University of Heidelberg

Department of Economics

Discussion Paper Series | No. 598

Learning by Negligence - Torts, Experimentation, and the Value of Information

Timo Goeschl and Tobias Pfrommer

July 2015

Learning by Negligence - Torts, Experimentation, and the Value of Information

Timo Goeschl*
Heidelberg University

Tobias Pfrommer†
Heidelberg University

July 31, 2015

How should tort law deal with agents that employ novel and imperfectly understood technologies that later turn out to involve harm? There is no agreement among different legal systems whether strict liability or negligence rules should govern these so-called 'development risks'. The law-and-economics literature, however, has predominantly favored strict liability. The present paper shows that the choice depends on the characterization of how society learns about technology risks. When experiential public data is an irreducible input into learning, theory justifies the use of specific negligence rules in order to govern development risks. We reconcile the existence of the negligence doctrine for development risks with the theoretical literature using a simple two-period unilateral care model. There, an optimally designed negligence rule can provide a better balancing of benefits, harm to third parties, costs of care effort, and the value of information from learning than strict liability. If feasible, the optimal negligence rule partitions the population of potential users into two groups. Only the high benefit group engages in the risky activity, subject to due care levels designed to deter the low benefit group.

*Email: goeschl@eco.uni-heidelberg.de. Postal address: Department of Economics, Bergheimer Str. 20, 69115 Heidelberg, Germany. Phone: +49 6221 548010, Fax: +49 6221 548020.

†Email: pfrommer@eco.uni-heidelberg.de. Postal address: Department of Economics, Bergheimer Str. 20, 69115 Heidelberg, Germany. Phone: +49 6221 548013, Fax: +49 6221 548020.

1 Introduction

The tension between the social gains from novel technologies and their potential risks is particularly palpable in the context of how tort law should deal with the fact that for such novel technologies, accident probabilities can, by definition, not be based on experiential data. Instead, the assessment of the risks emanating from their use can only rely on estimates derived from premarket testing and from reasoning by similarity. As a result, both product liability and environmental liability in the United States[1] and the European Union[2] raise questions about the extent to which an efficient system of torts should hold agents liable for employing novel technologies that later turn out to have harmed third parties.

Our paper shows that the choice between strict liability and negligence depends on the characterization of how society learns about technology risks. When experiential public data is an irreducible input into learning, theory justifies the use of specific negligence rules in order to govern development risks. We reconcile the existence of the negligence doctrine for development risks with the theoretical literature using a simple two-period unilateral care model. There, an optimally designed negligence rule can provide a better balancing of benefits, harm to third parties, costs of care effort, and the value of information from learning than strict liability. This balancing involves incentives for users to employ the technology who would refrain from adopting the technology under strict liability, leading to more public information about the technology's risks.

There is a small, but distinguished literature in law and economics that examines the efficiency of different liability regimes when information about the riskiness of an activity is incomplete. Shavell (1992) examines a setting in which a potential injurer does not know whether a certain activity is risky or not, but can acquire this information at a fixed cost, for example through additional laboratory tests. In such a setting, strict liability provides optimal incentives for both information acquisition and care, and so do a 'complete' negligence rule[3] and an 'optimal knowledge' negligence rule[4], but not alternative negligence rules of care levels that are optimal given the injurer's actual level of knowledge or given the true state of the world. In a similar vein, Kaplow and Shavell (1992) examine the acquisition of legal advice by agents imperfectly informed about the legal implication of a certain course of action that has not been embarked on yet. They

[1] See the Restatement (Third) of Torts: Product Liability (1998) and CERCLA (1980)

[2] See, in particular, the Development Risk Clause (the "DRC") provided in Articles 7 (e) and 15(b) of the Directive 85/374/EEC and Art. 8.4(b) of the Environmental Liability Directive 2004/35/EC, but also member state legislation such as the Environmental Liability Law in Germany which enshrines strict liability for development risks in Par. 1 UmwHG, Par. 89 WHG and USchadG (Schieber 1999; Teschabai-Oglu 2012).

[3] Under this rule, there is an optimal care level that is contingent on whether information acquisition was socially optimal or not. If the injurer operates below either the optimal care level or the optimal knowledge level, he is liable.

[4] Here, the negligence standard is based on the care level that is optimal under optimally acquired information

find that when legal advice helps the agent to learn the specific harm that its actions will impose on a victim, then strict liability provides optimal incentives for care and information acquisition, while a negligence rule will lead to excessive information acquisition. Ben Shahar (1998) examines producer liability for novel products and studies to what extent experiential data should be part of determining ex post whether a state-of-the-art product design was negligent when such data was not available to the manufacturer ex ante. Here, uncertainty regarding the risk properties of the novel technology can either be resolved through additional tests prior to marketing or managed through safety measures before or after marketing. The paper by Ben Shahar (1998) adds nuance to the negligence regimes studied by Shavell (1992) and Kaplow and Shavell (1992): It demonstrates that both hindsight- and state-of-the -art negligence rules induce either too much or too little care effort relative to the social optimum.[5] Finally, in a less technical mode, Dana (2010) develops a model of sequential research decisions to study the impact of liability rules on precaution by manufacturers of products that contain potentially harmful nanotechnology.[6]

The aim of the present paper is to further our understanding of the regulation of novel technologies through liability rules by extending the range of information environments used to study the question. The environment common in the existing literature is one in which the risk characteristics of the novel technology are learned through non-experiential private information acquisition. Here, 'non-experiential' refers to the common assumption that it is possible to eradicate any remaining uncertainties about the safety implications of the novel technology in the real world without ever experiencing its performance in a market setting. Instead, any relevant uncertainty can be resolved through additional laboratory tests before any harm has materialized (Shavell 1992, Ben Shahar 1998). 'Private' refers to the assumption that new information about the risk characteristics comes exclusively to the attention of the potential injurer. Experimentation with the novel technology generates no informational externalities.

Our paper, by contrast, presents an information environment in which the risk characteristics of the novel technology are learned through experiential public data. By experiential, we mean that in practice, even with considerable premarket testing, much of the information about novel technologies' risk characteristics still has to be acquired at the postmarket stage.[7] For many novel technologies, carrying out the potentially

[5] There is one exception to this summary statement: Negligence rules based on the state of the art provide the correct incentives for investments in safety devices prior to marketing.

[6] There is also an economics literature that considers the insurability of such risks when potential injurers are imperfectly informed, starting with Skogh (1998) and moving into issues of information acquisition and adverse selection in insurance (Crocker and Doherty 2000; Bajtelsmit and Thistle 2009).

[7] For example, in the pharmaceutical context, one study claims that more than half of drugs approved for release by the FDA are later shown to have serious adverse effects not detected during the three phases of pre-marketing clinical studies (Moore et al. 1998). As a result, there is considerable research effort on post-marketing surveillance of newly released drugs, with the results that a considerable number of products are subsequently withdrawn (Fontanarosa et al. 2004).

harmful activity in a market setting is an irreducible part of their exhaustive risk assessment. In contrast to laboratory tests, postmarket experiences with novel technologies typically generate public information: The presence or absence of harm is not private information for the manufacturer. It is observed by victims, third parties, interested scientists, and regulators. These informational externalities give rise to social returns to learning about the risk characteristics of the technology, impacting on the role of liability regimes in such settings: Different liability regimes will lead to differences in whether, and if yes by how many, the technology is used. This, in turn, leads to differences in the volume of experiential data that becomes available for future decisions on liability and care. Liability and learning are therefore linked, which matters for assessing the welfare impacts of different liability rules that determine not only the benefits, damages, and care costs associated with a novel technology, but also the value to society of learning more about the technology's risk characteristics in a market setting. This information value is at the heart of an economic argument for a possible role for a well-designed negligence rule that does better than strict liability in allocating the dynamic benefits and costs of novel technologies between society at large, users, and potential victims.

The main message of this paper obviously also relates to a recent literature on tort law and innovation. For example, Endres and Bertram (2006) and Endres and Friehe (2011) examine the impact of liability regimes for environmental damages on the rate and direction of technological change chosen by regulated firms regarding a care costs reducing technology, mainly finding strict liability to provide optimal incentives. Parchomovsky and Stein (2008) highlight the disincentive on innovation of the role of custom in determining negligence in court.

Furthermore, our findings are related to those of Schwartzstein and Shleifer (2013), who find a welfare-enhancing role for traditional regulation as an information device to complement tort litigation. Like in our setting, welfare is increased by encouraging socially desirable economic activity compared to strict liability. In contrast, this not achieved by a negligence rule alone, but through a combination of negligence and ex-ante regulation. Their assumption that firms can not appropriate full social returns is similar to the presence of information externalities in our paper, but they differ i) in their second fundamental assumption that courts (and regulators) make errors and ii) in regulators receiving an ex-ante signal regarding the safety of firm which is not experience-based.

To demonstrate that the consideration of an information environment focusing on experiential aspects is consequential, we develop a simple two-period model of unilateral care in which a population of possible users has heterogeneous private benefits. The social planner is charged with choosing, for each of the two periods, a liability regime for users of a novel technology that has undergone premarket testing, but still involves potential risks for third parties. All parties are risk-neutral. It is well known that in period 2, the absence of any further learning possibility makes the social planner optimally choose strict liability (Shavell 1980).[8] The core results of the paper concern the question of which regime the social planner optimally chooses for period 1. Relevant considerations for period 1 are expected benefits, expected harm, and care costs associated with the novel technology being used at a certain volume, but also the possibility of generating public data through experience on the market. All agents share a common belief about whether the technology is risky or not; with probability p there is a risk, with probability $1-p$ there is no risk. In case there is a risk, harm depends on care x chosen by the user, giving rise to an expected harm function $h(x)$ with the usual properties. The prior belief p about the riskiness of the technology in period 1 is an estimate based on laboratory pre-market tests or reasoning by similarity. We assume additional lab-based test will not provide a better estimate (or are too costly). This captures the important feature of emerging technologies that for users and regulators, their risks characteristics are uncertain. For simplicity we assume that uncertainty about the risk is completely resolved[9] if enough users employ the novel technology in period 1 and thus the volume of users is greater than a threshold parameter v_{info}.[10]

In this setting, we first derive the pattern of socially optimal use of the new technology in period 1 to provide an understanding of what outcome a liability regime should ideally generate. We find that independent from parameter choices, optimal use in period 1 always implies a volume sufficiently large for resolving information if only the stakes involved are large enough. By "stakes" we mean maximum benefits of using the technology and potential harm done in case there is a risk: Model parameters can be changed such that maximum benefits and total accident costs are increased by the same factor ("the scale") but all other characteristics remain unchanged.

[8] Strict liability would also be the regime a social planner would commit to in period 1 if there was no option for review.

[9] This assumption is obviously unrealistic, but our results are not affected if we assume partial revelation of uncertainty in line with the refinement described in Shavell (1992).

[10] Two examples can be used to provide a concrete illustration of the setting. One example is an agricultural context in which each farmer owns a unit tract of land with specific soil characteristics that is planted in a single crop variety. The novel technology consists of genetic modification of germplasm such that the farmer can choose to plant genetically modified organisms (GMOs) as crops. The potential harm associated with planting GMOs as a crop is that genetic material can be dispersed to an adjacent non-GMO farm, where a gene flow event could occur (Bouchie 2002). If such a gene flow occurs, a farmer that does not use GM crops can no longer sell his crop as GMO-free and suffers a price penalty (Bullock and Desquilbet 2002). Effective prevention of dispersion can only be carried out by GMO farmers through the construction of natural barriers or buffer zones around their crop. These efforts constitute the care level in this model. All agents know the damages, but the assessment of the potential of a gene flow event in the field is only based on experiments under controlled conditions and therefore fraught with uncertainty (Faure and Wibisana 2010). Another example is the use of nanotechnology in personal care products. For different applications, this technology offers different benefits. The possible harm associated with nanotechnology are adverse health effects of these ultrasmall particles passing through protective barriers in the human body (Dowling et al. 2004). Producers can influence, at a cost, the characteristics of the nanoparticles and thus decrease harm. As in the case of pharmaceuticals, however, existing studies under controlled conditions provide limited information on the toxicology of absorbed nanoparticles (Dana 2010).

We then compare strict liability and negligence in their ability to implement socially optimal behavior. We find that for settings in which some users should engage in the technology for static reasons alone, strict liability should implement the uncertainty resolving solution for large stakes. Contrary, if the technology should not be used at all under static conditions, a negligence rule is optimal for implementing the uncertainty resolving solution for large stakes. In order to ensure negligence to indeed be welfare-enhancing, the regulator has to deviate from the first-best due care level in order to control activity levels. Therefore, negligence can not implement the first-best solution, but is in this case still superior to strict liability under which the value of information would be foregone. The second-best care level may be higher or lower than the first-best, but will eventually, as stakes get bigger, be more stringent than the first-best.

In the next section, we define the model characteristics. We then derive the social planner's optimal behavior in the two-period setting in section 3 if care and activity levels are under his immediate control. Section 4 analyzes the conditions under which the choice of a strict liability and of a negligence rule maximize social welfare provide the better approximation to the social planner optimum when agents determine care and activity levels. Section 5 concludes.

2 The Model

We consider a social planner (SP) who is concerned with liability regulation of a novel technology for two consecutive periods, period 1 and period 2. The unilateral care setting involves a population of potential injurers that are heterogeneous with respect to the benefits of employing the technology. Every agent can decide in each period whether to make use of the novel technology or not. The continuum of agents can be ordered according to their private benefit (possibly negative) from switching to the new technology. We identify potential users with their position in the continuum $[0, \infty)$. This results in an aggregate marginal benefits curve $b(u) = B - u$, where B is the private benefit of the initial agent and u is the marginal user of the new technology. The SP knows the marginal benefit curve, but does not know the private benefits of single users. All agents are risk-neutral subjective expected utility maximizers.

At the beginning of period 1, it is not known with certainty whether the novel technology poses harm to third parties or not. However, all agents share a common belief regarding the likelihood that the novel technology is risky: With probability p there is a risk, with probability $1 - p$ there is no risk. In case there is a risk, expected harm done to third parties depends on a user's care level. The expected harm function $h(x)$ associated with care level x is the same for all users and known to all agents. We employ the usual assumptions $h'(x) < 0$ and $h''(x) > 0$. Uncertainty about the riskiness of the technology is resolved for period 2, if enough users employ the technology in period 1 and thus, a sufficient amount of experiential information is available. The activity of all users combined amounts to the volume v of all users: $v = \int a(u)\mathrm{du}$. If this aggregate experiential information exceeds a threshold v_{info}, information about the risk is revealed. Like the probability of risk p, the information threshold v_{info} is known to all agents.

The SP can choose between strict liability and negligence in each period. If he opts for negligence, he sets due care not only dependent on the state of information, but also taking into consideration effects on the volume of users, hence on information generation.

In short the givens of our model are:

1. The marginal private benefits curve $b(u) = B - u$
2. The care-harm relationship $h(x)$ in case there is a risk
3. The ex-ante belief p that the novel technology involves a risk
4. The minimum volume of users v_{info} needed in period 1 in order to resolve uncertainty about the risk

All givens are known to all agents including the SP. The SP can choose between strict liability and a negligence rule. In the latter case he is not confined to use the statically optimal care level as due care but can freely choose which due care level to set.

We employ the following assumptions in order to exclude the trivial cases of nobody or always somebody using the technology in the social optimum, irrespective of the state of information:

Assumption 1. *If it is known that the technology is riskless, some users want to engage in the technology*

$$B > 0$$

Assumption 2. *If it is known that the technology is risky, nobody should engage in the activity from a social point of view:*

$$B - x_1^S - h(x_1^S) \leq 0$$

3 The Social Optimum

The socially optimal management of the novel technology in period 1 involves balancing benefits, expected harm, care costs, but also the value of information. We first solve the static problem in period 2 without possibility of experiential learning and then turn towards the dynamic problem. Optimal behavior obviously depends on the state of information. We will indicate variables refering to situations under uncertainty with p, proven riskiness with 1 and proven risklessness with 0.

3.1 The Static Benchmark

Optimal static care depends on the state of information. There are three different states of information: Uncertainty, certainty of riskiness and certainty of risklessness. Optimal care is identical for all potential users of the technology, since they do not differ with respect to care costs, the harm function or access to information. In case of prevailing uncertainty, total expected accident costs per user are

$$x + p \cdot h(x).$$

The socially optimal x_p^S which minimizes these costs fulfills

$$1 + p \cdot h'(x) = 0.$$

If uncertainty is resolved and the technology turns out to be riskless, $x_0^S = 0$ is obviously optimal. If it turns out to be risky the optimal care level x_1^S fulfills

$$1 + h'(x) = 0.$$

It clearly holds that

$$x_0^S < x_p^S < x_1^S.$$

The socially optimal static activity level of a potential user depends on the state of information as well as on his private benefits $b(u)$. Socially optimal activity levels are derived given optimal care level for each state of information and each user. Under uncertainty a user with private benefit $b(u)$ should engage in the activity if

$$b(u) - x_p^S - p \cdot h(x_p^S) = B - u - x_p^S - p \cdot h(x_p^S) \geq 0.$$

Therefore, the activity level of a specific user u should be

$$a_p^S(u) = \begin{cases} 1 & \text{if } u \leq B - x_p^S - p \cdot h(x_p^S) \\ 0 & \text{if } u > B - x_p^S - p \cdot h(x_p^S) \end{cases}$$

Due to assumption 1, the optimal activity level is zero for all potential users in case of revealed riskiness:

$$a_1^S(u) = 0.$$

In case of revealed risklessness users should engage if they have a positve private benefit from using the technology, since precaution is not needed in known absence of risk.

$$a_0^S(u) = \begin{cases} 1 & \text{if } u \leq B \\ 0 & \text{if } u > B \end{cases}$$

The optimal volume of users is derived by aggregating all users who would optimally use the novel technology. Since all potential users are ordered with respect to their private benefits, the optimal volume coincides with the last user who should use the technology. In case of resolved uncertainty, this means: The optimal volume of users from a static point of view is

$$v_0^S = 0,$$

if the technology turned out to be risky (assumption 1) and

$$v_0^S = B$$

if the technology turned out to be riskless (assumption 2). Under uncertainty the optimal volume depends on whether marginal social benefits $B - x_p^S - p \cdot h(x_p^S)$ of the initial user are positive or not. If they are, nobody should use the technology and the optimal volume is zero. If they are not, the first $B - x_p^S - p \cdot h(x_p^S)$ users should use the technology:

$$v_p^S = \begin{cases} B - x_p^S - p \cdot h(x_p^S) & \text{if } B - x_p^S - p \cdot h(x_p^S) \geq 0 \\ 0 & \text{if } B - x_p^S - p \cdot h(x_p^S) < 0. \end{cases}$$

We introduce the notation $b_I^S = B - x_p^S - p \cdot h(x_p^S)$ for the initial user's statically social benefits, rewriting $v_p^S = \max[0, b_I^S]$.

Since period 2 is the last period, learning in period 2 is not consequential. Therefore, the socially optimal outcome in period 2 is solely determined by static outcomes:

Proposition 1. *Socially optimal behavior in period 2 is characterized by optimal static behavior.*

1. *Optimal care is the same for all potential users. Optimal care under proven riskiness is higher than optimal care under uncertainty which in turn is higher than optimal care under proven risklessness:* $x_0^S < x_p^S < x_1^S$.

2. *There is a private benefit threshold which divides potential users into two brackets. For the bracket with the higher private benefits activity is is optimal, for the lower value bracket no activity is optimal. The higher benefit bracket is empty in case of proven riskiness and under uncertainty if* $b_I^S \leq 0$.

3. *The volume of users entirely consists of the high-benefit bracket: It is* $v_1^S = 0$ *under proven riskiness,* $v_0^S = B$ *under proven risklessness and* $v_p^S = \max[0, b_I^S]$ *under uncertainty.*

The division into a high-benefit bracket and low-benefit bracket where the former consists of users and the latter of non-users of the technology persists for both the dynamical social optimum and all behavior under actual regulation: For the SP it is never optimal let a user with lower private benefits use the technology instead of one with higher private benefits but otherwise identical characteristics. Since a regulator does only know the distribution of private benefits, any liability rule applies to all potential users in the same way. Therefore, a user with higher private benefits will always have an incentive to use the technology if a user with lower private benefits has. Hence, the volume of users is sufficient information for knowing which users engage in the activity and which do not.

3.2 Dynamic Social Optimum

The only reason to deviate from statically optimal behavior is the value of information. Since the choice of care does not influence information, dynamically optimal care levels equal statically optimal care levels. Regarding aggregate activity, we will see that there are only two different potential volumes of users which may be optimal: The statically optimal volume v_p^S and the minimum volume required to reveal information v_{info}. In order to decide which volume is consistent with socially optimal dynamic behavior in period 1, we first have to attach a value to the information potentially obtained and then have to calculate the costs in terms of static inefficiency to be incurred for obtaining the value of information.

3.2.1 The Value of Information

The value of information of knowing p's true value is the difference in welfare between socially optimal behavior under uncertainty and (probability-weighted) welfare under socially optimal behavior after uncertainty is resolved. Welfare under proven riskiness is zero, since optimal volume is zero. Under proven risklessness optimal volume is B and social marginal benefits are $B-u$, therefore welfare is $\frac{1}{2}B^2$ in the social optimum. In the social optimum under uncertainty, the volume of users is $\max[0, b_I^S]$ and social marginal benefits are $b_I^S - u$, welfare being $\max[0, \frac{1}{2}b_I^S]^2$.Therefore, the value of information is

$$\text{VOI} = (1-p) \cdot \frac{1}{2}B^2 + p \cdot 0 - \max[0, \frac{1}{2}b_I^S]^2 = \begin{cases} (1-p)\frac{1}{2}B^2 & \text{if } b_I^S < 0 \\ (1-p)\frac{1}{2}B^2 - \frac{1}{2}(b_I^S)^2 & \text{if } b_I^S \geq 0 \end{cases}$$

In case the novel technology turns out to be riskless, not only does the technology become more valuable to a user of the technology, it also increases the number of users for whom the technology yields positive returns at all from a social point of view. Users of the technology produce valuable information not only for themselves but also for potential users whose private benefits are smaller and who are thus not using the technology under uncertainty.

3.2.2 The Costs of Obtaining Information

Obtaining information about the risk may require a volume of users beyond the statically optimal one. By definition this involves a deviation from statically optimal behavior. The static inefficiency incurred are the costs of obtaining information. This inefficiency of deviating to a certain volume v_{dev} is the difference in welfare between the static optimum v_p^S and v_{dev}. Welfare in the static optimum is

$$\max[0, \frac{1}{2}b_I^S]^2,$$

welfare from a given volume v_{dev} is

$$-\frac{1}{2}v_{dev}^2 + v_{dev} \cdot b_I^S.$$

Costs of deviating to v_{dev} are therefore[11]

$$C_{\Delta V}(v_{dev}) = \max[0, \frac{1}{2}b_I^S]^2 - (-\frac{1}{2}v_{dev}^2 + v_{dev} \cdot b_I^S).$$

This is increasing in v_{dev}. It follows that there can only be v_p^S or v_{info} dynamically optimal.

The costs of obtaining information are the costs of the smallest deviation that leads to resolving uncertainty, which is deviating to v_{info}. But if the statically optimal volume is already large enough to reveal information, $v_p^S \geq v_{info}$, there is no trade-off between statically optimal volume in period 1 and availability of information in period 2. Hence, there is no deviation from the statically optimal volume necessary and costs are defined as zero if this is the case.

$$C_{\Delta V}(v_{info}) = \begin{cases} 0 & \text{if } v_p^S \geq v_{info} \\ \max[0, \frac{1}{2}b_I^S]^2 + \frac{1}{2}v_{info}^2 - v_{info} \cdot b_I^S & \text{if } v_p^S < v_{info} \end{cases}$$

3.2.3 Dynamically Optimal Volume

The condition for the information revealing volume of users v_{info} in period 1 being superior to the statically optimal solution v_p^S is

$$\text{VOI} \geq C_{\Delta V}(v_{info}).$$

We summerize dynamically optimal behavior:

Proposition 2.

1. *If* $\text{VOI} \geq C_{\Delta V}(v_{info})$*, information acquisition is optimal. The dynamically optimal care level equals the statically optimal care level* x_p^S *and the optimal volume is* v_{info}*, being made up of all users* $u \geq B - v_{info}$.

2. *If* $\text{VOI} < C_{\Delta V}(v_{info})$*, information acquisition is not optimal. Both dynamically optimal care level and volume equal the statically optimal care level* x_p^S *and volume is* v_p^S*, the latter being made up of all users* $u \geq B - v_p^S$.

It is not difficult to construct both cases in which information information revelation is optimal and in which it is not, e.g. by varying v_{info}, p or coming up with accordingly tailored care-damage relationships $h(x)$. It does not seem possible to provide simple conditions for deciding which volume is optimal. Instead we want to introduce a feature of our model which will prove to dominate other characteristics, given any fixed choice of other parameters: the stakes of our problem. By "stakes" we mean the magnitude of both benefits and costs entailed by the novel technology: the maximum private benefit B and the determinants of the total accident costs x and $h(x)$.

[11] We index costs from a volume different from the statically optimal one with ΔV.

3.3 Optimality of Information Acquisition for High Stakes

We change stakes by introducing an additional parameter into the model: the scale. The aim is to leave the information structure, represented by p and v_{info}, unchanged, as well as the principle functional relationship between care and harm while increasing the maximum private benefit and total accident costs by the same factor. For that reason we introduce benefit and cost parameters depending on the scale σ: $B_\sigma := \sigma \cdot B$ and $h_\sigma(x) := \sigma \cdot h(\frac{x}{\sigma})$. Doing so means scaling up benefits and the care-harm relationship in both directions by σ, therefore leaving the proportions of private benefits, care expenditures and harm unchanged for every state of information.[12] Indeed, using h_σ, we obtain the total accident costs under uncertainty

$$x + p \cdot h_\sigma(x) = x + p \cdot \sigma \cdot h(\frac{x}{\sigma}).$$

Minimization yields

$$x_p^S(\sigma) = \sigma \cdot x_p^S$$

and

$$h_\sigma(x_p^S(\sigma)) = \sigma \cdot h(x_p^S).$$

As we can see, scaling up the care-harm relationship leads to total accident costs of a factor σ larger than before. The same holds of course true under resolved uncertainty.

Why is the scale important? It turns out that the value of information, the value of knowing the exact p, crucially depends on the scale: In case the novel technology turns out to be riskless, not only does the technology become more valuable to a user of the technology, it also increases the number of users for whom the technology yields positive returns at all from a social point of view. Users of the technology produce valuable information not only for themselves but also for potential users whose private benefits are smaller and who should thus not be using the technology under uncertainty. Since both of these effects get larger by increasing the scale, the value of information raises quadratically with the scale:

Lemma 1. *The value of information is homogeneous of degree 2 in the scale:*

$$\text{VOI}(\sigma) = \sigma^2 \cdot \text{VOI}(\sigma = 1) = \sigma^2 \cdot \text{VOI}$$

Information itself and the process of information acquiring remains unchanged, whereas the *impact* and thus the *value* of information changes in two ways, giving rise to the above result.

In order to learn the effects on optimal behavior, we have to analyze how increasing the scale impacts on the cost side of obtaining information. The cost side deals with

[12] We index the changed givens in our model with σ, whereas we write variables derived within the model as functions of σ. In both cases $\sigma = 1$ amounts to the same as with the original parameter choice without scale and we therefore identify the two. We always think of the original givens as a fixed baseline, refering to *the* maximum private benefits B and *the* care-harm relationship $h(x)$ although givens to which the scale is applied technically are just new givens.

static losses, therefore we first have to look at statically optimal behavior under changing scale. This behavior is defined by

$$v_p^S(\sigma) = \max[0, b_I^S(\sigma)] = \max[0, B_\sigma - x_p^S(\sigma) - p \cdot h_\sigma(x_p^S(\sigma))] = \sigma \cdot \max[0, b_I^S]$$

This means that socially optimal volume increases with the scale, if it had been positive in the first place and remains zero under changing scale if it had been zero anyways. Given $v_p^S(\sigma) < v_{info}$, the costs of obtaining information then are

$$C_{\Delta V}(\sigma, v_{info}) = \begin{cases} \frac{1}{2}v_{info}^2 - v_{info} \cdot \sigma \cdot b_I^S & \text{if } v_p^S = 0 \\ \frac{1}{2}(\sigma \cdot b_I^S - v_{info})^2 & \text{if } v_p^S > 0 \end{cases}$$

Since statically optimal volume increases with scale if it had been positive in the first place, obtaining information will eventually become costless as stakes increase and therefore of course optimal in that case. Conversely, if statically optimal volume is zero, costs of obtaining information rise linearly with the scale. This is due to the fact that the social costs of a single user rise, but the volume of users necessary to reveal information does not change. Since the value of information rises quadratically in this case, obtaining information will be optimal as well, if only the stakes are large enough.

Proposition 3. *Given any information structure p and v_{info}, care-harm relationship $h(x)$, and maximum private benefits B, obtaining information is optimal if and only if the scale is large enough: There exists a scale $\hat{\sigma}$, such that obtaining information is optimal if $\sigma \geq \hat{\sigma}$ and not obtaining information is optimal if $\sigma < \hat{\sigma}$.*

Remark. *If the statically optimal volume is zero this effect is driven by dynamic considerations, if it is positive it is driven by static considerations: If the statically optimal volume of users v_p^S is zero, the dynamically optimal volume in case of obtaining information being optimal is v_{info}, if it is positive the dynamically optimal volume in case of obtaining information being optimal is $\sigma \cdot v_p^S$.*

4 Optimal Regulation

We saw that the value of information can be large enough to impact on socially optimal behavior. We now turn toward optimal regulation. We briefly analyze optimal regulation in period 2 and investigate period 1 afterwards.

4.1 Second Period

Since the second is also the last period, there can not be any information effect, independent from whether uncertainty is resolved due to period 1 or not. Due to standard arguments strict liability leads to both optimal care and activity levels. A standard negligence rule leads to optimal care but excessive activity. A negligence rule employing a different due care level may or may not lead to suboptimal activity and hence volume of users, but it leads to suboptimal care for sure.

Proposition 4. *The optimal regulation in period 2 is strict liability, leading to optimal behavior described in proposition 1. Given the optimal volume is positive, any negligence rule is inferior to strict liabilty.*

4.2 First Period

Optimal regulation in period 1 is not that clear-cut. We first analyze behavior and total welfare under strict liability. Afterwards we will analyze optimal negligence rules, the behavior and welfare levels they induce, using strict liability as a baseline.

4.2.1 Strict Liability

From standard arguments we know that strict liability leads to the implementation of statically optimal behavior by all users. Therefore, it always implements socially optimal care levels.

If the statically optimal volume of users already suffices to obtain information, $v_p^S \geq v_{info}$, strict liability implements statically optimal behavior and leads to resolved uncertainty. In this case it implements socially optimal behavior and is therefore first-best. If $v_p^S < v_{info}$, the statically optimal volume is below the uncertainty resolving threshold and behavior under strict liability is still statically optimal but fails to resolve uncertainty. If this is first-best or not depends on whether resolving uncertainty is socially optimal or not: In case

$$\mathrm{VOI} < C_{\Delta V},$$

resolving information is too costly in terms of static welfare losses. Optimal behavior is then again characterized by static optimality and therefore strict liability still implemts the first-best solution. Contrary, if

$$\mathrm{VOI} \geq C_{\Delta V},$$

social optimality requires uncertainty to be resolved but strict liability fails to do so. From a dynamic point of view, this means strict liability presents the wrong incentives for all users $v_p^S < u \leq v_{info}$. Social optimality would demand them to use the technology in order to resolve uncertainty. Their incentives under strict liability do not reflect the social gains from resolving uncertainty but only the static expected harm which exceeds their private benefits. Does strict liability's failure to implement the first-best solution in this case leave room for negligence to enhance welfare?

4.2.2 Negligence Rules

Standard negligence rules automatically make use of the statically optimal level of care as due care standard. If a regulator chooses to employ a negligence rule in our setting, he does so because he wants to generate a volume of users large enough for resolving uncertainty about the technology's risk characteristic. Therefore, a negligence rule would optimally implement the smallest uncertainty resolving level v_{info}. Since the volume implemented under a standard negligence rule may either be higher or lower than v_{info}, the regulator might want to increase or lower the care standard in order to influence the volume of users. Doing so comes with a care efficiency loss. As the regulator does not know single agents' private benefits, he has to employ a homogeneous care standard. Therefore, the care inefficiency to be incurred is also homogeneous among users and a marginal change in the volume of users entails a change in care inefficiency for all users. The dynamically optimal care standard results from minimizing the sum of care and volume inefficiency.

Which is the due care level x_{imp} a regulator has to set, if he wants to implement some predefined volume of users v_{imp}? Since users will exercise due care in order to escape liability[13], users will only face the costs of taking due care. The user with the lowest private benefits $b(u)$ to use the technology will therefore be the one with

$$b(u) = B - u = x_{imp}$$

Therefore, the due care level to set for inducing a volume of v_{imp} is

$$x_{imp} = B - v_{imp}$$

The care inefficiency per user[14] entailed from changing from x_p^S to x_{imp} is

$$\begin{aligned} c_{\Delta x} &= [x_{imp} + p \cdot h(x_{imp})] - [x_p^S + p \cdot h(x_p^S)] = [x_{imp} - x_p^S] - p \cdot [h(x_p^S) - h(x_{imp})] \\ &= [B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})] \end{aligned}$$

Care inefficiency per user is the difference in total accident costs between the actual due care level x_{imp} and the statically optimal level x_p^S. Plugging in the actual care level in terms of volume of users aimed at and rearranging leads to the result. Costs from the care inefficiency are obtained by multiplying costs per user with the actual volume of users:

$$C_{\Delta x} = v_{imp} \cdot c_{\Delta x} = v_{imp} \cdot [B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})]$$

Combined with the costs of volume excessive of the statically optimal volume, we obtain the total costs of implementing a certain volume under a negligence rule:

$$\begin{aligned} &C_T(v_{imp}) = C_{\Delta v}(v_{imp}) + C_{\Delta x}(v_{imp}) \\ &= \begin{cases} \frac{1}{2} v_{imp}^2 - v_{imp} \cdot b_I^S + v_{imp} \cdot [[B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})]] & \text{if } b_I^S \leq 0 \\ \frac{1}{2}(b_I^S - v_{imp})^2 + v_{imp} \cdot [[B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})]] & \text{if } b_I^S > 0 \end{cases} \end{aligned}$$

[13] This is only true as long as minimal total accident costs are not lower than the due care level demanded to escape liability. But this only means that it is not possible to implement a smaller volume than the statically optimal one v_p^S, which is of course never desirable anyways. For that reason we ignore this possibility.

[14] We index costs from a care level different from the statically optimal one with Δx, using lower case c for costs per user and upper case C for absolute costs.

The implementation of a negligence rule always aims at information revelation. An optimal negligence rule therefore has to reveal information while minimizing the toal costs of doing so.

We know that strict liability is first best if either the statically optimal volume of users already suffices to obtain information ($v_p^S \geq v_{info}$) or if it is not desirable to resolve uncertainty due to too high costs in terms of static inefficiency (VOI $< C_{\Delta V}$). In those cases, any negligence rule must necesserily do worse than strict liability. Any negligence rule is statically suboptimal with respect to the optimal volume, to the care level or both.[15]

Contrary, if

$$\text{VOI} \geq C_{\Delta V},$$

social optimality requires uncertainty to be resolved, but strict liability fails to do so. Therefore, any negligence rule which implements a volume of users v_{imp} of at least v_{info} and fulfills the condition

$$\text{VOI} > C_T(v_{imp})$$

is superior to strict liability. Examples can be obtained by considering the special case of $v_{info} = B - x_p^S$. Beyond single examples, the question arises due to negligence's feature of increasing the volume of users: Does the result of social optimality of information revelation under increasing stakes translate to a negligence rule being the optimal regulation for large stakes?

4.3 Optimality of Information Acquisition for High Stakes

For all settings in which the statically optimal volume of users is strictly positive the answer is cleary no. We have

$$v_p^S(\sigma) = \sigma \cdot \max[0, b_I^S],$$

therefore the statically optimal volume will under increasing scale eventually cross the threshold of users required for resolving information. We have already seen that strict liability is first-best in those cases and that any negligence rule necesserily inferior to it.

For settings in which the statically optimal volume is zero, things are clearly different. Irrespective of the scale, the statically optimal volume remains zero and the only possibility to induce potential users to engage in the technology in sufficiently large numbers for resolving uncertainty is to employ a negligence rule. Recall from lemma 1 that the value of information is homogeneous of degree 2 in the scale:

$$\text{VOI}(\sigma) = \sigma^2 \cdot \text{VOI}$$

[15] Negligence is not necesserily strictly worse in this case: If a volume of zero users is socially optimal, any negligence which yields this outcome is as good as strict liability. However, these negligence rules would need to set an exremely high due care level and it would be nothing gained to switch to one of these negligence rules.

For evaluating the cost side, we first have to calculate the volume minimizing total static costs. Despite facing a trade-off between costs in terms of volume and in terms of care, it turns out that the former dominates the latter:

Lemma 2. *If the statically optimal volume is zero, total implementation costs are monotonically increasing in the implemented volume. This holds for any given scale:*

$$C_T'(\sigma, v_{imp}) \geq 0$$

The cost-minimizing negligence rule guaranteeing information revelation is therefore $x_{info} = B - v_{info}$*, implementing the minimum volume necessary to resolve uncertainty. This holds for any given scale.*

For any given scale the total static costs of resolving uncertainty are therefore given by:

$$C_T(\sigma, v_{info}) = C_{\Delta v}(\sigma, v_{info}) + C_{\Delta x}(\sigma, v_{info})$$

While the *impact* of information and hence its value increases with respect to the every single user *and* the number of users with increasing scale, the costs of acquiring information only change for single users, but the number of users required remains fixed. In this sense gains from information revelation grow faster than the costs of implementing a sufficiently large volume by means of an optimal negligence rule. This intuition proves true for all information structures and care-harm relationships, ensuring superiority of negligence in a broad class of settings:

Proposition 5. *Let the statically optimal volume of users be zero. For all information structures* p *and* v_{info} *and care-harm relationships* $h(x)$*, there exists a scale* $\bar{\sigma}$ *such that*

1. *strict liability is opimal if* $\sigma < \bar{\sigma}$
2. *negligence is opimal if* $\sigma > \bar{\sigma}$.

Remark. *Negligence proofs to be superior to strict liability if statically optimal volume is zero and the stakes involved large enough. However, since the implemented volume and the due care level demanded are directly linked, negligence can not implement*[16] *the first-best solution. Although an optimal negligence rule provides the right incentives in terms of activity for all users it does so by deviating from the first-best care level. The second-best care level might be either too low or too high, depending on whether* x_{info} *is higher or lower as first-best care* x_p^S*. Since this occors for large stakes and the statically optimal care level* $x_p^S(\sigma)$ *increases linearly with stakes, the care level implementing uncertainty resolving behavior* x_{info} *is smaller than the statically optimal care level* x_p^S *in most cases (i.e. for all* $\sigma \geq \frac{x_{info}}{x_p^S}$*), and the dynamically optimal due care level therefore stricter than the statically optimal one from a standard negligence rule.*

[16]Except for the special case that $v_{info} = v_p^S$.

Independent from the question of which regulation is optimal, we have seen that under increasing scale the optimal regulation always ensures information acquesition:

Proposition 6. *Given optimal regulation, information is always obtained for any given information structure p and v_{info}, care-harm relationship $h(x)$, and maximum private benefits B, if only the scale is large enough.*

Remark. *If the statically optimal volume is positive, strict liability is the optimal regulation for large enough scales. This is due to the fact that in this case social optimality of information acquisition is driven by static considerations. Contarary, if the statically optimal volume is zero, social optimality of information acquisition for large enough scales is driven by dynamic considerations, hence implemented by an optimal negligence rule.*

5 Conclusions

How should tort law deal with agents that employ novel technologies that later turn out to involve environmental harm? Different legal systems come to different, but often controversial answers. This is despite the economic literature consistently favoring strict liability, based on comparisons of the efficiency of different environmental liability regimes. The starting point of this paper was that a possible objection to this literature is it derives its conclusions in a peculiar information environment: Risk characteristics are learned by acquiring non-experiential private revelation. We believe that a more realistic information environment is one in which additional information about a technology's risk characteristics in the field is a result of learning from experiential public data. In such a setting, the liability regime and learning are tightly linked. We find that in a two-period unilateral care model, the superiority of strict liability for environmental harm is no longer guaranteed. Instead, an optimally designed negligence rule can provide a better balancing of benefits, environmental harm, care effort, and learning. This effect is guranteed if the novel technology would not be used at all under purely static considerations and the stakes involved, potential benefits and total accident costs, are very large. This gives rise to our main message: When in-situ experience with a technology is an irreducible part of risk assessment, then the economics of torts provide an argument for an important role of negligence rules.

The analysis presented in this paper offers a variety of avenues for further work and generalization. One area of further work concerns the optimal balance between laboratybased tests and in-situ experience. This optimal balance, and the ability of different liability regimes to implement it, is alluded to in Dana (2010), but technically hinges on assumptions about the data-generating processes in these two domains. A second area of further research is a departure from a two-period setting to a setting in which postmarket monitoring is a continuous process, giving rise to a stopping rule that triggers the introduction of strict liability. This includes a consideration of the dynamics of information diffusion among different agents in the population, for which our assumption of public observability was an obvious modelling shortcut.

Appendix

Proof of Lemma 1. We have

$$b_I^S(\sigma) = B_\sigma - x_p^S(\sigma) - h_\sigma(x_p^S(\sigma)) = \sigma \cdot B - \sigma \cdot x_p^S - \sigma \cdot h(x_p^S) = \sigma \cdot b_I^S$$

and

$$\text{VOI}(\sigma) = \begin{cases} (1-p)\frac{1}{2}B(\sigma)^2 & \text{if } b_I^S(\sigma) < 0 \\ (1-p)\frac{1}{2}B(\sigma)^2 - \frac{1}{2}(b_I^S(\sigma))^2 & \text{if } b_I^S(\sigma) \geq 0 \end{cases}$$

Since $b_I^S(\sigma) = \sigma \cdot b_I^S$, a change in the scale does not alter the case distinction. Therefore, we only have to show that each case itself is homogeneous of degree 2 in the scale. This directly follows from $B_\sigma = \sigma \cdot B$, $b_I^S(\sigma)) = \sigma \cdot b_I^S$ and the fact that p is not affacted by the scale. □

Proof of Proposition 3. First, suppose statically optimal volume v_p^S is positive. If $v_p^S(\sigma) = \sigma \cdot v_p^S \geq v_{info}$, costs of obtaining information are zero. If $v_p^S(\sigma) = \sigma \cdot v_p^S = \sigma b_I^S < v_{info}$, costs are $\frac{1}{2}(\sigma \cdot b_I^S - v_{info})^2$. Hence, costs are weakly decreasing in the scale. Additionally, costs converge to $\frac{1}{2}v_{info}$ if the scale goes to zero. The value of information is homgeneous of degree 2 in the scale. Hence it is strictly increasing in the scale and converging to zero if the scale goes to zero. Both costs and value of information are continuous in the scale. Therefore we have exactly one scale $\widehat{scale}$ s.t. costs and value of information equal. For any smaller scale costs outweigh the value of information and for any larger scale vice versa.

Now suppose statically optimal volume v_p^S is zero. Costs are now $\frac{1}{2}v_{info}^2 - v_{info} \cdot \sigma \cdot b_I^S$ with $b_I^S \leq 0$. The only difference to the first case is that costs are now increasing in the scale. Since they do so linearly, the value of information equals costs for some scale $\hat{\sigma}$. Again, this is the scale we were looking for. □

Proof of Lemma 2. We have

$$C_T(v_{imp}) = \frac{1}{2}v_{imp}^2 - v_{imp} \cdot b_I^S + v_{imp} \cdot [B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})]$$

Therefore

$$\begin{aligned} C_T'(v_{imp}) &= [v_{imp} - b_I^S] + [[B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})]] + v_{imp} \cdot [-1 - p \cdot h'(B - v_{imp})] \\ &= -b_I^S + [[B - x_p^S - v_{imp}] - p \cdot [h(x_p^S) - h(B - v_{imp})]] - v_{imp} \cdot p \cdot h'(B - v_{imp})] \end{aligned}$$

The first term is non-negative since $b_I^S \leq 0$, the second by definition (since x_p^S minimizes total accident costs), and the third since $h'(\cdot) < 0$. □

Proof of Proposition 5.

$$\begin{aligned} C_T(v_{info}, \sigma) &= C_{\Delta v}(v_{info}, \sigma) + C_{\Delta x}(v_{info}, \sigma) = \sigma^2 \cdot C_{\Delta v}(\frac{v_{info}}{\sigma}, 1) + \sigma^2 \cdot C_{\Delta x}(\frac{v_{info}}{\sigma}, 1) \\ &= \sigma^2 \cdot C_T(\frac{v_{info}}{\sigma}, 1) \end{aligned}$$

Comparing the value of information VOI(σ) and total implementation costs $C_T(v_{info}, \sigma)$ under increasing scale, the σ^2 cancels out. We have then to compare some positive VOI with $C_T(\frac{v_{info}}{\sigma}, 1)$. For $\sigma \to \infty$ we have $\frac{v_{info}}{\sigma}$ since v_{info} is fixed. The costs of obtaining information are continuous in the first argument and clearly approach zero if $\frac{v_{info}}{\sigma} \to 0$. $\square$

References

[1] Bajtelsmit, V. and Thistle, P. (2009). “Negligence, ignorance and the demand for liability insurance”. In: *The Geneva Risk and Insurance Review* 34.2, pp. 105–116.

[2] Ben-Shahar, O. (1998). “Should Products Liability Be Based on Hindsight?” In: *Journal of Law, Economics, and Organization* 14.2, pp. 325–357.

[3] Bouchie, A. (2002). “Organic farmers sue GMO producers”. In: *Nature Biotechnology* 20.3, pp. 210–210.

[4] Bullock, D. S. and Desquilbet, M. (2002). “The economics of non-GMO segregation and identity preservation”. In: *Food Policy* 27.1, pp. 81–99.

[5] Crocker, K. J. and Doherty, N. (2001). *Why people buy liability insurance under the rule of simple negligence.* Vol. 9. Emerald Group Publishing Limited.

[6] Dana, D. (2010). “When Less Liability May Mean More Precaution: The Case of Nanotechnology”. In: *UCLA Journal of Environmental Law and Policy* 28, pp. 153–199.

[7] Dowling, A., Clift, R., Grobert, N., Hutton, D., Oliver, R., O’neill, O., Pethica, J., Pidgeon, N., Porritt, J., Ryan, J., et al. (2004). “Nanoscience and nanotechnologies: opportunities and uncertainties”. In: *London: The Royal Society & The Royal Academy of Engineering Report*, pp. 61–64.

[8] Endres, A. and Bertram, R. (2006). “The development of care technology under liability law”. In: *International Review of Law and Economics* 26.4, pp. 503–518.

[9] Endres, A. and Friehe, T. (2011). “Incentives to diffuse advanced abatement technology under environmental liability law”. In: *Journal of Environmental Economics and Management* 62.1, pp. 30–40.

[10] Faure, M. and Wibisana, A. (2010). “Liability for Damage Caused by GMOs: An Economic Perspective”. In: *Geo. Int’l Envtl. L. Rev.* 23, p. 1.

[11] Fontanarosa, P. B., Rennie, D., and DeAngelis, C. D. (2004). “Postmarketing surveillance—lack of vigilance, lack of trust”. In: *Jama* 292.21, pp. 2647–2650.

[12] Kaplow, L. and Shavell, S. (1992). “Private versus Socially Optimal Provision of Ex Ante Legal Advice”. In: *Journal of Law, Economics, & Organization*, pp. 306–320.

[13] Moore, T. J., Psaty, B. M., and Furberg, C. D. (1998). “Time to act on drug safety”. In: *Jama* 279.19, pp. 1571–1573.

[14] Parchomovsky, G. and Stein, A. (2008). “Torts and innovation”. In: *Michigan law review*, pp. 285–315.

[15] Schieber, J. “Das Entwicklungsrisiko im Rahmen der Umwelthaftung und der Umwelthaftpflichtversicherung”. In: *Versicherungsrecht* 19.

[16] Schwartzstein, J. and Shleifer, A. (2013). “An Activity-Generating Theory of Regulation”. In: *Journal of Law and Economics* 56.

[17] Shavell, S. (1980). “Strict liability versus negligence”. In: *The Journal of Legal Studies*, pp. 1–25.

[18] Shavell, S. (1992). "Liability and the incentive to obtain information about risk". In: *The Journal of Legal Studies*, pp. 259–270.
[19] Skogh, G. (1998). "Development risks, strict liability, and the insurability of industrial hazards". In: *Geneva Papers on Risk and Insurance. Issues and Practice*, pp. 247–264.