional courtesy. Along similar lines, Blythe et al. [3] observe that users often see circumvention as a necessary means to get their job-related activities done, and not because they intend anything malicious.

Although researchers have talked about circumvention being common, we could not locate substantial literature that investigates the economics of this phenomenon and offers strategic insights. Addressing this gap is important, however, because it is well documented in prior literature that breaches have serious financial implications [6, 7] and that preventing circumvention can go a long way [3, 14, 19]. In order to address it, we develop a model to capture the essential elements mentioned in the NIST handbook [13]. In particular, we borrow the idea that both “the *dissemination* and the *enforcement* of policy are critical issues” [13, p. 146]. Enforcement requires, among other things, auditing—users are less likely to circumvent when they are afraid that their actions will be recorded and audited—and punitive actions when circumventions are detected.

Dissemination through employee training is equally important because circumvention often arises out of ignorance on behalf of the user. If users knew how easily such behavior could be exploited for malicious attacks and the true repercussions of such attacks on the organization, many of them, if not all, would have desisted from such activities. Curiously, we find that, at times, dissemination and enforcement work as substitutes and, at others, as complements.

3. Model Setup

We consider a sequential game in which the organization first chooses whether or not to implement a stricter security control and, if it does, the levels of training and enforcement to accompany this particular control. User training, the level of which is denoted x , may include but is not limited to [13, 28]:

- seminars and training sessions for the control,
- repeated reminders explaining the control, and
- videos and other links related to the control.

Likewise, y is a proxy for the enforcement level; anti-circumvention enforcement measures may include, among other things [4, 13]:

- physical inspection and monitoring,
- automated (real-time or batch) inspection and monitoring,
- analysis of users’ activity logs, and
- increasing the penalty level for violations.

Given the organization’s choice of training and enforcement levels, our users make a decision about whether they should circumvent the security control. As is customary, we traverse this timeline backwards, starting with the user behavior.

3.1. User Behavior

We consider a normalized user base of mass one. We assume that users are heterogeneous and have different valuations, w , for circumventing a specific security control. A user who faces a higher level of inconvenience from a stricter control should have a higher w . Similarly, a user with a greater urgency to engage in a prohibited activity—such as visiting a blocked site—is also likely to have a higher w .

Users are also heterogeneous in the cost or expected penalty, p , they incur when engaging in circumvention. There are many aspects that may appear as a part of this p ; we list a few below [17]:

- the expected penalty imposed on the user, that is, the probability of getting caught times the penalty on getting caught,
- the cost associated with learning the tricks to circumvent a specific control, especially when the control is not easy to bypass, and

- the moral cost in causing real harm to his or her own organization.

Therefore, a user $\langle w, p \rangle$ has a net benefit of $v = w - p$, and his individual rationality (IR) would dictate him to circumvent if and only if $v > 0$. We assume (see Figure 1):

ASSUMPTION 1. *The net benefit, v , is uniformly distributed over an interval $[a, b]$.*

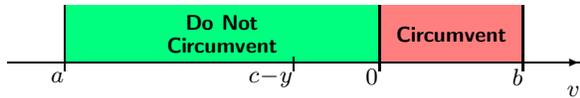


Figure 1. Users Choose to Circumvent (or Not Circumvent) Based on Their Net Benefit, v

We expect the eventual distribution of v to depend on the level of training, x , and the enforcement level, y . In particular, we let the end points, a and b , depend on x and y . To understand this dependence, we look at the underlying mechanism through which enforcement and training activities manifest themselves in the user’s net benefit, v . As mentioned earlier, standard enforcement activities involve physical and virtual monitoring and analysis; they essentially increase the chance of detection [4, 13]. In other words, when enforcement level increases, the probability for a user to get caught while circumventing increases. Enforcement activities could also involve imposing heftier penalties, including a termination of employment, when caught [27].

Since the expected penalty faced by a user is the probability of getting caught times the actual penalty when caught, all in all, higher enforcement levels imply a higher expected penalty, $\frac{\partial p}{\partial y} > 0$, and a correspondingly lower net benefit to all users. Therefore, it is reasonable to assume that the mean of v , denoted \bar{v} , decreases with y , that is, $\frac{\partial \bar{v}}{\partial y} = -\frac{\partial p}{\partial y} < 0$. Accordingly, in Figure 1, we capture this mean as $\bar{v} = \frac{a+b}{2} = c - y$.

In contrast, the impact of training on the users’ benefit is not as direct. To understand, we must recognize that a significant component of users’ heterogeneity in v arises out of imperfect or partial information.³ Uninformed or partially informed users are likely to under- or over-estimate, among other

³ Viewed differently, even if all users had the same value for v in reality, a lack of perfect information guarantees that they do not know this true value. Therefore, users with different levels of (mis-)information would have different perceptions of v , leading to a distribution around the true value.

things, the real cost borne by the organization, the difficulty in working with a new security control, or even the expected penalty imposed by an organization [13, 18, 28]. Now, since training activities can effectively reduce the information gap and bring users closer to the mean, \bar{v} , it is logical that training programs ought to reduce the overall variance in v , by reducing user’s uncertainty (lack of perfect information) about the true benefit.⁴ Put differently, x reduces the spread of the distribution by bringing a and b closer to the mean (without impacting the mean itself), while y shifts the distribution (and, hence, only its mean) to the left.⁵ We assume the following simple functional forms for a and b :

ASSUMPTION 2. *The endpoints of interval $[a, b]$ over which users are distributed are given by:*

$$a(x, y) = c - y - \frac{\alpha}{1+x}, \text{ and}$$

$$b(x, y) = c - y + \frac{\alpha}{1+x},$$

where $c < 0$ and $\alpha > 0$ are constants, and $c + \alpha > 0$.

The choice of $c < 0$ is reasonable. Since c represents the expected value of v for $y = 0$, a positive c would imply that a majority of the users find circumventing the control to be so desirable and so convenient that the control itself is of little value and the organization is perhaps better off not implementing it in the first place. On the other hand, we expect α , the parameter that captures the level of heterogeneity—the spread of the distribution—to be positive. We also expect $b(0, 0)$ to be positive. Otherwise, not a single user would engage in circumvention, even if the organization spends nothing on training and enforcement; the issue of circumvention then becomes moot. Hence, we only consider those parameter values for which $b(0, 0) = c + \alpha > 0$ holds.

We argued earlier that, although information provided during training can reduce the heterogeneity

⁴ Point to note here is that, although training can reduce the variance, it cannot completely remove the heterogeneity. This is because there is also an intrinsic part to users’ heterogeneity. Even when all users have perfect information, they will exhibit heterogeneous behavior simply because of their intrinsically different preferences.

⁵ Now, is it at all possible that training also influences the mean, and enforcement, the variance? Of course, it is possible. However, as discussed earlier, the primary impact of training ought to be on the variance, and that of enforcement, on the mean. Our conceptualization, in other words, considers only the primary impacts of training and enforcement, while abstracting away the secondary ones in order to avoid confounding factors.

among users, it can do so only to an extent; it can never fully eliminate this heterogeneity. An important point to note that our model specification conforms well to this requirement—in our setup, there is no finite x for which $a(x, \cdot) = b(x, \cdot)$, implying that, for all finite values of x , some level of heterogeneity does remain.

Clearly, the density function for the net benefit, v , can be expressed as:

$$f(v) = \begin{cases} \frac{1+x}{2\alpha}, & \text{if } v \in [c-y-\frac{\alpha}{1+x}, c-y+\frac{\alpha}{1+x}], \\ 0, & \text{otherwise.} \end{cases}$$

Since we know from users' individual rationality (IR) that every user with a $v \in (0, b]$ would engage in circumvention, we can easily find the size of this segment as a function of x and y :

$$s(x, y) = b(x, y) \times \frac{1+x}{2\alpha} = \frac{1}{2} + \frac{(1+x)(c-y)}{2\alpha}. \quad (1)$$

3.2. Organization's Problem

We assume that the organization's expected loss from the loophole it is trying to block using the stricter control is L ; in other words, if there were no circumvention, implementing the control is worth L to the organization. However, a portion of this saving is likely to be lost to circumvention; we assume it to be proportional to the fraction of users circumventing the control and write it as $L\mu s(x, y)$, where $\mu > 0$ is the constant of proportionality and $s(x, y)$ is as given in (1). Without loss of generality, we can normalize L to one:

ASSUMPTION 3. *The net value of the security control after circumvention is $(1 - \mu s)$, where $\mu > 0$ represents the relative magnitude of the damage caused by circumvention.*

To complete our model specification, we need to consider the costs associated with training, x , and enforcement, y . We assume a standard quadratic form for both:

ASSUMPTION 4. *The costs associated with training and enforcement levels x and y are $\frac{\beta x^2}{2}$ and $\frac{\gamma y^2}{2}$, respectively, where $\beta, \gamma > 0$.*

Combining all of these, we can write the organization's decision problem as:

$$(\mathbf{P}) \quad \max_{x, y \geq 0} \quad z = \left(1 - \mu s(x, y) - \frac{\beta x^2}{2} - \frac{\gamma y^2}{2} \right), \\ \text{s.t.} \quad z \geq 0 \text{ and } s(x, y) \geq 0,$$

where $s(x, y)$ is as in (1). The constraint $z \geq 0$ guarantees that a new security control, along with its

associated training and enforcement, is not deployed if there is no net benefit from doing so, when compared to doing nothing at all. In other words, if there are no feasible solutions to (\mathbf{P}) —that is, if the optimal value of the objective function becomes negative without the constraint $z \geq 0$ —the control is clearly not worth implementing, and the organization should not pursue it any further. Finally, the constraint $s(x, y) \geq 0$ ensures that, once an organization has eliminated circumvention completely, it should not want to spend any more on x and y .

It must be noted that the maximization problem in (\mathbf{P}) can also be transformed into a minimization problem as:

$$(\mathbf{P}') \quad \min_{x, y \geq 0} \quad z' = \left(s(x, y) + \frac{\beta' x^2}{2} + \frac{\gamma' y^2}{2} \right), \\ \text{s.t.} \quad \mu z' \leq 1 \text{ and } s(x, y) \geq 0,$$

where $\beta' = \frac{\beta}{\mu}$ and $\gamma' = \frac{\gamma}{\mu}$. Although the new formulation in (\mathbf{P}') involves fewer parameters in the objective function itself, it turns out that the parameter μ cannot be eliminated from (\mathbf{P}') as μ shows up in the constraint $\mu z' \leq 1$, which is logically equivalent to $z \geq 0$.

4. Results

When solved, (\mathbf{P}) or (\mathbf{P}') separates the parameter space into three distinct regions:

PROPOSITION 1. *Let $g_1(\gamma) = \gamma(2\alpha(\gamma c^2 + \mu) - \mu c) - \sqrt{\gamma(\gamma c^2 + 2\mu)(2\alpha\gamma c - \mu)^2}$ and $g_2(\gamma) = 8\alpha^2\gamma + \mu^2 - 4\alpha\gamma\mu(c + \alpha)$. Further, define:*

$$h_1(\gamma) = \begin{cases} \frac{\mu^3}{2\alpha g_1(\gamma)}, & \text{if } \gamma > \frac{\mu}{2\alpha(c+\alpha)} \\ \infty, & \text{otherwise,} \end{cases} \\ h_2(\gamma) = \begin{cases} \frac{\mu^2(2-\gamma c^2-\mu)}{g_2(\gamma)}, & \text{if } \gamma > \frac{\mu^2}{4\alpha(c\mu+\alpha(\mu-2))} \\ \infty, & \text{otherwise.} \end{cases}$$

Then, the following equilibrium outcomes emerge:

- **Circumvention Region:** *When $h_1(\gamma) \leq \beta \leq h_2(\gamma)$, the organization implements the control, but some users circumvent it.*

- **No Circumvention Region:** *When $\beta < h_1(\gamma)$, the control is implemented, and no users circumvent it.*

- **No Control Region:** *When $\beta > h_2(\gamma)$, the organization decides not to implement the control.*

The result in Proposition 1 is better visualized in Figure 2. As can be seen from this figure, when β is small or γ is small, or both, the organization can effectively eliminate all circumvention by users,

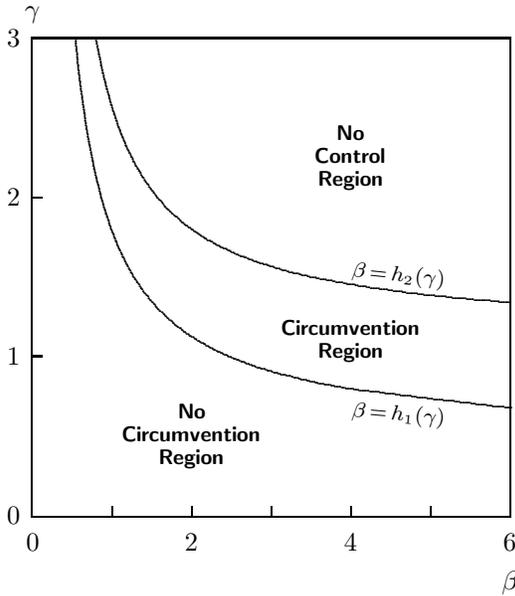


Figure 2. Relevant Partitions of the (β, γ) Space; $\alpha = 2$, $\mu = 3$, $c = -\frac{1}{3}$

either by providing sufficient training or by increasing the level of enforcement, or by using a combination of the two. When β and γ are both high, the organization simply cannot afford either approach and decides to not implement the control at all. In the middle, where β and γ take on moderate values, the control is adopted, along with a combination of training and enforcement; circumvention is controlled to an extent, but cannot be fully eliminated. Any organization struggling with the issue of circumvention should belong to this middle region.

We now look at the organization's optimal choices of training and enforcement levels:

PROPOSITION 2. *The optimal levels of training and enforcement can be expressed as follows:*

- **Circumvention Region** ($h_1(\gamma) \leq \beta \leq h_2(\gamma)$):

$$x^* = \frac{\mu(\mu - 2c\alpha\gamma)}{4\alpha^2\beta\gamma - \mu^2}, \text{ and}$$

$$y^* = \frac{\mu(2\alpha\beta - c\mu)}{4\alpha^2\beta\gamma - \mu^2}.$$

- **No Circumvention Region** ($\beta < h_1(\gamma)$): x^* is the only real and positive solution of:

$$\frac{\alpha\gamma(c(1+x) + \alpha)}{(1+x)^3} - x\beta = 0,$$

and $y^* = c + \frac{\alpha}{1+x^*}$.

- **No Control Region** ($\beta > h_2(\gamma)$): $x^* = y^* = 0$, trivially.

The results in Proposition 2 are presented in Figure 3. Proposition 2 and Figure 3 tell us an interesting story; they show that both the approaches, training and enforcement, are required for dealing with circumvention and that neither approach can achieve it on its own. For any values of β and γ satisfying $\beta \leq h_2(\gamma)$, both x^* and y^* are positive, implying that it is more effective to use the two approaches in combination, rather than in isolation. In other words, neither approach dominates the other in preventing circumvention. Of course, the correct mix depends on their relative costs, as parameterized by β and γ .

PROPOSITION 3. *The optimal level of training, x^* , is decreasing in β , and the optimal enforcement level, y^* , is decreasing in γ . Mathematically, $\frac{\partial x^*}{\partial \beta} \leq 0$ and $\frac{\partial y^*}{\partial \gamma} \leq 0$.*

The results in Proposition 3 are intuitive. As the cost for training (or enforcement) goes up, we would expect to see the organization cutting down on its level. On the other hand, as the marginal cost goes down, we should expect the level to increase. These trends are clearly discernible from Figure 3 as well.

Now, when x^* is reduced, what happens to y^* , and vice versa? That is, we now turn our attention to whether x and y act as substitutes or they complement each other. The answer to this question is curious and is found in our next result:

PROPOSITION 4. *The optimal levels of training and enforcement, x^* and y^* , complement each other in the circumvention region, but they act as substitutes in the no circumvention region. More specifically, $\frac{\partial y^*}{\partial \beta} \geq 0$ and $\frac{\partial x^*}{\partial \gamma} \geq 0$ in the no circumvention region, but $\frac{\partial y^*}{\partial \beta} \leq 0$ and $\frac{\partial x^*}{\partial \gamma} \leq 0$ in the circumvention region.*

This is counterintuitive. Training and enforcement are both means to the same end, that is, mitigation of circumvention. We would naturally expect that they are substitutes—when one becomes costlier, we expect it to be reduced with an accompanying increase in the level of the other. While this intuition remains valid in the no circumvention region, it no longer holds in the circumvention region, a situation an organization is most likely to find itself in. In the circumvention region, these two approaches complement each other. As a result, when β (or γ) increases, it not only results in a lower x^* (or y^*) but also in a lower y^* (or x^*). That these two approaches work hand-in-hand towards fulfilling an organizational goal has important implication for managers in charge of IT security in their organizations.

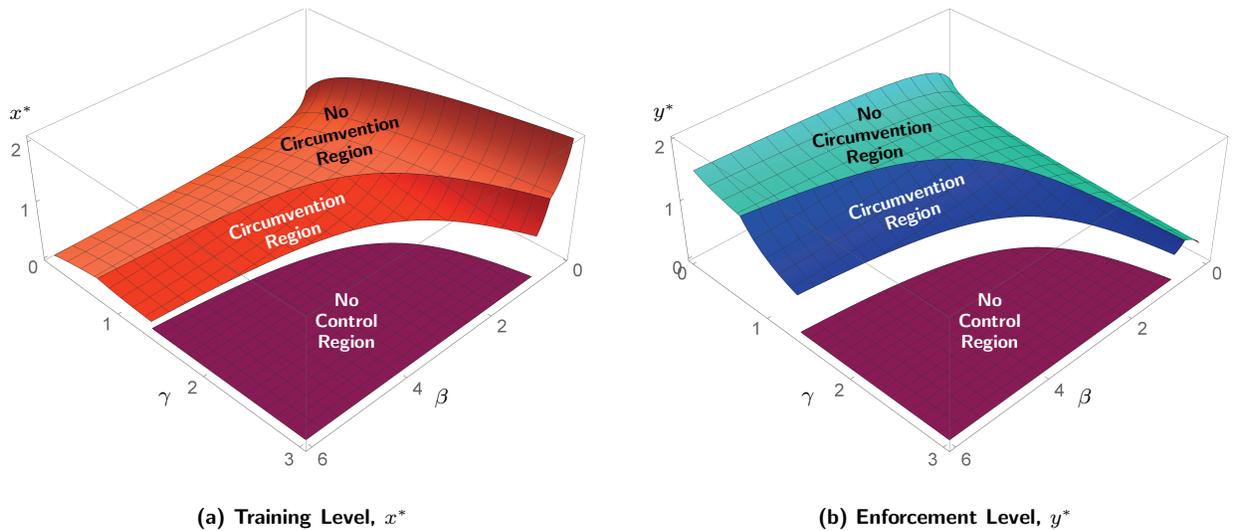


Figure 3. Optimal Levels of Training and Enforcement as Functions of β and γ ; $\alpha=2$, $\mu=3$, $c=-\frac{1}{3}$

We now turn our attention to the other parameters in the model. In particular, we are interested in the impacts of c and α , the parameters that define the distribution of the users' net benefit, as well as that of μ , the parameter representing the severity of circumvention. To that end, we look at the impact of these parameters on $\beta = h_1(\gamma)$, the boundary that separates the circumvention region from the no circumvention one.

PROPOSITION 5. *The circumvention region expands with c and α but shrinks with μ . Mathematically, $\frac{\partial h_1(\gamma)}{\partial c} \leq 0$, $\frac{\partial h_1(\gamma)}{\partial \alpha} \leq 0$, and $\frac{\partial h_1(\gamma)}{\partial \mu} \geq 0$.*

Put differently, when c and α increase, they pull the boundary between circumvention and no circumvention, $\beta = h_1(\gamma)$ in Figure 2, downward or to the left, thereby widening the region where circumvention happens. This can be clearly seen in Figure 4, where the original curve (in black) moves towards the red and blue curves as α and c increase, α from 2 to 3 and c from $-\frac{1}{3}$ to $-\frac{1}{10}$, respectively. This result is also along the expected lines. When c or α increases, it either moves the distribution to the right or expands its spread. In either case, $b(x, y)$ moves to the right, making it costlier for the organization to eradicate circumvention completely. In contrast, a higher μ makes user circumvention costlier for the organization, making it less tolerant towards such behavior. The net result is a right or upward shift in the boundary $\beta = h_1(\gamma)$, as shown by the shift to the green curve in Figure 4 when μ increases from 3 to 5.

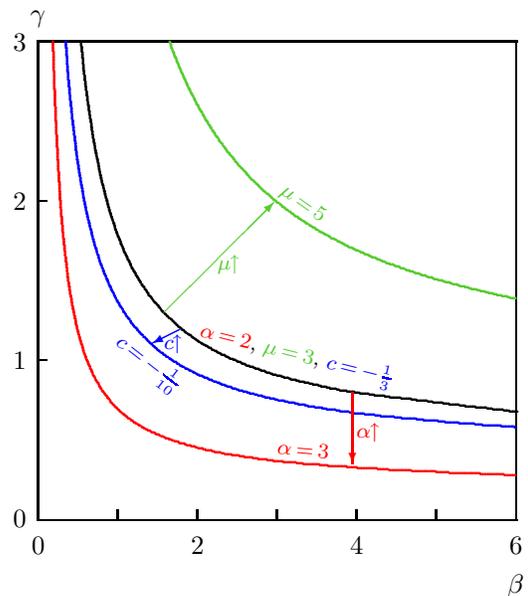


Figure 4. Comparative Statics on the Boundary, $\beta = h_1(\gamma)$

5. Discussion

Security and assurance of IT have taken the center-stage in an organization's IT policy and investment decisions. And, the issue of circumvention has simply added fuel to that fire. Typically, the failure to secure technology using technology has made organizations throw more money at acquiring even more technology [23]. Our work shows that such unidimensional approach to security might be pointless.

The role of education and training—creating awareness among users from different walks—has

long been recognized [4, 13, 28]. Our work clearly supports this point of view; we find that organizations are better off heeding this advice and investing in user training to raise the awareness level about security policies and their necessity.

What implementable insights do we find? First, in Proposition 1, we see that there is a significant portion of the parameter space in which it may be optimal for the organization to tolerate some level of circumvention. That certainly explains the situation observed in many organizations today. Facing higher costs for enforcement and training, organizations often recognize the futility in trying to eradicate circumvention fully. In fact, if the situation is sufficiently grim, it may be optimal for an organization to not deploy the security control in the first place.

In Proposition 2, we learn that neither approach to circumvention dominates the other, and they work best in combination, not in isolation. In other words, organizations need not immediately and fully shift their focus away from enforcement activities and adopt training as the only circumvention strategy. Circumvention is best addressed when both the approaches are used in a judicious mix.

Of course, this work, based on a positive modeling experiment, does not shed much light on what that judicious mix should constitute, but can certainly highlight some of its characteristics. Some of these characteristics are presented in Propositions 3 and 4. They are largely consistent with our basic understanding of how an organization might strategically behave in such situations; they give us the confidence that our setup carries a reasonable resemblance with the reality.

The most notable of the above characteristics—the non-monotonicity of x^* w.r.t. γ or, equivalently, the non-monotonicity of y^* w.r.t. β , in Proposition 4—is, however, counterintuitive. And, it also has a clear, actionable implication. It tells us that the strategy an organization might undertake in the face of circumvention can suddenly change once circumvention has been effectively dealt with. When dealing with user circumvention, an organization may initially invest in both enforcement and training. However, once it has achieved a full eradication of such behavior among its users, the organization may afterward relent in one of the approaches. In other words, an organization’s enforcement and training policies, along with its broader security policies, are likely to evolve with time.

Finally, Proposition 5 tells us how, depending on the context, an organization’s strategy may

shift between tolerating and not tolerating user circumvention, and whether training and enforcement should be treated as substitutes or complements of each other. For example, if a control is inherently annoying to begin with and most users are seriously affected, c would be relatively high. In this case, the boundary $h_1(\gamma)$ would shift left, making circumvention the likely equilibrium outcome. A good example of this would be a stringent password policy in an organization where each user is responsible for a number of passwords for logging on to a number of systems. However, if a single sign-on technology or key-chain technology is rolled out in the same organization, the users may no longer be as annoyed, and c could be low. What is interesting is that, the presence of a single sign-on or key chain technology could make training and enforcement substitutes for each other although, in their absence, training and enforcement would likely be complements. Similarly, one can compare situations involving different levels of α or μ . If a security control impacts a wide variety of users, we would have a high α to begin with. On the other hand, if the control is directed towards one particular group of employees in a closely knit unit of the organization, α would be low. There would accordingly be a bearing on the anti-circumvention strategy of the firm. Thus, Proposition 5, taken together with our other results, provides a manager in charge of IT security and assurance a well-rounded understanding of the issue of circumvention, adding to his ability to tackle the issue more effectively.

6. Conclusion

Deliberate circumvention by its user can pose significant security risks to an organization. Our work shows that investing only in enforcement activities to mitigate such behavior is futile, if it is not accompanied by proper training and education to increase the level of awareness among users.

How do our results relate to current industry practices? First, it has long been recognized that technology for the sake of technology does not work, and creating awareness among users is an important dimension towards effective security control [3, 4]. Our results seem to echo this sentiment by highlighting the need to have appropriate training programs. In recent times, there has been a growing recognition among industry professionals of the role played by user awareness and training. The search for effective training programs at a lower cost has led to the development of third-party training materials and modules—both generic and customized—many

of which can actually be offered online quite cheaply. Such developments are well in line with the observation from our modeling experiment.

Our model setup makes certain simplifying assumptions. For example, we assume that the impacts of the two approaches, enforcement and training, on the distribution of users' net benefit are very distinct. One impacts the mean, while the other, the variance. This abstraction is a simplification as, in reality, they both can influence the mean and variance at the same time. Although we can speculate how such a generalization might bias our results, we leave a complete, rigorous analysis to future research.

Further, we do not consider any budget constraints in dealing with circumvention, while in practice, organizations often contend with limited budget available for investing in enforcement and training. Once again, without proper analysis, it is difficult to speculate how that might impact our results. Despite these limitations, the purpose of this work would be amply served if it has succeeded in drawing attention to the need for user training in this interconnected business environment of today. Perhaps, McGowan [23] is correct when she concludes, "Institutions cannot hesitate in the goal to educate their employees."

References

- [1] R. Anderson and T. Moore, "The Economics of Information Security," *Science*, 2006 (314:5799), pp. 610–613.
- [2] A. Beautement, M. Sasse, and M. Wonham, "The Compliance Budget: Managing Security Behavior in Organisations," in *Proceedings of the 2008 New Security Paradigms Workshop*, Lake Tahoe, CA, September 2008, pp. 47–58.
- [3] J. Blythe, R. Koppel, and S. Smith, "Circumvention of Security: Good Users Do Bad Things," *IEEE Security and Privacy*, 2013 (11:5), pp. 80–1720.
- [4] P. Bowen, J. Hash, and M. Wilson, "Information Security Handbook: A Guide for Managers," NIST Special Publication 800-100, <https://www.nist.gov/publications/information-security-handbook-guide-managers>. October 2006. Accessed May 2, 2017.
- [5] F. Caldwell, "Why Sharing Passwords Is Now Illegal And What This Means for Employers And Digital Businesses," *Forbes*, <https://www.forbes.com/sites/ciocentral/2016/08/23/why-sharing-passwords-is-now-illegal-and-what-this-means-for-employers-and-digital-businesses/#3a46e3ff3a46>. August 2016. Accessed May 9, 2017.
- [6] K. Campbell, L. A. Gordon, M. P. Loeb, and L. Zhou, "The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market," *Journal of Computer Security*, 2003 (11:3), pp. 431–448.
- [7] H. Cavusoglu, B. Mishra, and S. Raghunathan, "The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers," *International Journal of Electronic Commerce*, 2004 (9:1), pp. 70–104.
- [8] ———, "The Value of Intrusion Detection Systems in Information Technology Security Architecture," *Information Systems Research*, 2005 (16:1), pp. 28–46.
- [9] H. Cavusoglu, S. Raghunathan, and W. T. Yue, "Decision-Theoretic and Game-Theoretic Approaches to IT Security Investment," *Journal of Management Information Systems*, 2008 (25:2), pp. 281–304.
- [10] J. Goodchild, "Workarounds: 5 Ways Employees Try to Access Restricted Sites," CSOnline, <http://www.csoonline.com/article/2125818/access-control/workarounds-5-ways-employees-try-to-access-restricted-sites.html>. August 2010. Accessed on May 2, 2017.
- [11] L. A. Gordon and M. P. Loeb, "The Economics of Information Security Investment," *ACM Transactions on Information and System Security*, 2002 (5:4), pp. 438–457.
- [12] ———, "Budgeting Process for Information Security Expenditures," *Communications of the ACM*, 2006 (49:1), pp. 121–125.
- [13] B. Guttman and E. Roback, "An Introduction to Computer Security: the NIST Handbook," NIST Special Publication 800-12, <https://www.nist.gov/publications/introduction-computer-security-nist-handbook>. October 1995. Accessed May 2, 2017.
- [14] M. Harrison, R. Koppel, and S. Barlev, "Unintended Consequences of Information Technologies in Health Care—An Interactive Sociotechnical Analysis," *Journal of the American Medical Informatics Association*, 2007 (14:5), pp. 542–549.
- [15] R. R. Heckle, "Security Dilemma: Healthcare Clinicians at Work," *IEEE Security and Privacy*, 2011 (9:6), pp. 14–19.
- [16] H. S. B. Herath and T. C. Herath, "Investments in Information Security: A Real Options Perspective with Bayesian Postaudit," *Journal of Management Information Systems*, 2008 (25:3), pp. 337–375.

- [17] C. Herley, “So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users,” in Proceedings of the 2009 New Security Paradigms Workshop, Oxford, UK, September 2009, pp. 133–144.
- [18] I. Kirlappos and M. Sasse, “What Usable Security Really Means: Trusting and Engaging Users,” in Proceedings of the Second International Conference on Human Aspects of Information Security, Privacy, and Trust, Crete, Greece, June 2014, pp. 69–78.
- [19] R. Koppel, S. Smith, J. Blythe, and V. Kothari, “Workarounds to Computer Access in Healthcare Organizations: You Want My Password or a Dead Patient?,” in Studies in Health Technology and Informatics, vol. 208, IOS Press, 2015, pp. 215–220.
- [20] R. Koppel, T. Wetterneck, J. Telles, and B.-T. Karsh, “Workarounds to Barcode Medication Administration Systems: Their Occurrences, Causes, and Threats to Patient Safety,” Journal of the American Medical Informatics Association, 2008 (15:4), pp. 408–423.
- [21] V. Kothari, J. Blythe, S. Smith, and R. Koppel, “Agent-Based Modeling of User Circumvention of Security,” in Proceedings of the 1st International Workshop on Agents and Cybersecurity, Paris, France, 2014.
- [22] C. Lee, X. Gang, and S. Raghunathan, “Mandatory Standards and Organizational Information Security,” Information Systems Research, 2016 (27:1), pp. 70–86.
- [23] J. McGowan, “Stop Throwing Money at Cybersecurity,” Banking Blog, <http://bankingblog.celent.com/2016/10/12/stop-throwing-money-at-cybersecurity/>. October 2016. Accessed May 2, 2017.
- [24] S. Morgan, “Worldwide Cybersecurity Spending Increasing To \$170 Billion By 2020,” Forbes, <https://www.forbes.com/sites/stevemorgan/2016/03/09/worldwide-cybersecurity-spending-increasing-to-170-billion-by-2020/#5f8913876832>. March 2016. Accessed May 4, 2017.
- [25] D. Upton and S. Creese, “The Danger from Within,” Harvard Business Review, <https://hbr.org/2014/09/the-danger-from-within>. September 2014. Accessed May 6, 2017.
- [26] H. R. Varian, “Managing Online Security risks,” The New York Times, 2000, <http://www.nytimes.com/library/financial/columns/060100econ-scene.html>.
- [27] L. Walsh, “Making an example: Enforcing company information security policies,” SearchSecurity.com, <http://searchsecurity.techtarget.com/Making-an-example-Enforcing-company-information-security-policies>. March 2004. Accessed May 2, 2017.
- [28] M. Wilson and J. Hash, “Building an Information Technology Security Awareness and Training Program,” NIST Special Publication 800-50, <https://www.nist.gov/publications/building-information-technology-security-awareness-and-training-program>. October 2003. Accessed May 2, 2017.