Information Acquisition and Liability for Unknowable Risk*

FRANCES XU LEE
Loyola University Chicago
December 16, 2022

Abstract. This paper studies a seller's incentive to acquire information on product risk both before and after a product has been sold to some buyers, recognizing that a higher amount of product adoption helps both the seller and the public to learn about the hidden risk and the seller has an advantage in learning faster. It examines the impact of strict liability with no regard of the actual knowledge of the seller or the availability of the data to the seller and the alternative of granting immunity for harm that is hard to be learned in the past. It shows that under strict liability, there is too much incentive to acquire information post-marketing to influence the public's learning and to avoid past liability. Foreseeing its own strong incentive to acquire information down the road, the seller under-invests in information acquisition pre-marketing. Even though under both regimes, the seller under-invests in pre-marketing information acquisition relative to the socially optimal behavior, such distortion is stronger under strict liability than under immunity for scientifically unknowable risk.

Keywords: information acquisition, strict liability, unknown risk, the state of the art, timing of information acquisition

JEL Classification. K13; D83

*francesxu312@gmail.com. I thank the audience of the 32nd annual meeting of ALEA, MLEA conference at University of Chicago, the law and economics seminar at The University of Texas at Austin School of Law, the economics seminar at University of Illinois Urbana-Champaign, the LET X at Georgetown Law School.

1. Introduction

Much attention has been given to the seller's incentive to acquire information about product risk *prior* to selling the product. The seminal paper Shavell (1992) focused on the incentive to acquire information under different liability rules and Polinsky and Shavell (2012) examined the incentive to acquire information under different liability and disclosure rules. However, these formal incentive models (and several others to be discussed in detail later) only studied the incentive to acquire information *prior* to selling the product.[1] In practice, very often the seller faces the question of whether to acquire information about a product after the product has been marketed and consumed for a significant amount of time. For example, Johnson & Johnson's Baby Powder was first sold in 1894 and Johnson & Johnson started testing whether it contained asbestos, a cancer-causing mineral, in 1971, when people were becoming more aware of a type of cancer called mesothelioma.

A key debate in the legal literature is whether a seller should be given immunity if at the time of marketing or distributing the product, the seller had no way of knowing the risk. This issue frequently arises in practice because consumer adoption over time generates valuable data on the product that enables a better chance of discovering the harm, and many types of harm, such as cancer, take a long time to realize. Such discovery of information is simply not possible to replicate in a laboratory environment within a short period of time prior to actual consumer adoption. For example, Facebook Inc. conducted research on whether Instagram, a social media platform it acquired in 2013, causes mental health issues among teenagers around 2018 (Wells, Horwitz, and Seetharaman 2021). Such product research would not have been possible without the data on the teenagers who have used Instagram for a period of time. That is, the detection of the harm and the association of the harm to a particular product are often only possible after enough consumers have suffered the harm.

Broadly speaking, there are two distinct liability regimes. On one side, strict liability is imposed with no regard to the seller's knowledge. On the other hand, immunity is carved out for "unknowable" harm in the language of Keeton (1970). That is, if at the time of marketing there is no data or technology available to discover the harm, the seller is not held liable. The so-called "state-of-the-art" defense, used by some courts but not all, gives a more concrete definition to what constitutes "unknowable" risks. It gives immunity to

[1] In Spier (2011), the seller learns of product risk after selling to the buyers, but the learning is passive: the observable accidents inform the seller of the risk.

sellers who have conducted research that is considered standard and comprehensive at the time but failed to discover the harm.[2] For the purpose of this paper, I will refer to the former as "strict liability" and the latter as "unknowable" or "state-of-the-art" defense. See Wade (1983) for a survey of the different liability approaches.

While granting immunity for unknowable risks seems intuitively fair and strict liability intuitively unfair, there is no formal economic model to compare the two policies in a setting where the seller may acquire costly information *after* a product has been sold. In this paper, I examine these two liability rules in a setting where the learning ability about the product improves over time for both the seller and the public as more consumers adopt the product and suffer the hidden harm, while the seller maintains an advantage in this learning and indirectly also influences the public's learning by influencing the amount of consumption data in the market.

I first analyze a model where no pre-marketing learning is possible to highlight the post-marketing incentives. Think of the Instagram example where no lab experiment can mimic the actual gradual impact of social media on teenagers. In this setup, I show that strict liability induces some socially undesirable incentives after some buyers have purchased. First, under strict liability, marketing a product without learning about it exposes the seller to not only the liability on the harm upon the future buyers, but also the harm upon the past buyers. This may induce a seller to drop a positive expected value product to obstruct the public's learning and detection of the potential harm. Second, the seller can hide a bad discovery and mitigate the product harm to reduce the chance for the public to discover the harm, which creates too much incentive to acquire post-marketing information compared to the socially optimal outcome. In other words, the value of information post-marketing is higher for the seller than for the society. Interestingly, both forms of inefficiency, dropping a positive expected value product and acquiring too much information, would disappear if the seller is given the unknowable defense. Unknowable defense allows the seller to fully internalize only the impact of its decision on the future buyers.

Then I enrich the model to allow for a state-of-the-art test to be available prior to the marketing of the product. If the product is harmful, the state-of-the-art test may reveal this

[2]Several U.S. states have enacted statutes expressly providing for the state-of-the-art defense. See, e.g., Ariz. Rev. Stat. Ann. § 12-683 (Supp. 1980); N.H. Rev. Stat. Ann. § 5507-D:4 (Supp. 1979). See also, Murphy v Owen-Illinois, Inc, 779 F.2d 340, 342-3 (6th Cir. 1985). Article 7(e) of European Union Product Liability Directive says: "The producer shall not be liable as a result of this Directive if he proves that the the state of scientific and technical knowledge at the time when he put the product into circulation was not such as to enable the existence of the defect to be discovered."

harm to the seller with some probability, but may also fail to uncover it. The immunity kicks in if the state-of-the-art test was run but did not alert the seller to any harm. On the other hand, if the state-of-the-art test was not run, the two sides of a product liability litigation would have to fight it out in court on whether the harm was "knowable". The weaker the state-of-the-art test was, the more likely that the defense can convince the court that the risk was unknowable. Under such an enriched setup, I contrast the socially optimal outcome, the outcome under strict liability, and the outcome under the state-of-the-art defense.

Foreseeing its own strong incentive to acquire information down the road to avoid the retrospective liability on the early buyers, the seller rationally reduces information acquisition pre-marketing. Interestingly, even though the seller always has an incentive to under-invest in pre-marketing information acquisition relative to the socially optimal behavior because the seller only cares about the product harm if it is caught, the seller reduces such information acquisition more under strict liability than under the state-of-the-art defense because the action of running a state-of-the-art test increases the chance of being granted immunity in the future should the harm be exposed to the public.

In the legal literature, the opponents to strict liability often criticize from a mere sense of ex post justice which we know is not an argument grounded in theoretical modeling.[3] A standard economic model of optimal care has shown that strict liability is in many settings socially optimal because it makes the seller internalize all the impact of the product. Therefore, even though ex post a seller who has taken the socially optimal action may get punished under strict liability (think of a seller who has taken the socially optimal level of care but whose product still unfortunately caused an accident), strict liability provides the right ex ante incentive. This paper re-examines this insight about strict liability when the action of concern (acquiring information about the product after more data become available) happens after some harm might have realized and the action can influence the chance that the public learns of the harm eventually. This paper shows that strict liability does not generate the socially optimal incentives in both the decision to test product and the decision to keep selling the product.

There are a few arguments in favor of strict liability, although most of these arguments are comparing strict liability to different versions of negligence, rather than unknowable defense. Unknowable defense conditions liability on the data availability or technology

[3] Schwartz (1985) repeatedly expresses views such as "requiring a firm to do what seemingly cannot be done is unfair."

of testing in the past and the information acquisition decision of the seller, whereas negligence typically conditions liability on whether the seller has exerted the optimal amount of care, a choice of the seller that I greatly simplified in this paper to focus on the information acquisition incentives. The literature finds that negligence causes excess incentive to acquire information (Shavell 1992, Kaplow and Shavell 1992) because the seller escapes liability as long as the due care is taken. However, this paper finds that the seller underacquires information prior to marketing the product under unknowable defense (and strict liability as well) due to an interaction between the information acquisition incentives before and after the product has been sold to some buyers. Casual reasoning also suggests that strict liability provides more incentive to acquire information prior to marketing the product than unknowable defense (Faure, Visscher and Weber 2016). However, this paper finds there to be less incentive to acquire information prior to marketing the product under strict liability than under unknowable defense.

This paper's result complements and differs from the existing theoretical literature. In Shavell (1992), strict liability and a carefully designed negligence rule are both perfect in inducing socially optimal behavior. This paper however shows that strict liability may induce too much or too little information acquisition. An important incentive present in this paper, but absent in the literature, is the seller's incentive to influence the public's learning. A seller can reduce the public's learning by dropping or fixing a bad product, which brings a bigger benefit under strict liability than under unknowable defense.

The supporters for strict liability also cite simplicity as a benefit. Keeton (1970) argues "One of the most important justifications for eliminating the negligence theory altogether is that doing so eliminates several complicated questions concerning whether the maker knew or should have known in the exercise of ordinary care about the extent of the dangers involved in the use of the particular product." Schwartz (1985) states that "The seemingly inevitable arbitrariness involved in distinguishing between knowable and unknowable risks has led some commentators to argue for the rule that imposes risks on firms regardless of what they knew at production time, and some courts have been persuaded." These forms of costs of implementing unknowable defense is practical and realistic in some cases, which should definitely be weighed against the benefit of unknowable defense identified in this paper. However, a process of finding out whether there was enough realized data to discern a pattern at a certain time in the past may not be costly. The analysis in Wagner (1997), which favors strict liability over unknowable defense, operates under the assumption that the seller actually has a way to know the risk pre-marketing but chooses

to not acquire that information because the action of not learning can fool the court into thinking that the information is not learnable. The court in this paper is assumed to be sophisticated enough to implement unknowable defense correctly.

This paper's inclusion of the information acquisition incentive *after* some products have been sold sets us apart from the literature on information acquisition on product safety in general. Poggi and Strulovici (2021) studies liability design that influences the incentive to acquire information acquisition only *prior* to marketing the product. Their model is rich in many other dimensions: the seller has private information and the liability can fine tune the damage amount according to part of the evidence obtained by the seller during information acquisition. Friehe and Schulte (2017) study information acquisition *prior* to marketing the product when firms require the approval of a regulator to launch a potentially risky product. Ben-Shahar (1998) examines the impact of strict liability (called "hindsight" in his paper) on various incentives, including the incentive to acquire information *prior* to marketing the product. Friedman and Wickelgren (2017) study the incentive to research drug's side effects prior to selling the drugs under different liability regimes with regard to brand name drug makers and generic drug makers.

Also recognizing how actual consumption data can generate information, Goeschl and Pfrommer (2015) shows that a regulator would want to encourage enough users to adopt a novel technology to resolve uncertainty about the technology's risk characteristic. They show that strict liability may discourage sellers' early adoption and information generation. No agent in their model makes an information acquisition choice. They only make a technology adoption choice which is influenced by the liability rule. We examine the information acquisition incentive in a product market where a seller has a long term interest and may drop a product and may also hide the findings of its information acquisition. Unlike Goeschl and Pfrommer (2015), which found that strict liability discourages information generation because it leads to less early adoption, we found that strict liability may induce too much information acquisition because of the seller's desire to reduce the public's learning. In Henry, Loseto and Ottaviani (2022), there are both ex ante and ex post learning and one of the regimes the paper looks at is strict liability. There the ex post learning takes a different form: information arrives exogenously and the seller only chooses whether or not to withdraw the product. There is no incentive for the seller to influence the public's discovery of harm.

The rest of the paper is organized as follows. Section 2 formally analyzes the model where a test becomes available to the seller only after some consumer adoption. I de-

rive the socially optimal outcome and the outcomes under strict liability and unknowable defense. Section 3 extends the model to also allow a weaker test to be available prior to any consumer adoption in order to study the state-of-the-art defense. I compare the socially optimal outcome, the outcome under strict liability, and that under unknowable defense. I also provide a numeric example to illustrate the case where strict liability under-incentivizes pre-marketing information acquisition because of the existence of the chance to do post-marketing information acquisition. Section 4 discusses alternative interpretations of the state-of-the-art defense and a "should-have-known" formulation of liability. Section 5 concludes.

2. Only post-marketing test is available

2.1. Setup

The court commits to a regime of liability: strict liability or unknowable defense.

There are two possible states about a product: the product is either harmful or not. The common prior about the product is that it is harmful with probability $\lambda \in (0,1)$. If the product is not harmful, it causes no harm. The value of the product to any buyer without the harm is $v > 0$.[4] If the product is harmful, aside from the positive value v, it causes a harm $H > v$ to a buyer. Normalize the production cost to zero.

At the beginning (time 0), there is no way to test. The technology is just not there or the data is not available before any consumer adoption. This assumption will be relaxed in Section 3. A monopoly seller decides whether or not to market the product. If the seller markets it, the seller offers a take-it-or-leave-it price p_1 to Q_1 amount of early buyers that would potentially buy. If these early buyers do not buy, they will wait in the market, that is, the timing stays at the beginning.

If these early buyers buy, they leave the market and at this time (time 1) a test becomes available privately to the seller. One can think of Q_1 is the amount of consumer data needed for the seller to be able to research on the harm. The seller chooses whether or not to run this test covertly which costs $c \geq 0$ to the seller. The test, if run, can discover evidence of harm if harm exists. Assume that the lack of evidence of harm is only verifiable in court, so it is not possible to disclose to the consumers that no harm is found. If the test shows the

[4] This assumption can be relaxed with no impact on the analysis: the early buyers and late buyers, defined later, can have different values.

existence of harm privately to the seller, the seller can choose whether or not to pay a fixed cost k to fix the harmful product and the action of fixing the product is not observable.[5] For example, a content-regulating mechanism can be added to Facebook to remove harmful contents or a producer can stop buying ingredients from contaminated sources or find substitutes for certain ingredients. Then, the seller makes a take-it-or-leave-it price offer p_2 to Q_2 amount of late consumers that would potentially buy. If the late buyers do not buy, they leave the market. There are no other potential buyers. Assume $k \in [0, \nu Q_2)$, so the product is worth fixing if found to be harmful. Assume the late buyers are important enough: $\nu Q_2 > \lambda H(Q_1 + Q_2)$, so that an uninformed seller would not want to drop the product before selling to the late buyers even if that triggers full liability for all buyers. The buyers know whether they are early or late.

If both the early and late buyers are harmed by the product, eventually over time, the public will learn the harm for sure. However, if only the early buyers are harmed by the product, the public has chance $1-\pi$ of eventually learning the harm on the early buyers. Assume $\pi \in (0,1)$. That is, more data of the harm in the market place increases the public's ability to learn by π.[6] The parameter π is the amount of influence that the seller has on the public's learning: if the seller let the harm accumulate in the market, the public will learn for sure, but if the seller can somehow reduce the data in the market to only the early buyers, the public will only learn with probability $1-\pi$. Therefore, π is the amount of the influence the seller has on the public's learning.

When the public learns that the product is harmful, liability for all realized harm is imposed on the seller under strict liability. However, under unknowable defense, immunity is granted for harm (and only harm) caused to the early buyers based on the fact that the seller had no way of knowing the harm prior to selling to the early buyers.

The parameter Q_1 captures how much product adoption is needed to allow the seller to learn at a cost. The parameters Q_2 and π capture the difference in the ability to learn

[5]For simplicity, I do not allow the seller to disclose a product is fixed. This can be endogenized. Under both liability regimes, the late buyers are sure of being compensated for the harm so announcing the product as fixed will not increase the willingness-to-pay of the late buyers but will alert the public to the harm on the early buyers, which is bad for the seller under strict liability.

[6]The assumption $\pi > 0$ also assumed that the public cannot eventually catch up on learning based on only Q_1 amount of hidden harm. One can make the counter-argument that the passing of time will help the public to learn as technology advances and other products who share similar features enter the market. However, on the other hand, the passing of time would also reduce the memory and evidence on whether consumption actually occurred or whether certain ingredients in the production were contaminated, and the potential defendant may not even continue to exist.

about the product between the seller and the public. Assume no discounting and all are risk neutral. Liability imposed on the seller results in compensation to the buyers with no waste. Not buying gives buyers payoff 0. The buyers maximize value from the product minus any price paid and plus any compensation from court. The seller maximizes profit minus any costs of information acquisition, harm mitigation, and legal liability. Assume the tie-breaking rules that buyers buy when indifferent, the seller markets the product when indifferent, and the seller acquires information when indifferent. The equilibrium concept is Perfect Bayesian Nash Equilibria.[7]

2.2. Socially optimal seller behavior

First, we look for the seller's behavior that would maximize the social welfare.

Litigation transfers and product prices do not matter for social welfare because they are transfers within the society. If the seller's test at time 1 reveals that the product is harmful, it is socially optimal for the seller to fix the product and continue selling it because $k < HQ_2$. If the test shows no harm or if no test is run, it is socially optimal to continue selling the product because $v > \lambda H$. The subgame social welfare of running the seller's private test is:

$$\lambda(vQ_2 - k) + (1-\lambda)vQ_2 - c = vQ_2 - c - \lambda k > 0$$

The subgame social welfare of not running the seller's private test is:

$$\lambda(v-H)Q_2 + (1-\lambda)vQ_2 = (v-\lambda H)Q_2 > 0$$

Taking the difference shows that running the test is socially optimal if

$$c \leq \lambda(HQ_2 - k) \tag{1}$$

That is, testing at time 1 is socially optimal if the expected net gain of fixing the harm for late buyers exceeds the information cost. Given these socially optimal decisions at time 1, the social value of selling the product at the very beginning is:

$$(v-\lambda H)Q_1 + \max\{vQ_2 - c - \lambda k, (v-\lambda H)Q_2\} > 0$$

[7] Note that the late buyers are always compensated for any harm under both strict liability and unknowable defense, so their beliefs about the seller's hidden past actions and the hidden chance of product harm is irrelevant to their purchase decision. When making price offer p_1, the seller has no private information relative to the early buyers. Therefore, we do not introduce notations of the buyers' beliefs here.

Not marketing the product has a social value of 0. Therefore, the socially optimal outcome is that the seller sells to early buyers and then the seller acquires information if and only if $c \leq \lambda(HQ_2 - k)$. The *socially optimal social welfare* is then:

$$U_{SO} = \begin{cases} (v - \lambda H)Q_1 + vQ_2 - \lambda k - c & \text{if } c \leq \lambda(HQ_2 - k) \\ (v - \lambda H)(Q_1 + Q_2) & \text{if } c > \lambda(HQ_2 - k) \end{cases}$$

Note that, without liability, because the action of fixing the product and testing the product are covert, the seller will have no incentive to do either. Even if we allow voluntary disclosure, any statements announcing the product fixed or the testing not discovering harm are cheap talk.

2.3. Strict liability

In this section, we study the seller's incentive under strict liability. First notice that strict liability and unknowable defense both hold the seller liable for any harm done to the late buyers. Moreover, if any harm is done to the late buyers, which can only happen when late buyers purchase, the public is sure to find out because of the abundance of data available to the public. Therefore, regardless of the regimes of liability, the late buyers are guaranteed the compensation for any harm caused to them. Therefore, the late buyers are always willing to pay up to v for the product. That is, they are charged v by the seller in any equilibrium. We start the analysis with the subgame after the test of time 1 is run privately by the seller.

If the test result indicates no harm, the seller charges the willingness-to-pay v from the late buyers, resulting in a payoff vQ_2.

If the test result uncovers the harm, the seller can fix the product, which gives the seller a subgame payoff $vQ_2 - k - (1 - \pi)HQ_1$. Here, the late buyers will not be harmed and the public may go after the harm on the early buyers with probability $1 - \pi$. The seller can also choose to not fix the product and continue selling the harmful product as it is. Then the public will learn of the harm for sure in the future and the seller's payoff is $vQ_2 - HQ_2 - HQ_1$, reflecting the liability on all buyers, which is worse than fixing the product.

Then, taking into account the possibilities of good news as well as bad news from the

test, the seller's subgame expected payoff of testing is:

$$\nu Q_2 - \lambda k - \lambda(1-\pi)HQ_1 - c$$

If the seller does not test at time 1 and sell to late buyers while being uninformed, the seller's future payoff is $\nu Q_2 - \lambda H(Q_1+Q_2)$. Taking the difference shows that the seller runs the test if and only if:

$$c \leq \lambda(\pi HQ_1 + HQ_2 - k) \tag{2}$$

Comparing this with the socially optimal information acquisition incentive (1) shows that the seller's incentive to acquire information is stronger than the socially optimal incentive.

Proposition 1. *The incentive of the seller to test the product at time 1 is too high under strict liability.*

The private incentive to test is too high because the test allows the seller to (partially) prevent the public's learning of any bad news and thus avoid the liability on past sales, while the seller's ignorance will give the public the chance to learn because of the continued accumulation of harm in the market.

Next, we go backward to time 0 to characterize the seller's equilibrium behavior and payoffs for the whole game. Note that whenever early buyers do not expect the seller to compensate them for any potential harm, the price they pay gets the seller to internalize this expected harm.

Case 1. $c \leq \lambda(\pi HQ_1 + HQ_2 - k)$. In this case, the seller will test and the early buyers know that the secret fixing of the product in the future means that they will get compensated for any harm incurred with only probability $1-\pi$, so they are at most willing to pay $\nu - \lambda\pi H$. Then, the seller's payoff of selling at time 0 is $(\nu - \lambda H)Q_1 + \nu Q_2 - \lambda k - c$. Recall that $\nu Q_2 > HQ_1 + HQ_2 > \lambda k - c$, so the payoff of selling at time 0 is strictly greater than 0, so the seller sells.

Case 2. $c > \lambda(\pi HQ_1 + HQ_2 - k)$. In this case, the seller will not test and will sell to late buyers while being uninformed. The accumulation of data means that the early buyers will be compensated for any harm for sure, which implies that the early buyers are willing to pay only ν for the product. Then, the seller's payoff of selling at time 0 the product is $(\nu - \lambda H)(Q_1 + Q_2) > 0$, so the seller sells at time 0. To summarize, the ex ante payoff of

the seller under strict liability is:

$$U_S = \begin{cases} (v-\lambda H)Q_1 + vQ_2 - \lambda k - c & \text{if } c \leq \lambda(\pi HQ_1 + HQ_2 - k) \\ (v-\lambda H)(Q_1+Q_2) & \text{if } c > \lambda(\pi HQ_1 + HQ_2 - k) \end{cases}$$

The seller's payoff is sometimes strictly less than the socially optimal social welfare. Comparing U_S to U_{SO} gives the following immediate result:

Proposition 2. *Under strict liability, the seller's payoff, U_S, is sometimes strictly below the socially optimal social welfare U_{SO}. This happens when the seller over-acquires information under the parameter case $\lambda(HQ_2 - k) < c \leq \lambda(\pi HQ_1 + HQ_2 - k)$.*

Here, some value is destroyed because the seller sometimes over-invests in acquiring information. This is in contrast with the previous literature (Shavell 1992) where strict liability perfectly aligns the seller's interest with the socially optimal. This is because here the early harm is not automatically discovered by the public. If a product is fixed early enough, there is not enough data to fully inform the public to link the injury of the early buyers to the product. Strict liability makes the seller care about this early liability and fixing the product obstructs the public's discovery of harm. Therefore the value of the information is higher to the seller than to the society by the time the early harm is already sunk.

This result is robust to some variations. First, if the value of the late buyers are not significant, the seller may want to drop the product, which, if done secretly, is another way of obstructing the learning of the public. Second, the basic model assumes the product is completely fixable. If the product is not fully fixable, the testing and fixing equilibrium outcome may get replaced by testing and dropping outcome if dropping can be done secretly, and if dropping the product may get noticed and trigger public investigation, the seller may prefer to partially fix. It remains true that under strict liability the seller has too much incentive to test.

2.4. Unknowable defense

In this section, I study the seller's incentive under unknowable defense. Here, because the product risk is impossible to learn at time 0, the court will never impose liability for harm caused to the early buyers under unknowable defense. Therefore, once the early buyers have purchased, the seller's information acquisition action does not affect the early buyers

any more, but only affects the late buyers. Any harm on the late buyers will eventually be fully exposed, so the seller is fully liable for such harm under unknowable defense. Therefore, the seller's decision on the test is always socially optimal. The following propositions then follow immediately without the need for any mathematical proofs.

Proposition 3. *The incentive of the seller to test the product at time 1 is socially optimal under unknowable defense.*

Notice that if the early buyers do not expect the seller to compensate them for any portion of the potential harm, they reduce the price they are willing to pay accordingly, so the price they pay gets the seller to internalize any uncompensated harm. This is true under strict liability or unknowable defense, so the casual conclusion that unknowable defense will leave the early buyers uncompensated is false: they are compensated through the lower price they pay for a riskier purchase. Moreover, all the late harm will be discovered by the public because of the accumulation of data in the market place, so the seller fully internalizes the late buyers' harm as well.

Proposition 4. *Under unknowable defense, the seller's behavior is socially optimal.*

That is, the payoff of the seller under unknowable defense is the same as the socially optimal social welfare:

$$U_U = U_{SO}.$$

In summary, Section 2 makes a case for granting unknowable defense to align the seller's incentive to acquire information and fix product defect when there is no early test available. The next section incorporate the availability of an early test.

3. Pre-marketing testing is also available

In Section 2, the assumption was that it is commonly known that there is no test available at all before selling to the early buyers, so it is clear that the harm on the early buyers was unknowable. In some cases, a more realistic assumption is that, prior to marketing the product, a statistical analysis could be done with limited data from lab experiments rather than actual consumption data, but such a test is less able to detect the harm than a test using the actual consumption of the product. In this section, I allow there to be an imperfect test (the state-of-the-art test) prior to the adoption of the early buyers at time 0.

The supporters of strict liability have argued that strict liability incentivizes the seller to acquire information relative to unknowable defense *prior* to marketing the product to the benefit of the society. One court opinion in an asbestos case granting the plaintiff's motion to strike the state-of-the-art defense argues that "(b)y imposing on manufacturers the costs of failure to discover hazards, we create an incentive for them to invest more actively in safety research."[8] By allowing there to be an early pre-marketing test, we can examine the logic of such an argument. In particular, is it true that when socially optimality calls for pre-marketing information acquisition, strict liability incentivises such pre-marketing testing relative to unknowable defense?

Assume the seller can run a pre-marketing test (the early test) prior to making the price offer to the early buyers at the same cost $c \geq 0$ as the post-marketing test. But this early test is less productive than the late test. If the product is harmful, the early test can reveal the harm with probability $\delta \in [0,1)$ and is mute with probability $1-\delta$. The early test is always mute if the product is not harmful.[9] After the early test (if run), the seller has the option to fix, sell or drop the product covertly. It is verifiable to the court whether the test result is mute or revealing.[10] Assume the tie-breaking rule that the seller will run the early test when indifferent. To focus on the trade-offs between the tests, in this section we will restrict the attention to the case of a highly-promising product as defined in Section 2: $vQ_2 \geq \lambda H(Q_1+Q_2)$ so dropping the product is a dominated action that we do not need to consider. The equilibrium concept requires the early buyers' beliefs about the covert actions of the seller to be correct.

3.1. Socially optimal seller behavior

Note that the early test is informative even when it is mute. The fact that it does not detect harm tells the seller something: the probability that the product is harmful is revised down to $\hat{\lambda} \equiv \lambda \frac{1-\delta}{1-\lambda\delta}$ which is less than the prior λ. The analysis of Section 2.2 has shown that the benefit of running the late test depends on the probability of harm before the test. Therefore, it is more beneficial run the late test if the early test is not run than if the early test is mute. The following proposition is proved in the appendix. It is intuitive that the socially optimal behavior calls for more testing when the test cost is lower (c is lower) and

[8] Beshada v. Johns-Manville Corp., 447 A.2d 539, 548 (N.J. 1982).

[9] The model in Section 2 is the special case of $\delta = 0$.

[10] Alternatively, if the private information is verifiable to the buyers (possibly through misrepresentation litigation), the information unraveling would imply that it is equivalent to mandatory disclosure, which is discussed in the subsection 4.3.

also relies more on the early test when the early test is stronger (δ is higher).

Proposition 5. *1. When $c \leq \min\{\frac{\lambda\delta}{1-\lambda\delta}HQ_1, \hat{\lambda}(HQ_2-k)\}$, the socially optimal behavior is to first run the pre-marketing test, and then run the post-marketing test if the pre-marketing test does not detect harm.*

2. *When $\delta < \frac{HQ_2-k}{H(Q_1+Q_2)-k}$ and $\frac{\lambda\delta}{1-\lambda\delta}HQ_1 < c \leq \lambda(HQ_2-k)$, the socially optimal behavior is to not run the pre-marketing test, but always run the post-marketing test.*
3. *When $\delta \geq \frac{HQ_2-k}{H(Q_1+Q_2)-k}$ and $\hat{\lambda}(HQ_2-k) < c \leq \lambda\delta(H(Q_1+Q_2)-k)$, the socially optimal behavior is to only run the pre-marketing test.*
4. *When $c > \max\{\lambda\delta(H(Q_1+Q_2)-k), \lambda(HQ_2-k)\}$, the socially optimal behavior is to not run any test.*

The left panel of Figure 1 shows the socially optimal outcomes in the space of the cost of tests (c) and the strength of the early test (δ). The outcome "both tests" means that the early test is always run, and the late test is run if the early test is mute.

3.2. Strict liability

Under strict liability, the seller's subgame incentives are governed by the analysis in Section 2.3. The following proposition is proved in the appendix.

Proposition 6. *Under strict liability,*

1. *When $c \leq \min\{\frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1, \hat{\lambda}(\pi HQ_1 + HQ_2 - k)\}$, the seller first runs the pre-marketing test, and then runs the post-marketing test if the pre-marketing test does not detect harm.*
2. *When $\delta < \frac{\pi HQ_1+HQ_2-k}{HQ_1+HQ_2-k}$ and $\frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1 < c \leq \lambda(\pi HQ_1+HQ_2-k)$, the seller does not run the pre-marketing test, but always runs the post-marketing test.*
3. *When $\delta \geq \frac{\pi HQ_1+HQ_2-k}{HQ_1+HQ_2-k}$ and $\hat{\lambda}(\pi HQ_1+HQ_2-k) < c \leq \lambda\delta(H(Q_1+Q_2)-k)$, the seller only runs the pre-marketing test.*
4. *When $c > \max\{\lambda\delta(H(Q_1+Q_2)-k), \lambda(\pi HQ_1+HQ_2-k)\}$, the seller does not run any test.*

Figure 1 compares the seller's behavior under strict liability to the socially optimal one. It shows a few deviations from the socially optimal behavior. First, when the late test is run, the impact of the early test to the seller is limited to only the part of the public's learning of the early harm that the late test will not be able to eliminate, while the social impact is the entire harm on the early buyers, so we see that sometimes the socially optimal "both

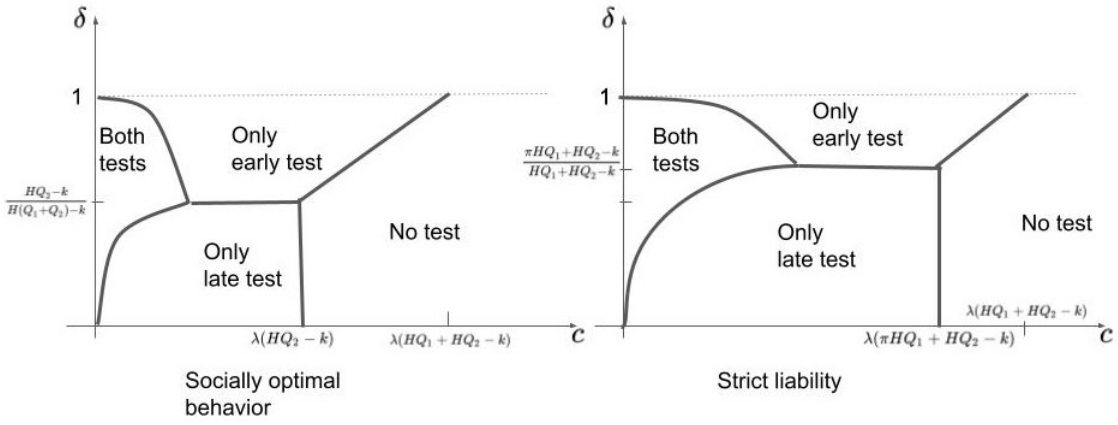


Figure 1. *The impact of strict liability when there is an early test*

tests" outcome is replaced by "only late test" in the left lower corner of the figure by strict liability, and sometimes the "only early test" outcome is replaced by the "only late test" in the middle of the figure, both of which represent a disincentive to run the early test. The seller only cares about the harm that will be found out, not the entire harm.

Second, given the analysis in Section 2, it is not surprising that strict liability creates too much incentive to run the late test in the subgame at time 1 because of the incentive to avoid past liability. This implies that sometimes the socially optimal "no test" outcome is replaced by "only late test" in the lower right region of the figure, and sometimes the socially optimal "only early test" is replaced by "both tests" in the upper left region of the figure, both of which represent an incentive to over-run the late test. To conclude the comparison with the socially optimal behavior, we have:

Corollary 1. *Under strict liability, the seller has too much incentive to run the post-marketing test and too little incentive to run the pre-marketing test.*

3.2.1. The state-of-the-art defense

The early test is a state-of-the-art test. If the seller has run it and has found no evidence of harm, this can convince the court that the harm was unknowable. If the seller has not run the early test, given that the early test is weak, it is debatable whether the harm was "knowable". Because the early test can only reveal the harm with probability δ, it is plausible that the defendant can convince the court that the knowledge was knowable with probability δ. That is, the court will grant the immunity with probability $1-\delta$ in

the event that the public learns the harm and the seller did not run the early test.[11] The following table of the probability of liability for the harm on the early buyers illustrates the difference between strict liability and the state-of-the-art defense:

	no early test	*early test revealing*	*early test mute*
strict liability	1	1	1
state-of-the-art	δ	1	0

The following proposition is proved in the appendix.

Proposition 7. *Under the state-of-the-art defense,*

1. *When* $c \leq \min\{\frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1, \hat{\lambda}(HQ_2-k)\}$*, the seller first runs the pre-marketing test, and then runs the post-marketing test if the pre-marketing test does not detect harm.*
2. *When* $\delta < \frac{HQ_2-k}{(1-\pi)HQ_1+HQ_2-k}$ *and* $\frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1 < c \leq \lambda(\delta\pi HQ_1+HQ_2-k)$*, the seller does not run the pre-marketing test, but always runs the post-marketing test.*
3. *When* $\delta \geq \frac{HQ_2-k}{(1-\pi)HQ_1+HQ_2-k}$ *and* $\hat{\lambda}(HQ_2-k) < c \leq \lambda\delta(H(Q_1+Q_2)-k)$*, the seller only runs the pre-marketing test.*
4. *When* $c > \max\{\lambda\delta(H(Q_1+Q_2)-k), \lambda(\delta\pi HQ_1+HQ_2-k)\}$*, the seller does not run any test.*

Figure 2 compares the seller's behavior under the state-of-the-art defense to the socially optimal one. It shows a few deviations from the socially optimal behavior. Just as in strict liability, the seller only cares about the harm when it is found out, not the entire harm, so we see that sometimes the socially optimal "both tests" outcome is replaced by "only late test" in the left lower corner of the figure, and sometimes the "only early test" outcome is replaced by the "only late test" in the middle of the figure, both of which represent a disincentive to run the early test. Also similar to pure strict liability, when the seller has not run the early state-of-the-art test, it faces strict liability with probability δ, so we see the over-running of the late test in lower right corner.

[11] One can imagine alternative definition of the state-of-the-art defense that imposes full liability (regardless of the strength of the early test) when the court finds out that the seller did not run an early test that was available. Section 4 analyzes such an alternative, called AD. A big problem with AD is that there is no continuity in the liability imposed. The lack of a test gives full immunity on the harm caused to the early buyers, but not running an useless test with $\delta = 0$ (if it is the best test available) rips off any immunity when the harm is found out by the public. This would cause a practical problem: the court will implement very different liability based on two practically non-distinguishable scenarios: the best test available is extremely weak or the test does not exist.

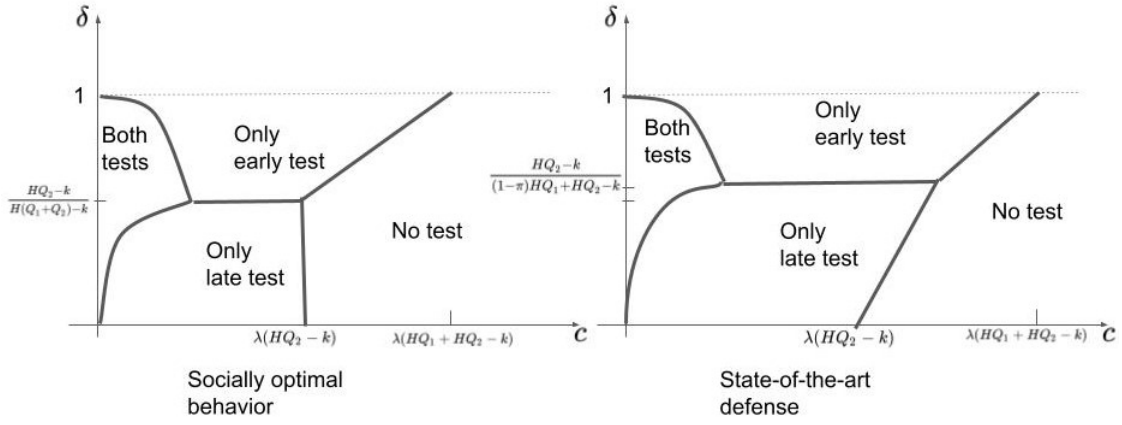


Figure 2. *The impact of the state-of-the-art defense when there is an early test*

Corollary 2. *Under the state-of-the-art defense, the seller has too much incentive to run the post-marketing test and too little incentive to run the pre-marketing test.*

3.3. Comparing strict liability and the state-of-the-art defense

While both regimes' departure from the socially optimal behavior is qualitatively similar, comparing the strict liability with the state-of-the-art defense in Figure 3, we see several differences. In each of the three regions they differ, the state-of-the-art defense achieves the socially optimal outcome but strict liability does not.

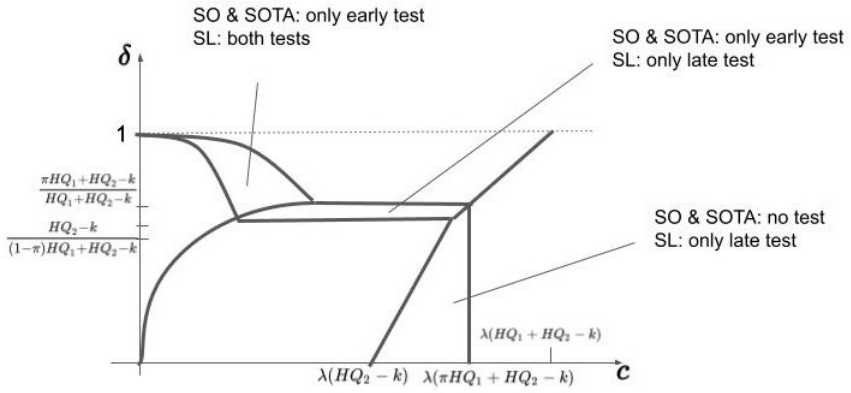


Figure 3. *Comparing strict liability with the state-of-the-art defense*

Proposition 8. *Granting a the state-of-the-art defense improves social welfare over strict liability.*

In other words, imposing liability on the unknowable risk is bad for the social welfare. It is true that unknowable defense discourages early testing relative to the socially optimal benchmark, however strict liability also discourages early testing, and on top of that strict liability over-encourages the late test. In other words, the criticism on unknowable defense in terms of a lack of incentive to acquire information prior to marketing the product is flawed by using the wrong benchmark.

There is one parameter case where the socially optimal outcome is "both tests" while the unknowable defense causes the early test to be dropped. In this parameter case, the strict liability also causes the early test to be dropped. In this parameter case, contrary to the common misconception, the incentive to run the early test is actually the same under strict liability and unknowable defense because the subgame incentive under both rules is to run the late test when the early test is mute. Strict liability imposes full liability on the early buyers. When the early test works (with probability δ), the harm and thus the liability is removed, so the benefit of the early test is proportional to δ at $\lambda\delta(1-\pi)HQ_1$. Under unknowable defense, if the seller does not run the early test, liability on the early buyers attaches with probability δ. If the seller runs the early test, there is no liability on early buyers at all. Therefore, the benefit of the early test is also proportional to δ at $\lambda\delta(1-\pi)HQ_1$.

There is another parameter case where the socially optimal outcome and the unknowable defense outcome are both "only early test", but strict liability outcome is "only late test". In a subset of this parameter case, the seller under strict liability wants to run the late test if the early test is mute, but the socially optimal outcome calls for no late test if the early test is mute. Foreseeing the seller's own strong incentive to run the late test in the future, the seller optimally chooses not to run the early test under strict liability: even by running the early test, the seller cannot fully eliminate the late test and its cost. In the other subset of this parameter case, the seller does not run the late test if the early test is mute under both regimes and under socially optimality. Here, the benefit of the an early test is actually higher under unknowable defense than under strict liability. An early test under unknowable defense eliminates the expected liability associated with the early buyers because the early test is considered the state of the art, but not so under strict liability.

This analysis complements the work of Ben-Shahar (1998) which, among many contributions, examined the impact of strict liability (termed "hindsight") on the incentive to acquire information prior to selling the product when information will be revealed later

for free. Ben-Shahar (1998) found strict liability will cause over-investment in information acquisition *prior* to marketing the product, I show that the opposite can be true. Ben-Shahar's definition of the state of the art regime is different from this paper's definition of the state-of-the-art defense. In Ben-Shahar, even though the risk was perfectly knowable pre-marketing at a cost, the state-of-the-art defense sets the standard of care based on an uninformed seller, which means the state-of-the-art rule in Ben-Shahar (1998) deems the risk as not knowable or in some sense as not worthy to be learned prior to marketing the product. This paper requires the court to have a correct understanding of how likely the risk was technologically knowable. The state-of-the-art rule in this paper does not attempt to infer whether the seller should have learned prior to marketing the product, but only focus on the strength of the early test if it is run.

3.4. Illustrative example

Assume that the ex ante probability that a product is harmful is 10%. If the harm exists, a state-of-the-art lab test (the early test) costing $20,000 may reveal this harm to the seller with probability 62.5%. That is, the test however may fail to detect the harm with probability 37.5%. There are potentially 1 million buyers for this product. The value of a non-harmful product is $0.50. A harm is equivalent to $1 per consumer. After selling to half of the all potential buyers (the early buyers), there are enough consumption data such that a $20,000 test (the late test) can now perfectly reveal any harm to the seller. Once the seller learns the harm, the seller can spend $200,000 to eliminate the harm by redesigning the product. The seller can choose to sell to the remaining buyers (the late buyers). If only half of the buyers bought the product, the public will eventually learn the harm with probability 50%. If all the potential buyers bought the product, the public will eventually learn the harm for sure. The two liability regimes are:

1. Strict liability: the seller has to compensate any buyers for the harm when the public learns of the harm.
2. State-of-the-art defense: when the public learns of the harm, the seller always has to compensate any late buyers for the harm. (1) If the seller has run the early test and the test did not reveal harm, i.e., is mute, the seller does not have to compensate the early buyers because the harm was "unknowable"; (2) if the seller has not run the early test, the harm would be deemed as knowable prior to selling to the early buyers with probability 62.5%, therefore, the seller has to compensate the early buyers with probability 62.5%.

For reasons that will be clear in the formal setup, the sellers always sells to all buyers even if the seller is uninformed and always fixes the harm upon discovering it.

Consider the scenario that the early test is run but the test is mute. This is nevertheless informative and informs the seller that the chance of the existence of the harm is lower at $10\%(1-62.5\%)/(1-10\%\times 62.5\%) = 4\%$. What decision on the late test is socially optimal? A late test can completely eliminate any harm to the late buyers after the seller pays to fix the product, so the expected benefit of the late test is $4\%(\$500,000-\$200,000) = \$12,000$. Because this amount is smaller than the cost of the late test \$20,000. The late test is not socially optimal to run after a mute early test.

How would a seller make this decision about the late test under strict liability? If the seller runs the late test, with probability 4%, the test will reveal a harm after which the seller can fix the product. Not fixing the product will surely alert the public of the harm and causes a liability on all past buyers at \$1 million. Fixing it will reduce the chance of the public becoming aware of the harm to 50%. The liability saving of \$500,000 clearly outweighs the fixing cost \$200,000, so the product will be fixed upon seller's discovery of harm. Fixing the product does not mean that the seller will completely have no liability, because there is 50% chance that the public will learn of the harm on early buyer's harm, which will be imposed on the seller under strict liability. Not running the late test means that the public will eventually find out the harm and impose an liability of \$1 million, so the expected benefit of the late test is $4\%\times(\$1,000,000-50\%\times\$500,000-\$200,000) = \$22,000$, which is greater than the \$20,000 cost of the late test, so the seller will run the late test even though it is socially optimal not to run it. Secretly learning through the late test allows the seller to obstruct the public's learning of the harm on the early buyers: the seller over-acquires post-marketing information. This misalignment of incentives on the late test will in turn affect the incentive on the early test.

Is it socially optimal to run the early test in this example? If the early test is not run, the belief about the chance of harm remains at 10%. Then, it is socially optimal to run the late test because the expected benefit of the late test in removing harm on the late buyers outweighs the test cost: $10\%(\$500,000-\$200,000) = \$30,000 > \$20,000$. By running the early test, conditional on a harmful product, the society incurs the cost of fixing the product if the early test is not mute and the cost of harm on all buyers if the early test is mute, which add up to $62.5\%\times\$200,000 + 37.5\%\times\$1,000,000 = \$500,000$. By not running the early test, conditional on a harmful product, the society incurs the cost of fixing the product (after the late test) and the harm on the early buyers, which add up to

$\$500,000 + \$200,000 = \$700,000$. Whether or not the early test is run, exactly one test cost will be incurred, so the test cost does not enter into the comparison. Therefore, it is socially optimal to run the early test.

How would a seller make this decision about the early test under strict liability? First, if the seller does not run the early test, the seller wants to run the late test because the expected benefit of the late test is $10\% \times (\$1,000,000 - 50\% \times \$500,000 - \$200,000) = \$55,000$ is greater than the $20,000 cost of the late test. That is, the seller runs the late test if the early test is not run and also if the early test is mute, so the seller is always informed prior to selling to the late buyers. By running the early test, the seller incurs early test cost right away, but avoids the test cost of the late test when the early test is not mute, which happens with probability $10\% \times 62.5\%$. By running the early test, the seller can fix the product before harm is done to the late buyers and thus reduce the public's learning and reduce the probability of being liable for the early buyers by 50% when the early test reveals the harm. The expected reduction of liability and expected reduction in testing cost associated with the late test add up to $10\% \times 62.5\% \times 50\% \times \$500,000 + 10\% \times 62.5\% \times \$20,000 = \$16,875$. This is smaller than the cost of the early test at $20,000. Therefore, the seller does not want to run the early test even though it is socially optimal to do so. The key difference stems from the fact that a seller under strict liability will overrun the late test. This implies that harm never occurs to the late buyers, which reduces the value of the early test. Moreover, running the early test leads to a duplication of the test costs because with a large probability the early test is mute and the seller will run a second test.

What would a seller do in this example under the state-of-the-art defense? First, without an early test, here the seller does not want to run the late test. A late test in this situation can obstruct the public's learning to avoid liability on the late buyers and to reduce the liability by 50% chance on the early buyers. This liability would stick with only probability 62.5% without any reduction because the harm on the early buyers is deemed unknowable with probability 37.5%. That is, the expected benefit of the late test in this situation is $10\% \times (62.5\% \times 50\% \times \$500,000 + \$500,000 - \$200,000) = \$45,625$, which is greater than the late test's test cost at $20,000. That is, the seller will run the late test if the early test is not run. Second, if the early test is mute, the seller is more optimistic about the product thinking the chance of harm is only 4% before the late test. Then, the expected benefit of the late test becomes $4\% \times (62.5\% \times 50\% \times \$500,000 + \$500,000 - \$200,000) = \$18,250$, which is less than the late test's test cost at $20,000. That is, the seller will not run

the late test if the early test is mute. Therefore, the late test decision is socially optimal even though the trade-offs are not exactly the same as the society's trade-offs. By running the early test, conditional on a harmful product, the seller incurs the cost of fixing the product if the early test is not mute and the cost of harm on only the late buyers if the early test is mute because of the state-of-the-art defense, which add up to $62.5\% \times \$200,000 + 37.5\% \times \$500,000 = \$312,500$. By not running the early test, conditional on a harmful product, the society incurs the liability on the early buyers with probability $62.5\% \times 50\%$ because of the chance of the court considering the risk to be unknowable and the public's ability to detect the early harm, and the cost of fixing the product (after the late test), which add up to $62.5\% \times 50\% \times \$500,000 + \$200,000 = \$356,250$. Whether or not the early test is run, exactly one test cost will be incurred, so the test cost does not enter into the comparison. Therefore, it is socially optimal to run the early test.

This example shows how imposing liability for unknowable risk can create inefficient information acquisition incentives.

4. Discussion

4.1. Alternative state-of-the-art formulation

In Section 3, the definition of the state-of-the-art defense based the liability on the strength of the early test. One can think of a definition that does not depend on the strength of the test: it imposes full liability if the early test is not run. We call this the AD definition. The following table of probability of being liable for the harm on the early buyers when the public discovers the harm contrasts AD with the definition of strict liability and the definition of the state-of-the-art defense in Section 3:

	no early test	*early test revealing*	*early test mute*
strict liability	1	1	1
state-of-the-art	δ	1	0
Alternative definition (AD)	1	1	0

While this AD definition seems intuitive, there is a big problem of the lack of continuity in the liability imposed. As explained when defining the state-of-the-art in the main model, such a AD definition will cause a practical problem: the court will implement very different liability based on two practically non-distinguishable scenarios: the best test available is extremely weak in detecting this particular harm or a test to detect this harm does not

exist. In this section, we nevertheless can examine the impact of a state-of-the-art defense using the AD definition. The proof of the following proposition is in the appendix.

Proposition 9. *Under AD,*

1. *When* $c \leq \min\{\frac{\lambda}{1-\lambda\delta}(1-\pi)HQ_1, \hat{\lambda}(HQ_2-k)\}$*, the seller first runs the pre-marketing test, and then runs the post-marketing test if the pre-marketing test does not detect harm.*
2. *When* $\delta < 1-\frac{(1-\pi)HQ_1}{HQ_2-k}$ *and* $\frac{\lambda}{1-\lambda\delta}(1-\pi)HQ_1 < c \leq \lambda(\pi HQ_1+HQ_2-k)$*, the seller does not run the pre-marketing test, but always runs the post-marketing test.*
3. *When* $\delta \geq 1-\frac{(1-\pi)HQ_1}{HQ_2-k}$ *and* $\hat{\lambda}(HQ_2-k) < c \leq \lambda HQ_1 + \lambda\delta(HQ_2-k)$*, the seller only runs the pre-marketing test.*
4. *When* $c > \max\{\lambda HQ_1 + \lambda\delta(HQ_2-k), \lambda(\pi HQ_1+HQ_2-k)\}$*, the seller does not run any test.*

Because the seller is rewarded with immunity by running the test regardless of the strength of the test, now the seller actually overruns the early test relative to the socially optimal outcome. As a result, the AD-style state-of-the-art does not always improve on the strict liability.

4.2. "Should-have-known" liability

The state-of-the-art rule requires the court to only ascertain the quality of the potential test prior to the marketing of the product. One can envision a more costly rule where the court finds out the optimal pre-marketing information acquisition decision, which involves knowledge of all the parameters. If the socially optimal decision is to run the early test, then the court imposes full liability for the harm done to the early buyers when the public becomes aware of the harm regardless of whether the seller has run a test and whether the test was revealing or mute, and if the socially optimal decision is to not run the early test, the seller is not liability for any harm done to the early buyers regardless of whether the seller has run a test and whether the test was revealing or mute. One can term this liability rule a "should-have-known" rule.[12]

First, observe that for parameter case where the socially optimal outcome is "both tests" or "only early test", this should-have-known liability is equivalent to strict liability,

[12]The language of "should have known" has been used in court, e.g. "[a] manufacturer of a product has a duty to warn foreseeable users of dangers in the use of that product of which he knows or should have known" (Mitchell v. Sky Climber Inc., 396 Mass. 629, 631, 1986).

which implies that it suffers from the same inefficiency of over-running the late test and under-running the early test and similarly under-performs the state-of-the-art defense.

Second, for parameter case where the socially optimal outcome is "only late test" or "no test", this should-have-known liability provides full immunity for the early harm. By unhinging the subgame decision on the late test from the sunk past harm, the late test decision is always socially optimal. The immunity for the early harm further favors not running the early test. Therefore, in this parameter case, the outcome is socially optimal.

Proposition 10. *When the socially optimal outcome is to run the early test, the should-have-known rule gives the same outcome as strict liability. When the socially optimal outcome is not to run the early test, the should-have-known rule gives the socially optimal outcome.*

Therefore, the should-have-known rule sometimes under-performs the state-of-the-art rule (just as strict liability does when under-running the early test), but sometimes outperforms the state-of-the-art rule when by discouraging over-running the late test.

4.3. Mandatory disclosure

One can think of different levels of mandatory disclosure. The most comprehensive level of mandatory disclosure covers all private information: whether a test is run, the outcome of the test, and whether the product is fixed. A more realistic and limited mandatory disclosure (implementable by providing strong financial incentive to whistle blow for example) is to only force the seller to disclose a bad test outcome.[13]

A limited mandatory disclosure is enough to remove the seller's incentive to manipulate the public's learning when deciding whether to run the late test under strict liability. To see that, denote the seller's assessment of the probability of harm be θ right before making the decision on the late test. Running the late test gives a seller payoff $-c-\theta HQ_1+vQ_2-\theta k$ because the seller has to disclose a bad test outcome and the public is fully alerted to the harm on the early buyers while the harm on the late buyers is avoided by paying the fixing cost. Not running the late test gives a seller payoff $-\theta HQ_1+(v-\theta H)Q_2$ because the accumulation of data will also eventually alert the public to the early harm. That is, regardless of whether the late test is run, the seller is responsible for the early harm under strict liability, so the potential existence of the past harm does not distort the seller's decision about the future. Moreover, the early buyers are fully compensated for any harm

[13] This is different from the post-sale duty to warn because once the seller fixes a product there is no duty to warn any more, but mandatory disclosure still requires disclosure of harm that existed in the past.

to them, so they are always willing to pay v, so the private information of the seller prior to selling to the early buyers do not matter. Therefore, the seller is efficiently under strict liability.

For the same reason as under strict liability, mandatory disclosure on the late test's bad outcome will also make the seller's decision on the late test socially optimal under the state-of-the-art defense. To see that, consider the two scenarios. First, after the early test turned out to be mute, the seller no longer has liability on any early harm. The seller's payoff of running the late test is $-c+vQ_2-\hat{\lambda}k$, while the seller's payoff of not running the late test is $(v-\hat{\lambda}H)Q_2$, so the decision is based on comparing c with $\hat{\lambda}(HQ_2-k)$. Second, after the early test is not run, the seller is liable for any early harm with probability δ. The seller's payoff of running the late test is $-c-\lambda\delta HQ_1+vQ_2-\lambda k$, while the seller's payoff of not running the late test is $-\lambda\delta HQ_1+(v-\lambda H)Q_2$, so the decision is based on comparing c with $\lambda(HQ_2-k)$. The following proposition's proof is available upon request:

Proposition 11. *Under mandatory disclosure on any harm uncovered from the late test, strict liability and the state-of-the-art defense are both socially optimal.*

4.4. A free signal after some consumption

A plausible situation is that the seller is initially unaware of the potential harm, until some other free signal alerts the seller to the potential harm. In a traditional Bayesian setup, this corresponds to a very low (but positive) prior λ and a free informative signal after the early consumption. In this situation, the socially optimal course of action is to not run the early test, and then to run the late test only if the signal realization is bad enough: there is a high enough number of accidents for example. Extending our analysis, it is easy to see that any retrospective liability regime (strict or state-of-the-art) can cause the late test to be over-run in the sense that the seller chooses to run the late test after a less negative signal realization.

5. Conclusion

This paper compares the seller's behaviors under strict liability and state-of-the-art defense to the socially optimal behavior and also with each other. The paper recognizes that product harm is not always discovered or correctly attributed and often the causality underlying an injury is only discovered when there has been a sufficient amount of accumulated harm in the market place, and the seller's actions of information acquisition pre-

and post-marketing has an influence on its own learning of the product risk as well as on the public's learning of the product risk.

This paper shows that under strict liability, there is too much incentive to acquire information post-marketing in order to obstruct the public's learning and avoid past liability. Foreseeing its own strong incentive to acquire information down the road, the seller rationally reduces information acquisition pre-marketing. Interestingly, even though the seller always has an incentive to under-invest in pre-marketing information acquisition relative to the socially optimal behavior because the seller only cares about the product harm if it is caught, the seller reduces such information acquisition more under strict liability than under the state-of-the-art defense because the action of running a state-of-the-art test increases the chance of the seller being granted immunity in the future should the harm be exposed to the public.

This paper provides a criticism of strict liability and a vindication of the state-of-the-art defense grounded in the strategic incentive of the seller in a dynamic setting.

References

Ben-Shahar, O., 1998. Should products liability be based on hindsight? JL Econ. & Org., 14, p.325.

Faure, M., Visscher, L. and Weber, F., 2016. Liability for Unknown Risks–A Law and Economics Perspective. Journal of European Tort Law, 7(2), pp.198-227.

Friedman, E. and Wickelgren, A.L., 2017. Who (if Anyone) Should be Liable for Injuries from Generic Drugs? The Journal of Law, Economics, and Organization, 33(3), pp.541-577.

Goeschl, T. and Pfrommer, T., 2015. Learning by Negligence-Torts, Experimentation, and the Value of Information (No. 598). Discussion Paper Series.

Henry, E., Loseto, M. and Ottaviani, M., 2022. Regulation with experimentation: Ex ante approval, ex post withdrawal, and liability. Management Science, 68(7), pp.5330-5347.

Kaplow, L. and Shavell, S., 1991. Private versus socially optimal provision of ex ante legal advice (No. w3868). National Bureau of Economic Research.

Keeton, W.P., 1970. Products Liability–Inadequacy of Information. Tex. L. Rev., 48, p.398.

Poggi, F. and Strulovici, B., 2021. Liability Design with Information Acquisition. Northwestern University working paper.

Polinsky, A.M. and Shavell, S., 2012. Mandatory versus voluntary disclosure of product risks. The Journal of Law, Economics, & Organization, 28(2), pp.360-379.

Schwartz, A., 1985. Products liability, corporate structure, and bankruptcy: toxic substances and the remote risk relationship. The Journal of Legal Studies, 14(3), pp.689-736.

Shavell, S., 1992. Liability and the incentive to obtain information about risk. The Journal of Legal Studies, 21(2), pp.259-270.

Spier, K.E., 2011. Product safety, buybacks, and the post-sale duty to warn. The Journal of Law, Economics, & Organization, 27(3), pp.515-539.

Wade, J.W., 1983. On the Effect in Product Liability of Knowledge Unavailable Prior to Marketing. NYUL rev., 58, p.734.

Wagner, W.E., 1996. Choosing ignorance in the manufacture of toxic products. Cornell L. Rev., 82, p.773.

Wells, G., Horwitz, J. and Seetharaman, D., 2021. Facebook knows Instagram is toxic for teen girls, company documents show. The Wall Street Journal.

Appendix

Proof of Proposition 5.

Proof. The benefit of the strong late test after early buyers have already purchased and no prior test has been run is the expected value of the test being able to remove the harm for the future late buyers: $\lambda(HQ_2 - k)$. The benefit of an early test alone (without any late test) is the expected value of the test being able to remove the harm for all future buyers: $\lambda\delta(H(Q_1 + Q_2) - k)$. A mute outcome of the early test suggests that (through Bayesian updating) that the probability of harm is lower than the prior at $\hat{\lambda} = \lambda\frac{1-\delta}{1-\lambda\delta} < \lambda$. Based on the analysis in Section 2.2, if the early test is not run, the late test should be run if and only if $c \geq \lambda(HQ_2 - k)$. A similar logic implies that, if the early test is run but mute, the late test should be run if and only if $c \geq \hat{\lambda}(HQ_2 - k)$. It is clear that if the early test detects the harm, the socially optimal course of action is fixing the product with no more test.

Case 1. $c > \lambda(HQ_2 - k)$. It is socially optimal to *not* run the late test regardless of whether the early test is mute or not run. The cost of the early test is the expected cost of fixing the harm plus the test cost itself $\lambda\delta k + c$ and the benefit of the expected harm avoided on all future buyers $\lambda\delta H(Q_1 + Q_2)$. Therefore, the socially optimal outcome is "no test and sell" for $c > \lambda\delta(H(Q_1 + Q_2) - k)$ and "only early test" otherwise.

Case 2. $c \leq \hat{\lambda}(HQ_2 - k)$. It is socially optimal to run the late test whether the early test is mute or not run. The cost of the early test is the testing cost itself c. The benefit of the early test are twofold: there is the expected harm avoided on the early buyers and there is also the benefit of potentially saving the cost of the late test, which add up to $\lambda\delta HQ_1 + \lambda\delta c$. Note that because the late test will be run if the early test does not uncover harm, whether or not the early is run has no impact on the late buyers: they are always saved from any harm. Comparing c and $\lambda\delta HQ_1 + \lambda\delta c$ shows that the socially optimal outcome is "only late test" for $c > \frac{\lambda\delta}{1-\lambda\delta}HQ_1$ and "both tests" (where the late test is run only if the early test is mute) otherwise.

Case 3. $\hat{\lambda}(HQ_2 - k) < c \leq \lambda(HQ_2 - k)$. It is socially optimal to run the late test if the early test is not run, but to *not* run the late test if the early test is mute. That is, exactly one test is always run. Running the early test means that the early buyers are expected to escape harm with some chance. However, this has to be weighed against the fact that the late test can always save the late buyers from the harm, but the early test (as a substitute of the late test) only saves the late buyers with probability δ when

the product is harmful. Therefore, it boils down to the comparison between $\lambda\delta HQ_1$ and $\lambda(1-\delta)(HQ_2-k)$. This implies that the socially optimal behavior is to not run a very weak early test ($\delta < \frac{HQ_2-k}{H(Q_1+Q_2)-k}$), which results in "only late test", but to run a moderately weak early test ($\delta \geq \frac{HQ_2-k}{H(Q_1+Q_2)-k}$), which results in "only early test". ■

Proof of Proposition 6.

Proof. Case 1. $c > \lambda(\pi HQ_1 + HQ_2 - k)$. In this case, the seller does not run the late test when the early test is not run and also when the early test is mute because $\hat{\lambda} < \lambda$. (1) Consider an equilibrium where the seller does not run the early test. Expecting no test is ever run, the early buyers pay a price of v. The seller's payoff of not running the early test is $(v-\lambda H)(Q_1+Q_2)$. Consider the deviation of running the early test. If the test uncovers the harm, the seller fixes it right away. If the test is mute, the seller runs no future test and has liability on all buyers. The payoff of running the early test is therefore $-c-\lambda\delta k+(v-\lambda(1-\delta)H)(Q_1+Q_2)$. This implies that the seller will deviate to running the early test if $c < \lambda\delta(H(Q_1+Q_2)-k)$. Therefore, if $\delta \leq \frac{\pi HQ_1+HQ_2-k}{HQ_1+HQ_2-k}$, then this parameter case is not possible, so the seller will not deviate to running the test. On the other hand, if $\delta > \frac{\pi HQ_1+HQ_2-k}{HQ_1+HQ_2-k}$, there is a non-empty parameter case where the seller deviates to running the early test: $\lambda(\pi HQ_1+HQ_2-k) < c < \lambda\delta(H(Q_1+Q_2)-k)$. (2) Consider an equilibrium where the seller runs the early test. The early buyers expect the future harm on the late buyers will help to fully reveal any harm to the public, so the early buyers are willing to pay v. The difference in the payoffs of the proposed equilibrium and the deviation payoff is the same as that reached under (1). This implies that when equilibrium of (1) does not exist, equilibrium of (2) exists.

Case 2. $c \leq \hat{\lambda}(\pi HQ_1 + HQ_2 - k)$. In this case, the seller runs the late test whether the early test is mute or not run. (1) Consider an equilibrium where the seller does not run the early test. The seller will run the late test and fix the product so the public will learn the harm on the early buyers with probability $1-\pi$. Expecting that, the early buyers pay up to $v-\lambda\pi H$. The seller's expected liability on the early buyers is $\lambda(1-\pi)HQ_1$. The seller's payoff of not running the early test (and then running the late test) is $(v-\lambda H)Q_1+vQ_2-\lambda k-c$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller runs the late test, and has $1-\pi$ chance of liability on any harm on the early buyers. The payoff of running the early test is therefore $-c+\lambda\delta[-k+(v-\lambda\pi H)Q_1+vQ_2]+(1-\lambda\delta)[(v-\lambda\pi H)Q_1+vQ_2-\hat{\lambda}k-c-(1-\pi)\hat{\lambda}HQ_1] = -(2-\lambda\delta)c-\lambda k+(v-\lambda H)Q_1+$

$\nu Q_2 + \lambda\delta(1-\pi)HQ_1$. The early testing brings a benefit to the seller by eliminating the harm on the early buyers with some probability, which would otherwise create a liability. This benefit amounts to $\lambda\delta(1-\pi)HQ_1$, which is weighed against the extra testing cost $(1-\lambda\delta)c$ caused by the early testing. This equilibrium of not running the early test (so only the late test) exists under this parameter case if $c > \frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1$. (2) It follows that an equilibrium of running the early test (and then running the late test if the early test is mute) exists if $c \leq \frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1$.

Case 3. $\hat{\lambda}(\pi HQ_1 + HQ_2 - k) < c \leq \lambda(\pi HQ_1 + HQ_2 - k)$. The seller runs the late test if the early test is not run, but the seller does not run the late test if the early test is mute. (1) Consider an equilibrium where the seller does not run the early test. The seller will run the late test and fix the product, so the public will learn the harm on the early buyers with probability $1-\pi$ if harm exists. That is, the expected liability on the seller is $\lambda(1-\pi)HQ_1$. Expecting that, the early buyers pay $\nu - \lambda\pi HQ_1$ and the seller's payoff is $(\nu-\lambda H)Q_1 + \nu Q_2 - \lambda k - c$. Consider the deviation of running the early test. If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller does not run any more test and has full liability on all buyers. The payoff of running the early test is therefore $-c+\lambda\delta[-k+(\nu-\lambda\pi H)Q_1+\nu Q_2]+(1-\lambda\delta)[(\nu-\lambda\pi H)Q_1+(\nu-\hat{\lambda}H)Q_2-\hat{\lambda}HQ_1]$. Running the early test causes an expected loss on the late buyers because should the test be mute, the seller gives up learning and thus the chance to fix the product for the late buyers: $\lambda(1-\delta)(HQ_2-k)$. In the mean time, the early test may or may not reduce the liability on the early buyers. Early-buyers-associated liability is $\lambda(1-\pi)HQ_1$ under no early test, but $\lambda(1-\delta)HQ_1$ under the early test. Since in either case, exactly one test is run, the test costs cancel out. Taking the difference of the payoffs shows that the equilibrium with no early test (and then late test) exists under this parameter case if $\delta < \frac{\pi HQ_1+HQ_2-k}{HQ_1+HQ_2-k}$. (2) It follows that an equilibrium of running the early test (and then running the late test if the early test is mute) exists if $\delta \geq \frac{\pi HQ_1+HQ_2-k}{HQ_1+HQ_2-k}$.

■

Proof of Proposition 7.

Proof. The incentive to run the late test after the early test is mute depends on the comparison of c and $\hat{\lambda}(HQ_2-k)$ under full immunity for early harm, while the incentive to run the late test after no early test depends on the comparison of c and $\lambda(\delta\pi HQ_1 + HQ_2 - k)$ reflecting that with probability δ the early liability attaches. Based on the subgame incentive on whether to run the late test, there are three parameter cases.

Case 1. $c > \lambda(\delta\pi HQ_1 + HQ_2 - k)$. In this case, the seller does not run the late test when the early test is not run and also when the early test is mute because $\hat{\lambda} < \lambda$ and $\delta HQ_1 > 0$. The continued selling of the product will ensure that any harm is eventually revealed. (1) Consider an equilibrium where the seller does not run the early test. There is strict liability on the early buyers with probability δ when harm is discovered. Expecting that, the early buyers pay a price of $v - \lambda(1-\delta)H$. The seller's payoff of not running the early test (which means no test is ever run) is $(v - \lambda H)(Q_1 + Q_2)$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller runs no future test and has no liability on the early buyers. The payoff of running the early test is therefore $-c + \lambda\delta[-k + (v - \lambda(1-\delta)H)Q_1 + vQ_2] + (1-\lambda\delta)[(v - \lambda(1-\delta)H)Q_1 + (v - \hat{\lambda}H)Q_2] = -c - \lambda\delta k + (v - \lambda(1-\delta)H)(Q_1 + Q_2)$. This implies that the seller will deviate to running the early test if $c < \lambda\delta(H(Q_1+Q_2)-k)$. That is, under Case 1 an equilibrium with no early test exists if $c > \lambda\delta(H(Q_1 + Q_2) - k)$ (2) It follows that an equilibrium where the seller runs the early test exists if $c \leq \lambda\delta(H(Q_1 + Q_2) - k)$.

Case 2. $c \leq \hat{\lambda}(HQ_2 - k)$. In this case, the seller runs the late test whether the early test is mute or not run. (1) Consider an equilibrium where the seller does not run the early test. The seller will run the late test and fix the product so the public will learn the harm on the early buyers with probability $1 - \pi$. Expecting that, the early buyers pay up to $v - \lambda H + \lambda(1-\pi)\delta H$. The seller's payoff of not running the early test (and then running the late test) is $(v - \lambda H)Q_1 + vQ_2 - \lambda k - c$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller runs the late test, and immunity kicks in and there is no liability on the early buyers. The payoff of running the early test is therefore $-c + \lambda\delta[-k + (v - \lambda H + \lambda(1-\pi)\delta H)Q_1 + vQ_2] + (1-\lambda\delta)[(v - \lambda H + \lambda(1-\pi)\delta H)Q_1 + vQ_2 - \hat{\lambda}k - c] = -(2-\lambda\delta)c - \lambda k + (v - \lambda H + \lambda(1-\pi)\delta H)Q_1 + vQ_2$. The early testing brings a benefit to the seller by eliminating the harm on the early buyers, which would otherwise create an expected liability of $\lambda\delta(1-\pi)HQ_1$, which is weighed against the extra testing cost $(1-\lambda\delta)c$ caused by the early testing. This equilibrium of not running the early test (so only the late test) exists under this parameter case if $c > \frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1$. (2) It follows that an equilibrium of running the early test (and then running the late test if the early test is mute) exists if $c \leq \frac{\lambda\delta}{1-\lambda\delta}(1-\pi)HQ_1$.

Case 3. $\hat{\lambda}(HQ_2 - k) < c \leq \lambda(\delta\pi HQ_1 + HQ_2 - k)$. The seller runs the late test if the early test is not run, but the seller does not run the late test if the early test is mute. (1) Consider

an equilibrium where the seller does not run the early test. The seller will run the late test and fix the product so the public will learn the harm on the early buyers with probability $1-\pi$, which creates an expected liability of $\lambda(1-\pi)\delta HQ_1$ on the seller. Expecting that, the early buyers pay up to $v-\lambda H+\lambda(1-\pi)\delta H$. The seller's payoff of not running the early test (and then running the late test) is $(v-\lambda H)Q_1+vQ_2-\lambda k-c$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller does not run any more test and has no liability on the early buyers due to the state-of-the-art defense. The payoff of running the early test is therefore $-c+\lambda\delta[-k+(v-\lambda H+\lambda(1-\pi)\delta H)Q_1+vQ_2]+(1-\lambda\delta)[(v-\lambda H+\lambda(1-\pi)\delta H)Q_1+(v-\hat{\lambda}H)Q_2]$. Running the early test causes an expected loss on the late buyers because should the test be mute, the seller gives up learning and thus the chance to fix the product for the late buyers: $\lambda(1-\delta)(HQ_2-k)$. This should be weighed against the benefit for the seller in avoiding the liability on the early buyers: $\lambda\delta(1-\pi)HQ_1$. Since in either case, exactly one test is run, the test costs cancel out. This equilibrium of not running the early test (so only the late test) exists under this parameter case if $\delta > \frac{HQ_2-k}{(1-\pi)HQ_1+HQ_2-k}$. (2) It follows that an equilibrium of running the early test (and then running the late test if the early test is mute) exists if $\delta \leq \frac{HQ_2-k}{(1-\pi)HQ_1+HQ_2-k}$.

■

Proof for Proposition 10.

Proof. We only need to check the parameter case where the socially optimal outcome contains no early test.

First, consider the parameter case $\delta < \frac{HQ_2-k}{H(Q_1+Q_2)-k}$ and $\hat{\lambda}(HQ_2-k) < c \leq \lambda(HQ_2-k)$. The society's ex ante expected cost of running the early test is $\lambda\delta k+\lambda(1-\delta)H(Q_1+Q_2)$, reflecting that if the early test is mute, the late test will not be run. Because should-have-known rule removes the liability on the early harm, the seller's ex ante expected cost of running the early test is $\lambda\delta k+\lambda(1-\delta)HQ_2$. That is, the should-have-known rule removes the liability on the early harm when the early test is mute, amounts to $\lambda(1-\delta)HQ_1$. The society's ex ante expected cost of not running the early test is $\lambda(HQ_1+k)$ reflecting the early harm and the cost of fixing after the late test. Because should-have-known rule removes the liability on the early harm, the seller's ex ante expected cost of not running the early test is λk. That is, the should-have-known rule removes the liability on the early harm for sure, which is equal to λHQ_1. Therefore, the seller further favors not running the early test under should-have-known rule, so the outcome is "only late test".

Similar comparison between $\lambda(1-\delta)HQ_1$ and λHQ_1 implies that the seller chooses not to run the early test in the parameter case where $c > \max\{\lambda\delta(H(Q_1+Q_2)-k), \lambda(HQ_2-k)\}$. Therefore, the outcome is "no test". ■

Proof of Proposition 9.

Proof. The incentive to run the late test after the early test is mute depends on the comparison of c and $\hat{\lambda}(HQ_2-k)$ under full immunity for early harm, while the incentive to run the late test after no early test depends on the comparison of c and $\lambda(\pi HQ_1 + HQ_2 - k)$ just as under strict liability. Based on the subgame incentive on whether to run the late test, there are three parameter cases.

Case 1. $c > \lambda(\pi HQ_1 + HQ_2 - k)$. In this case, the seller does not run the late test when the early test is not run and also when the early test is mute because $\hat{\lambda} < \lambda$ and $\delta HQ_1 > 0$. The continued selling of the product will ensure that any harm is eventually revealed. (1) Consider an equilibrium where the seller does not run the early test. There is still full liability on the early buyers when harm is discovered. Expecting that, the early buyers pay a price of v. The seller's payoff of not running the early test (which means no test is ever run) is $(v-\lambda H)(Q_1+Q_2)$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller runs no future test and has no liability on the early buyers. The payoff of running the early test is therefore $-c+\lambda\delta[-k+vQ_1+vQ_2]+(1-\lambda\delta)[vQ_1+(v-\hat{\lambda}H)Q_2] = -c-\lambda\delta k+vQ_1+(v-\lambda(1-\delta)H)Q_2$. This implies that the seller will deviate to running the early test if $c < \lambda HQ_1 + \lambda\delta(HQ_2-k)$. That is, under Case 1 an equilibrium with no early test exists if $c > \lambda HQ_1 + \lambda\delta(HQ_2-k)$ (2) It follows that an equilibrium where the seller runs the early test exists if $c \leq \lambda HQ_1 + \lambda\delta(HQ_2-k)$. Compared to strict liability, the seller has more incentive to run the early test here because running the early test removes liability on the early buyers whether the test is revealing or mute.

Case 2. $c \leq \hat{\lambda}(HQ_2-k)$. In this case, the seller runs the late test whether the early test is mute or not run. (1) Consider an equilibrium where the seller does not run the early test. The seller will run the late test and fix the product so the public will learn the harm on the early buyers with probability $1-\pi$. Expecting that, the early buyers pay up to $v-\lambda H+\lambda(1-\pi)H = v-\lambda\pi H$. The seller's payoff of not running the early test (and then running the late test) is $(v-\lambda H)Q_1+vQ_2-\lambda k-c$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller

fixes the product right away. If the test is mute, the seller runs the late test, and immunity kicks in and there is no liability on the early buyers. The payoff of running the early test is therefore $-c+\lambda\delta[-k+(v-\lambda\pi H)Q_1+vQ_2]+(1-\lambda\delta)[(v-\lambda\pi H)Q_1+vQ_2-\hat{\lambda}k-c]=-(2-\lambda\delta)c-\lambda k+(v-\lambda\pi H)Q_1+vQ_2$. The early testing brings a benefit to the seller by eliminating the harm on the early buyers, which would otherwise create an expected liability of $\lambda(1-\pi)HQ_1$, which is weighed against the extra testing cost $(1-\lambda\delta)c$ caused by the early testing. This equilibrium of not running the early test (so only the late test) exists under this parameter case if $c>\frac{\lambda}{1-\lambda\delta}(1-\pi)HQ_1$. (2) It follows that an equilibrium of running the early test (and then running the late test if the early test is mute) exists if $c\leq\frac{\lambda}{1-\lambda\delta}(1-\pi)HQ_1$. This also reflects a stronger incentive to run the early test than under strict liability.

Case 3. $\hat{\lambda}(HQ_2-k)<c\leq\lambda(\pi HQ_1+HQ_2-k)$. The seller runs the late test if the early test is not run, but the seller does not run the late test if the early test is mute. (1) Consider an equilibrium where the seller does not run the early test. The seller will run the late test and fix the product so the public will learn the harm on the early buyers with probability $1-\pi$, which creates an expected liability of $\lambda(1-\pi)HQ_1$ on the seller. Expecting that, the early buyers pay up to $v-\lambda H+\lambda(1-\pi)H=v-\lambda\pi H$. The seller's payoff of not running the early test (and then running the late test) is $(v-\lambda H)Q_1+vQ_2-\lambda k-c$. Consider the deviation of running the early test (unobservable to the early buyers). If the test uncovers the harm, the seller fixes the product right away. If the test is mute, the seller does not run any more test and has no liability on the early buyers due to the state-of-the-art defense. The payoff of running the early test is therefore $-c+\lambda\delta[-k+(v-\lambda\pi H)Q_1+vQ_2]+(1-\lambda\delta)[(v-\lambda\pi H)Q_1+(v-\hat{\lambda}H)Q_2]$. Running the early test causes an expected loss on the late buyers because should the test be mute, the seller gives up learning and thus the chance to fix the product for the late buyers: $\lambda(1-\delta)(HQ_2-k)$. This should be weighed against the benefit for the seller in avoiding the liability on the early buyers: $\lambda(1-\pi)HQ_1$. Since in either case, exactly one test is run, the test costs cancel out. This equilibrium of not running the early test (so only the late test) exists under this parameter case if $\delta>1-\frac{(1-\pi)HQ_1}{HQ_2-k}$. (2) It follows that an equilibrium of running the early test (and then running the late test if the early test is mute) exists if $\delta\leq 1-\frac{(1-\pi)HQ_1}{HQ_2-k}$.

■