International Review of Law & Economics 73 (2023) 106116

Contents lists available at ScienceDirect

International Review of Law & Economics

journal homepage: www.elsevier.com/locate/irle

Incentives for investments in defensive technology: An economic analysis of the Safety Act

Mattias K. Polborn *,1

Department of Economics, Vanderbilt University, United States of America
Department of Political Science, Vanderbilt University, United States of America
Department of Economics, University of Cologne, Germany

ARTICLE INFO

Keywords:
Safety Act
Liability limits

ABSTRACT

Civilian targets of terrorist or criminal attacks (e.g., sport stadiums, chemical or nuclear industry; infrastructure such as ports or pipelines) are often owned by the private agents who choose how to guard against potential attacks. This creates an important externality problem, as some of the benefits of better protection accrue to other private agents who would suffer from an attack. We analyze a model in which a social planner wants to provide incentives for the deployment of defensive technologies. Our results show that some features of the Safety Act, enacted after the 2001 terror attacks, are probably counterproductive.

1. Introduction

When terrorists attack in order to cause mass casualties, they usually do this in settings where security is managed, at least in part, by private firms. For example, in the 9/11 attacks, the terrorists used planes owned by United Airlines and American Airlines to attack the World Trade Center, a building owned by a private sector entity.

Thus, private sector firms play a large role in protecting against terrorist attacks, both in terms of loss avoidance (e.g., making it harder to hijack planes) or loss mitigation (e.g., improving evacuation plans for buildings). However, the efficient provision of anti-terror defense is complicated by an obvious externality problem: Most of the damage avoided by the successful thwarting of a terrorist attack benefits other agents, not the firms that have to pay for the expenses necessary to use the defensive technology. Thus, we would expect that the equilibrium level of anti-terrorism defense chosen by private firms is suboptimal.

A similar externality problem arises also in some non-terrorist cybersecurity attacks that are undertaken for purely criminal purposes. On May 7th of 2021, Colonial Pipeline, a major American oil pipeline company, suffered a ransomware cyberattack that forced Colonial Pipeline to halt the flow of fuel for 5 days, leading to the shutdown of thousands of gas stations throughout the Southeastern United States, as well as massive gasoline shortages and dramatic price increases in the affected markets, with substantial consequences for millions of consumers. While Colonial Pipeline paid $5 million in Bitcoin as ransom, it is fairly clear that most of the economic damages caused by the cyberattack actually accrued to third parties.

While third parties can attempt to recuperate on those losses by suing, as over 10 thousand gas stations tried to by bringing a class action lawsuit against Colonial Pipeline, this is not feasible for the millions of customers who suffered harm from Colonial Pipeline's inability to safeguard their network. As such, the incentives for security improvements are weaker than socially optimal.

While the standard economic policy response to positive externalities is to subsidize the underprovided activity, the U.S. Congress chose to take a different path in the aftermath of 9/11. Specifically, in 2002, Congress passed the Support Anti-terrorism by Fostering Effective Technologies Act (henceforth, the SAFETY Act) as an inexpensive way to subsidize the development and usage of anti-terrorism technologies.

This law empowers the Secretary of Homeland Security to certify certain anti-terrorism technologies (henceforth ATTs), and to limit the liability of users and suppliers of approved ATTs if the technology fails to stop an attack. Third-party liability in case of a successful attack is a substantial concern for many firms. For example, in 2005, a New York court found the owner of the World Trade Center, the Port Authority of New York and New Jersey, more than 65% liable for third party damages due to the 1993 World Trade Center bombing (Spence et al., 2012). By offering a safe harbor from such litigation, the SAFETY Act is essentially intended to act as an implicit subsidy for ATT investment.[2]

A broad range of often industry-affiliated publications lay out the benefits of the SAFETY Act for different industries (e.g., Carpentier and

[*] Correspondence to: 2301 Vanderbilt Place, Nashville, TN 37235, United States of America.
E-mail address: mattias.polborn@vanderbilt.edu.

[1] I would like to thank Tigran Polborn and two anonymous referees for very helpful comments.

[2] A more detailed description of the SAFETY Act is provided in Section 3.

https://doi.org/10.1016/j.irle.2022.106116
Received 27 July 2022; Received in revised form 5 November 2022; Accepted 18 November 2022
Available online 24 November 2022
0144-8188/

Finch, 2012 for the natural gas and electricity industry; Biagini, 2008 for the airport management industry; Bryant, 2016 for the sports industry). Moreover, Harter (2006) and Knake (2016) suggest that SAFETY Act-like protections should be extended to other hazards, such as natural disasters or (non-terrorist) cyber attacks. It is therefore important to understand the economic effects of this type of liability-limiting law on loss-avoidance and loss-mitigation technology adoption.

We model the incentives for ATT provision by a firm that is a target of a potential attack. The firm chooses, at a cost, the ATT technology it wants to use, from a set of available technologies. The technology determines the probability with which an attack will be defeated. The firm also submits its chosen technology for certification by a regulatory agency. The probability of certification is increasing in the quality of the technology, and if the technology is certified, then the firm's liability is limited to a prespecified level.

Because the liability limitation has an economic value and the probability of certification is increasing in the quality of the ATT, the firm has an incentive to choose a better technology. However, conditional on receiving the certification, the firm faces a reduced loss from a successful attack, and therefore has reduced incentives to choose a better ATT. In general, either effect can dominate. In particular, we show that, if the ATT has a sufficiently large probability of successful defense, then the SAFETY Act is counter-productive in the sense that the ATT chosen is lower than in the absence of the SAFETY Act. Furthermore, we show that these problems are exacerbated in a dynamic setting: Liability reductions which the firm acquired in the past reduce its incentive for further innovation.

Note that our analysis of the efficiency effects of the SAFETY Act focuses on anti-terrorism technology development and/or deployment, taking as given the probability of a terrorist attack. In addition to affecting the choice of ATT, the implicit subsidy through the liability reduction might also lead to excessive entry in the affected industries.[3] However, this effect is quite indirect, and there might also be equilibrium effects on the behavior of terrorists (e.g., better defenses of a particular firm may redirect attacks towards other firms). It is beyond the scope of this article to analyze all general equilibrium effects of the SAFETY Act and alternative such as subsidies for defensive technology for private sector firms.

As argued by Shavell (1984), direct regulation is a frequent and relatively attractive policy option when some part of the damages remains uncompensated under liability. A standard objection to regulation — namely, that the regulatory agency cannot judge what makes sense for the private agent — does not apply here because the SAFETY technology designation or certification represents a procedure where the agency does exactly that: judge what makes sense for the private agent.

In fact, while there is no direct safety regulation in our model, some of the industries that are taking advantage of SAFETY Act provisions are already heavily regulated, such as the nuclear industry. In those industries, if the Office for the Implementation of the Safety Act just certifies the standard of care that firms have to choose to comply with the agency that regulates them directly, the liability limitations from the SAFETY Act essentially just constitute a lump-sum transfer to those firms, equal in value to the expected liability avoided. Yet, the scope of the SAFETY Act is considerably larger, and firms that applied successfully for certification or designation under the Act include many industries where safety against terrorist attacks is generally not directly regulated.[4] Our model is most relevant for those industries.

Our paper contributes to the existing literature on the effects of liability rules (Shavell, 2009). Essentially, a liability limitation such as the one provided by the SAFETY Act can be interpreted as being similar to a combination of a strict liability rule (for the amount up to the limit) and a negligence rule for damage amounts that exceed the limit (Shavell, 1980). Furthermore, while the required level of care that leads to a liability exclusion under a negligence rule is usually up to interpretation by the court, under the SAFETY Act the technology is pre-approved by the regulatory agency, and so the firm is fairly certain about the level of liability it faces in case of a successful attack.

The paper proceeds as follows. Section 2 reviews related literature. Section 3 provides more detailed information about the SAFETY Act. The model is presented in Section 4, and analyzed in Section 5. Section 6 considers a dynamic extension of the model, and Section 7 concludes.

2. Related literature

Liability in tort law is usually assigned according to either strict liability rule (under which the injurer must compensate the victim, even if the injurer was not at fault) or negligence rule, under which there is a certain "standard of care" that absolves the injurer from liability (see Posner, 1973 and Shavell, 1980 for classical contributions; Cooter, 1991 for an excellent short review; and Shavell, 2009 for a monographic treatment).

The SAFETY Act in our model is essentially an opportunity for a potential injurer to pre-certify a standard of care that absolves it from liability for damages exceeding a limit chosen by the regulator. Below the limit, the injurer remains liable under a strict liability rule. Furthermore, the SAFETY Act only defines an option for the injurer, who can also choose to operate under the existing legal system.

Our main argument that the SAFETY Act has an ambiguous effect on the equilibrium level of care is related to the classic argument about a random standard of care under negligence rule (see, for instance, Miceli, 1997, pp.45) where an injurer escapes liability completely with a probability that is increasing in its level of care. The prospect of achieving such an exemption increases care incentives, while the possibility of already having reached the required level of care diminishes the incentive for further care.

Like our paper, the seminal contribution of Endres and Bertram (2006) is also concerned with how liability rules affect technology development. Specifically, they analyze how strict liability and negligence rule affect an injurer's incentives to develop improvements to its accident prevention technology. In their setting, strict liability implements the first best; in contrast, negligence rule only does so if the court knows not only the technology that the injurer chose to implement, but also the cost of implementing different technologies. Similarly, in our setting, the informational requirements for the SAFETY Act to implement the efficient technology are very demanding.

Polborn (1998) analyzes the choice between strict liability and negligence rule in a setting where, if the accident occurs, the injurer cannot pay for the losses because of limited assets. Likewise, inability to pay for losses may be a principal reason in our setting why firms may not choose the efficient antiterrorism technology, if there is no additional incentive provided.

Demougin and Fluet (1999) and Nell and Richter (2003) provide additional justifications for using negligence rules over strict liability, such as asymmetric information in the management of the injurer firm, and risk allocation in the presence of incomplete insurance markets. None of these effects are present in our model framework, but it is clearly conceivable to argue that they also provide some arguments in favor of the incentives provided by the SAFETY Act.[5]

[3] See, for example, Faure and Fiore (2009) for a discussion of how the nuclear liability subsidy induces an artificial competitiveness of nuclear energy.

[4] Lieberman et al. (2019) mention, for example, industrial manufacturing, real estate, healthcare and financial services firms as beneficiaries of the SAFETY Act.

[5] For example, the liability limitation for the firm using a SAFETY Act approved technology moves the economic harm from the firm to (usually a multitude of) other agents who now cannot sue the firm. In the presence of incomplete insurance markets, such a spread of the risk may be beneficial.

3. The safety act

In an attempt to foster the development and deployment of defensive anti-terrorist technologies, Congress passed the *Support Anti-terrorism by Fostering Effective Technologies (SAFETY) Act* as past of the Homeland Security Act of 2002. The following information is mostly based on the Department of Homeland Security's "SAFETY Act 101 briefing", available at https://www.safetyact.gov/externalRes/refDoc/refGroup/8/SAFETY%20Act%20101%20Briefing.pdf

The Act provides legal liability protections for sellers and users of anti-terrorism technologies. These liability protections apply against claims that satisfy both of the following two conditions. First, the firm must have applied for and been granted such a liability protection for the technology by the Office of SAFETY Act implementation; and second, the Secretary of Homeland Security must have declared that the claim resulted from an Act of Terrorism. Several hundred applications are submitted each fiscal year, and approximately half of them are approved.

In the context of the SAFETY Act, the definition of anti-terrorism technology is quite broad and includes physical products (e.g., blast mitigation materials), services (e.g., screening services at a sports stadium), and software and other forms of intellectual property.

Technologies can either be given a "designation" (which comes with a particular liability limit in case of a terrorist attack), or a "certification" (which provides essentially full protection similar to the "government contractor defense" that shields military contractors against liability when their products fail to perform). In addition to limited liability, both designation and certification also confer certain procedural benefits, such as exclusive action in federal court, no joint and several liability for non-economic damages, and a prohibition against punitive damages.

The technological standards required for the two levels are somewhat opaque ("proven effectiveness" in the case of designation, and "high confidence it will continue to be effective" in the case of certification). Observe that, in contrast to the patent system, it is not explicitly required that a technology surpasses the performance of the currently best available technology.

4. The model

We consider the problem of a firm F that faces the risk of an *attack*. If the attack is successful, the firm has a direct loss of S. Furthermore, third parties are also harmed by a successful attack and experience a loss equal to $Z+X$; here, Z is the part of the loss for which F is liable, while X is the externality, i.e., the additional loss that is beyond what the firm is legally liable for.

For example, suppose that F operates a pipeline that is shut down for two weeks by a cyber attack because F did not sufficiently prepare for the possibility of an attack. In this case, refineries (who were directly contracting with F) might have a strong case of contract nonfulfillment against F so that their damages are contained in Z. In addition, it is highly likely that, in this scenario, gas stations would run out of gas and gas prices would spike, harming consumers. These losses are captured by X as they would be very difficult to recover through a legal process, either because courts might be hesitant to recognize a legal obligation of F to supply to consumers who do not have a contract with F, or, in case that such an obligation was construed, the damage then would likely exceed F's assets so that X is unrecoverable. Of course, if F is not held liable for the losses (for example, because its prevention effort is deemed non-negligent), then all third party losses are captured in X.[6]

Let $p \in [\underline{p}, 1]$ denote the probability that F is able to defeat the attack, in which case no harm is caused. Here, $\underline{p}$ is interpreted as the baseline technology. Alternatively, F can, at a cost $K(p)$, choose a technology $p > \underline{p}$, where $K(\cdot)$ is a strictly increasing and convex function, with $\lim_{p\to 1} K'(p) \to \infty$. That is, all marginal improvements are costly, and increasingly so for technologies that defeat attacks with a higher probability; in particular, achieving total security against an attack is extremely costly.[7] We furthermore assume that $K(\underline{p}) = 0$, i.e., costs are measured relative to the baseline technology.

In the spirit of the institutions set by the Safety Act, we assume that the task of incentivizing the deployment of defensive technology is given to a regulatory agency that can certify a technology. Certification of F's technology is completely voluntary (i.e., the firm is free simply to choose some defensive technology p and not pursue certification); if certification is granted, it reduces F's legal liability for third party losses in case of a successful attack by $\ell \in [0, Z]$ to $Z - \ell$, (and, correspondingly, increases X by the same level); thus, $\ell = 0$ corresponds to an ineffective regulation (where certification does not change F's liability), while $\ell = Z$ means that F is completely shielded from liability.[8] The primary question of interest of our analysis is how the level of ℓ affects F's incentives to improve its defensive technology.

The probability of certification is given by a nondecreasing function $\Phi(\cdot)$ that maps p into a probability of certification in $[0, 1]$. In principle, one could also think of Φ as another tool for optimization for the regulatory agency. In fact, we will show in Section 5.2 that, if this is the case, the regulatory agency can often implement the first best, and that the best policy involves a step function; that is, there is a level $\hat{p}$ such that $\Phi(p) = 0$ for $p < \hat{p}$ and $\Phi(p) = 1$ for $p \geq \hat{p}$.

However, it often appears realistic that, from the perspective of when the firm has to choose its technology p, there is some uncertainty as to whether any given technology would be certified. In practice, only about 55 percent of certification applications between 2016 and 2021 were granted, and costly rejections would not occur if the certification process was not random from the firm's point of view.

Defensive technologies are often unique to each firm's specific situation so that the agency cannot establish a track record of which technologies it would certify and which ones it would reject. Furthermore, the standards for certification set out in the Safety Act ("Proven effectiveness" or "high confidence that technology will continue to be effective" for designation and certification, respectively) are somewhat vague, but reference only technical aspects of the technology. Note that the standards do not make any reference to the "state of the art", and it is explicitly not the case that a certified technology needs to represent an improvement (or even a significant improvement) over existing technology, as would be the standard for the issuance of a patent.

In contrast, the law states that the regulatory agency has considerable leeway in choosing the liability limit. In particular, it may consider "[t]he possible effects of the cost of insurance [for the liability limit] on the price of the product, and the possible consequences thereof for development, production, or deployment of the technology".

[6] We should emphasize that, in our setting, internalizing this externality is not an option for the policy maker. In other contexts, the joint use of liability and fines has been discussed (see, for instance, Goerke (2003)). In principle, imposing a fine X in the event of a successful attack can induce the first-best level of p. However, imposing additional fines on firms after having been victims of a terror attack may be politically difficult. Furthermore, specifying the correct amount of the fine X ex-ante would be necessary and might be difficult in many contexts (e.g., in the Colonial Pipeline example).

[7] The latter assumption is useful as it generates interior solutions, but is not fundamentally necessary for our qualitative results to obtain.

[8] Such a complete shield is sometimes called the *government contractor defense*.

5. Analysis

5.1. Social optimum

We start with an analysis of the socially optimal level of defensive technology. A social planner would minimize the sum of all types of expected losses from a successful attack and the cost of developing defensive technology p,

$$(1-p)(S+Z+X)+K(p) \tag{1}$$

Differentiating with respect to p and rearranging yields the following optimality condition

$$K'(p)=S+Z+X. \tag{2}$$

The straightforward interpretation is that, in an optimum, the marginal cost of developing a slightly better technology must equal the marginal damage avoided through a slightly better technology. Observe, furthermore, that the convexity of the cost function K is sufficient for global optimality of the solution of (2). For reference in the following, let p^* denote the solution of (2).

5.2. Best case scenario: Clear certification standards

In this section, we assume that the regulatory agency can specify a level $\hat{p}$ such that $\Phi(p)=0$ for $p<\hat{p}$ and $\Phi(p)=1$ for $p\geq\hat{p}$, as well as set the liability reduction level $\ell\in[0,Z]$.

The firm always has the option not to participate in the certification process; in this case, it is completely liable for Z in case of a successful attack. Thus, it minimizes

$$(1-p)(S+Z)+K(p), \tag{3}$$

by choosing p such that

$$K'(p)=S+Z. \tag{4}$$

Denote the solution of this first-order condition p_0.[9]

Given a liability reduction ℓ, the maximum $\hat{p}$ that the firm would be willing to accept gives the firm the same utility level as they could get by opting out, and thus satisfies

$$(1-p_0)(S+Z)+K(p_0)=(1-\hat{p})(S+Z-\ell)+K(\hat{p}), \tag{5}$$

where the left-hand side is the firm's total expected cost when they get no liability reduction, and the right-hand side is their expected cost if they develop technology $\hat{p}$.

The following Proposition 1 shows that higher values of liability reduction ℓ allow for the maximum implementable safety level to be larger. In particular, the largest possible liability reduction $\ell=Z$ maximizes the largest value $\hat{p}$ such that the firm is still willing to choose $\hat{p}$ and be certified.

Proposition 1. *For any $\ell\in[0,Z]$, there is a value of $\hat{p}\in[p_0,1)$ that solves (5), and it is increasing in ℓ.*

Proof of Proposition 1. Note that, except when $\ell=0$, the minimizing value of p for the right-hand side is smaller than p_0, which implies that the right-hand side is increasing in p for all values of p greater than p_0.

The right-hand side of (5) evaluated at $\hat{p}=p_0$ is smaller than the left-hand side. Furthermore, the right-hand side is continuous in $\hat{p}$, and at $\hat{p}=1$, the right-hand side is larger than the left-hand side evaluated at p_0. This implies the existence of a unique value of $\hat{p}$ such that (5) holds with equality.

The right-hand side of (5) is decreasing in ℓ, which implies (together with the fact that the right-hand side of (5) is increasing in $\hat{p}$) that $\hat{p}$ is increasing in ℓ. ■

Proposition 1 shows the promise of the Safety Act for improving security. Specifically, if there is a generous reduction in liability for certified technology, then the firm is willing to implement challenging improvements in their defensive technology in order to obtain the liability reduction. In many cases, it may be possible to incentivize the firm to adopt the socially optimal level of precaution.

Note, however, that the informational requirements for the regulatory agency implementing such a scheme are quite high. The regulatory agency has to know the cost function K for developing defensive technologies, and must be able to convey the certification threshold $\hat{p}$ to the firm before they start with their development. In practice, the regulatory agency (i.e., the Office of Safety Act Implementation in the case of the Safety Act) does not set ex-ante targets for firms, but waits until firms have developed defensive technologies for which they seek certification, and only then chooses whether or not to approve them.

Ganuza and Gómez (2008) argue that, when injurers have limited assets so that in case of an accident, they are judgment-proof, then the second-best standard of care in a negligence rule should be set to reflect this problem. In particular, if the injurer is unwilling to implement the first-best level of care under a negligence rule, a lower negligence standard of care should be set that makes the injurer indifferent between the implementing this 'realistic' standard, and the level of care they would choose under strict liability when they expect to be judgment-proof in case of an accident. This is, of course, the same argument as our Eq. (5).

5.3. Firm behavior under unclear certification standards

We now turn to an analysis of the firm's defensive technology choice when the regulatory agency's certification action is ex-ante uncertain for the firm and described by a continuous and increasing function $\Phi()$, rather than a step function.

The firm minimizes the sum of their expected own losses, its liability losses, and the development and deployment cost,

$$(1-p)(S+Z)-(1-p)\Phi(p)\ell+K(p). \tag{6}$$

Observe that (6) differs from (1) in two ways. First, the firm does not consider X, the external harm, that is not subject to liability. Second, with probability $\Phi(p)$, the firm's liability in case of a successful attack is reduced by ℓ to $Z-\ell$.

Differentiating (6) with respect to p, the first-order optimality condition is

$$-(S+Z)+[\Phi(p)-(1-p)\Phi'(p)]\ell+K'(p)=0,$$

which can be rearranged as

$$K'(p)=S+Z+[(1-p)\Phi'(p)-\Phi(p)]\ell. \tag{7}$$

As we show in Proposition 2, a sufficient condition for (7) to characterize an optimum is that $\Phi()$ is weakly concave in p.

We can interpret the case without a Safety Act like regulation as setting $\Phi(p)=0$ for all p, or, alternatively, $\ell=0$. In this case, the third term on the right-hand side of (7) simply drops out, and (7) coincides with (4) which implies that in the absence of the Safety Act, technology level p_0 would be chosen. Comparing with (2), we have the standard result that the firm chooses a socially suboptimal level of precaution because it does not consider the external harm X when choosing p. The central question is whether the Safety Act increases or decreases the firm's incentives for defensive technology investment.

Proposition 2. *Let p_S denote the level of precaution chosen by the firm, i.e., the solution of (7), and define the term in square brackets in (7) as a function*

$$H(p)=(1-p)\Phi'(p)-\Phi(p).$$

Then, $p_S<p_0$ (i.e., the Safety Act decreases the level of precaution) if and only if $H(p_S)<0$.

[9] It is clear that the second order condition for minimization holds, as $K''>0$ by assumption.

Proof. Suppose that $H(p_S) < 0$ so that $K'(p_S) < S + Z$. Since K is convex and $K'(p_0) = S + Z$, this implies $p_S < p_0$. Necessity follows from an analogous argument. ■

Observe that, for p sufficiently close to 1, we always have $H(p) < 0$. This follows from the fact that Φ is bounded ($\Phi(1) = 1$, and thus Φ' is also bounded), and therefore $\lim_{p\to 1} H(p) = -1$. Thus, if a technology is actually quite successful in foiling attacks, then it follows that the liability reduction regulation is counterproductive because it lowers the level of precaution, $p_S < p_0$. Because p_0 is already lower than the socially optimal level of precaution, p^*, this decrease lowers social welfare.

In contrast, for the opposite implication to be true, i.e., the Safety Act leads to the firm choosing a higher level of precaution, it is necessary that the probability of certification increases strongly at the firm's optimal choice, *and* that this probability is not too close to 1 at the firm's optimal choice.

In a similar vein, we can analyze how the level of liability reduction ℓ affects the optimal level of care. Implicit differentiation of (7) yields

$$\frac{dp}{d\ell} = \frac{[(1-p)\Phi'(p) - \Phi(p)]}{K''(p) + [2\Phi'(p) - (1-p)\Phi''(p)]\ell}. \tag{8}$$

In an optimum, the denominator of the right-hand side is positive; this is simply the second-order condition of minimization. Thus, the sign of $dp/d\ell$ is determined by the sign of $H(p)$, the numerator. In particular, whenever $H(p) < 0$, then a marginally more generous liability limit will lead to a weaker defensive technology, and vice versa.

Intuitively, why is the Safety Act likely to decrease the optimal level of precaution? There are two effects that go in opposite directions. First, for the firm, the liability reduction provides an additional value. An increase in p improves the chance that the firm will obtain this prize, and so there is an additional incentive to improve the defensive technology. On the other hand, conditional on receiving the liability reduction, an increase in ℓ reduces the loss in case of an attack, and therefore decreases the marginal incentive to increase p.

Which of these effects dominates depends on the specific situation. However, for a technology that is already relatively good ($p \approx 1$) and has a high probability of being certified ($\Phi(p) \approx 1$), the second effect is likely to outweigh the first one. This is because, when the probability that the technology is certified is large, the marginal probability increase from quality improvement must be relatively small. In contrast, the effect that reduced liability has on incentives for further safety technology improvement remains large.

So far, we have assumed that the safety technology only affects the probability of a successful defense against the terror attack. What would change if the defense technology, instead, mitigates the damage in case of a successful attack? In this case, a liability limit completely eliminates the firm's marginal incentives to invest in reducing the damages — any reduction of damages beyond the liability limit has no private value for the firm. In this sense, our model provides a relatively favorable setting for the SAFETY Act because the gross benefit of defeating a terror attack always remains positive; it is only that liability limits reduce the amount of the gross benefit from further increases in p.

6. Dynamics

So far, we have analyzed the effects of liability limitations in a static setting where the firm chooses its safety technology once and for all. In that setting, it is theoretically possible to use liability limitations as a reward that optimally incentivizes defensive technology development and deployment, even though we have argued that, in terms of practical implementation, the effect might be more likely to be detrimental.

We now consider a simple dynamic framework in which the defensive technology in the second period can build upon the technology of the first period. In this setting, a liability reduction used to incentivize defensive technology usage in the early period reduces the incentive to achieve further improvements in the second period.

Specifically, consider a setting in which the first period looks exactly like in our basic model (with the cost function now denoted $K_1(p)$), and is followed by a second period, in which the cost of defensive technology p_2 is given by $K_2(p_2; p_1)$.

The cost $K_2(p_2; p_1)$ of defensive technology p_2 depends on p_1. We assume that p_1 becomes the new baseline technology; that is, continuing to use the same technology as the one chosen in period 1 is free (i.e., $K_2(p_1, p_1) = 0$), while "upgrading" to a level $p_2 > p_1$ costs additional money.

Consider the following simple example that illustrates the problems that can be created for second period innovations if the regulator generates incentives for first-period innovation by giving the firm liability limitations. In case of a successful attack, the firm faces an external liability risk of 1000. In addition, there is an additional external damage that is not subject to liability of again 1000.

In the first period, the baseline risk is 20%, but it can be reduced to 10%, at a cost of 150 (and, for simplicity, we assume that this is the only risk reduction option that is available for the firm).

The social value of this technology is $(0.2 - 0.1) \times 2000 = 200$. However, without any liability limitation incentives, the private value of the risk reduction is equal to the expected liability cost reduction of $(0.2 - 0.1) \times 1000 = 100$, and because this is less than the cost of acquiring the technology, the firm will not use the technology if not provided additional incentives to do so. In fact, the minimum liability reduction required to incentivize the firm is $\ell = 500$. Providing such an incentive therefore appears beneficial from a social point of view.

Now suppose that, in the second period, a technology upgrade becomes available that would reduce the risk by another 5 percentage points, at a cost of 60. This upgrade is, clearly, socially worthwhile as it reduces the expected losses by $0.05 \times 2000 = 100 > 60$.

However, given that the firm's remaining liability in case of a successful attack is only 500, the technology's private value for the firm is only $0.05 \times 500 = 25$. Moreover, there is no *further* liability reduction that makes implementation of the upgrade attractive for the firm — the value of a complete removal of liability is only $0.1 \times 500 = 50$ for the firm.

Intuitively, what happens in this example is that a liability reduction provides incentives for technology adoption in the first period, but this also reduces the possibility to provide additional incentives in future periods.

The "ammunition" that liability reduction can provide over all time periods is necessarily limited, and equal to the full legal liability Z. Thus, even if the regulator uses the Safety Act framework optimally in the first period to incentivize defensive technology adoption, the ability to provide additional incentives necessarily decreases. This would not be the case if the regulator instead used direct subsidies in order to incentivize the firm's technology adoption.

Note that this problem appears particularly acute if there is some kind of arms race between terrorists and defenders. For example, suppose that, in every period, the probability that an attack is successful is 20 percent if last period's technology is used, but that an upgraded technology can reduce this to 10 percent.

In such a world, incentives for adopting the upgraded technology have to be given in every period, and this therefore cannot be done by permanent liability reductions. In principle, persistent incentives can be provided by granting liability reductions only for one period, i.e., they have to be re-earned every period. However, this requires that the regulatory agency can determine the appropriate "period length" for which the liability reduction can be granted, and do so in advance. This likely creates practical problems; it is conceptually much easier for the regulatory agency to test a technology submitted for certification than it is to assess for how long it is going to be effective, because the latter depends on the potential attacker's technological progress which is inherently uncertain.

7. Discussion and conclusion

The fact that many targets of terrorist or criminal attacks are owned by private agents who choose how to guard against potential attacks creates an important externality problem. The benefits of better protection accrue to other private agents who would suffer from an attack.

The justifications provided for the passage of the SAFETY Act emphasize this positive externality and promote the act as providing an effective subsidy (which, nevertheless, does not cost the taxpayer any money). These are, at least politically, attractive arguments. Unfortunately, our model shows that there are significant problems when it comes to the incentives for technology development and deployment under the SAFETY Act.

A perfect regulatory agency (i.e., one that knows the firm's cost function for providing safety, and can thus calculate the efficient level of protection) could, in principle, use liability limitations to incentivize the firm to provide a higher level of protection. However, in practice, the Office for the Implementation of the Safety Act is only tasked with certifying ‘‘effective’’ anti-terrorism technologies. In this setting, our model shows, the incentives provided by the Act are ambiguous, because, conditional on certification of the technology, a liability limitation means that there is less at stake for the firm. Thus, for those firms that have a very good chance that their technology will be certified, the Safety Act will, in effect, be counter-productive.

Finally, we also show that a problem with liability limitations in a dynamic setting is that they can provide positive incentives for technology development only once, but, in following periods, diminish the incentives for further improvements because less is at stake for the firm. Overall, this leads us to being very skeptical about the SAFETY Act, as well as proposals to extend its principles to new domains.

Data availability

No data was used for the research described in the article.

References

Biagini, Raymond, 2008. Mitigating airport terrorism tort liabilities through the US SAFETY Act. J. Airport Manag. 2 (4), 318–324.

Bryant, Karen Bunso, 2016. The Safety Act, terrorism, and the National Football League. Ariz. State Univ. Sports Entertain. Law J. 6, 203.

Carpentier, Deborah A., Finch, Brian E., 2012. Tort immunity for cyber attacks through SAFETY Act. Natural Gas Electricity 29 (4), 1–7.

Cooter, Robert D., 1991. Economic theories of legal liability. J. Econ. Perspect. 5 (3), 11–30.

Demougin, Dominique, Fluet, Claude, 1999. A further justification for the negligence rule. Int. Rev. Law Econ. 19 (1), 33–45.

Endres, Alfred, Bertram, Regina, 2006. The development of care technology under liability law. Int. Rev. Law Econ. 26 (4), 503–518.

Faure, Michael G., Fiore, Karine, 2009. An economic analysis of the nuclear liability subsidy. Pace Environ. Law Rev. 26, 419.

Ganuza, Juan José, Gómez, Fernando, 2008. Realistic standards: Optimal negligence with limited liability. J. Legal Stud. 37 (2), 577–594.

Goerke, Laszlo, 2003. Road traffic and efficient fines. European J. Law Econ. 15 (1), 65–84.

Harter, Ava A., 2006. Encouraging corporate innovation for our homeland during the best of times for the worst of times: Extending Safety Act protections to natural disasters. Duke L. Tech. Rev. 6, 1.

Knake, Robert K., 2016. Creating a federally sponsored cyber insurance program. Council on Foreign Relations, Cyber Brief, November, JSTOR.

Lieberman, Joseph, Riddle, Clarine, Robertson, Mark, 2019. SAFETY Act decreases private sector risk and liability. Available at https://www.kasowitz.com/media/3585/safety-act-decreases-private-sector-risk-and-liability.pdf.

Miceli, Thomas J., 1997. Economics of the Law: Torts, Contracts, Property, Litigation. Oxford University Press.

Nell, Martin, Richter, Andreas, 2003. The design of liability rules for highly risky activities—Is strict liability superior when risk allocation matters? Int. Rev. Law Econ. 23 (1), 31–47.

Polborn, Mattias K, 1998. Mandatory insurance and the judgment-proof problem. Int. Rev. Law Econ. 18 (2), 141–146.

Posner, Richard A., 1973. Strict liability: A comment. J. Legal Stud. 2 (1), 205–221.

Shavell, Steven, 1980. Strict liability versus negligence. J. Legal Stud. 9 (1), 1–25.

Shavell, Steven, 1984. Liability for harm versus regulation of safety. J. Legal Stud. 13 (2), 357–374.

Shavell, Steven, 2009. Economic Analysis of Accident Law. Harvard University Press.

Spence, Shannon, Ali, Akmal, Taylor, Vance, 2012. The SAFETY Act: A powerful benefit for the water sector. J. Am. Water Works Assoc. 104 (7), 22–25, URL https://awwa.onlinelibrary.wiley.com/doi/epdf/10.5942/jawwa.2012.104.0113.