Smart Products: Liability, Investments in Product Safety, and the Timing of Market Introduction*

Herbert Dawid[†] Gerd Muehlheusser[‡]

February 21, 2020

Abstract

We analyze the role of product liability for the emergence and development of smart products such as autonomous vehicles (AVs). We develop, and calibrate to the U.S. car market, a dynamic model where a (monopolistic) innovator chooses safety stock investments, the timing of market introduction, and the product price. Inducing higher long-term product safety through a strict (partial) liability rule reduces short-term safety investments and slows down AV market penetration. By contrast, negligence-based liability fosters initial investments without hampering long-term product safety. However, too stringent liability might forestall investments in the development of AVs and their market introduction.

JEL-Codes: O31, K13, L11, L62

Keywords: Product Innovation, Product Liability, Smart Products, Autonomous Vehicles, Optimal Investment Dynamics

*We gratefully acknowledge very valuable comments and suggestions by Shay Lavie, Kathy Spier, and Avi Tabbach. We also received useful feedback from participants of the Viennese Vintage Workshop 2019 as well as from seminar audiences at Bielefeld University, National Taiwan University, and Tel Aviv University. Fabienne Bovier provided excellent research assistance. This article is a substantial revision of a previous version circulated as CESifo Working Paper No. 7673. Dawid gratefully acknowledges financial support from the German Research Foundation (DFG) through the Collaborative Research Centre (SFB 1283) "Taming uncertainty and profiting from randomness and low regularity in analysis, stochastics and their application".

[†]Bielefeld University, Department of Business Administration and Economics, and Center for Mathematical Economics, Universitätsstrasse 25, D-33615 Bielefeld, Germany. hdawid@uni-bielefeld.de

[‡]University of Hamburg, Department of Economics, Von-Melle-Park 5, D-20146 Hamburg, Germany; CESifo and IZA. gerd.muehlheusser@uni-hamburg.de

1 Introduction

Motivation The advent of digitization has already changed our everyday lives in many fundamental ways, and it is widely expected that this development will continue in the future at an even accelerated speed. As a consequence, we will move more and more towards a *digital economy*, with far-reaching consequences for many areas such as workplace relationships, product markets, international trade or mobility. These developments also constitute major policy challenges, for example with respect to infrastructure investments, standardization, regulation, and legislation.[1]

From a technological point of view, the key drivers for these developments are recent advances in areas such as robotics, artificial intelligence (AI), cloud computing, or big data analytics (see e.g. UN Conference on Trade and Development, 2017), which enable the development of so-called *smart products*. According to Porter and Heppelmann (2015), the capabilities of smart products can be grouped in four areas: monitoring, control, optimization, and autonomy. This article focusses on autonomy. Products exhibiting autonomy rely less and less on human decision-making, and major decisions are instead being taken by the product's operating system, based on algorithms. One example of particular topicality are Autonomous Vehicles (AVs), which is the prime application in this article, but similar developments are also observed for other smart products ranging from household appliances to heavy military weapons.[2]

One alleged major benefit of AVs is that in the long run, when fully matured, they are potentially much safer than conventional vehicles in terms of their accident rate.[3] The main reason is that computerized decision-making based on AI algorithms is potentially less error-prone than decision-making by human drivers, who are often impaired by poor sight and slow reactions times, for example due to fatigue, distraction, alcohol or drug consumption. While (conventional) cars have become considerably safer over the last decades, car crashes are still a significant problem and create a huge social cost, so that any significant reduction in accident risks would generate potentially large social benefits.[4]

However, these long-run benefits of AVs cannot be expected to be reaped immediately upon their market introduction. In the meantime, as long as the technology is not yet fully matured, AVs are likely to be involved in crashes in the same way as conventional cars are, potentially at even a higher rate. This is

[1]For an overview over some of these policy challenges, see e.g. OECD (2017); UN Conference on Trade and Development (2017).

[2]See Schellekens (2015, p.507) for a more detailed classification of different degrees of automation for vehicles. Sassone et al. (2016) provide a survey of different areas where smart products are expected to be available in the near and intermediate future.

[3]Apart from a lower accident rate, AVs are expected to give rise to further benefits such as reducing traffic congestion or improving mobility of the elderly (see e.g. Douma and Palodichuk, 2012).

[4]For example, according to a recent report by the U.S. National Highway Traffic Safety Adminstration (Blincoe et al., 2015), in 2010 there were 32,999 people killed in car crashes in the U.S., and 3.9 million were injured. Moreover, for that year the total social costs due to car accidents is estimated to be $ 836 billion.

highlighted by the recent fatal crashes caused by AVs of Uber and Tesla during test drives, triggering considerable public attention.[5]

Under the assumption that the safety of AVs increases over time as the technology improves, this raises the question of *when* and with which safety level AVs should be brought to the market. For example, it is argued that upon introduction, AVs should be at least as safe as conventional cars, if not much safer (see e.g. Schellekens, 2015, p. 510). In a similar vein, in a recent panel discussion Dieter Zetsche, then CEO of Daimler (one of the world's major car manufacturers), argued that even if AVs caused only a tenth of the number of fatalities as conventional cars, this would still not be tolerated by the public. As a result, "they will have to be a hundred times better, but then it gets really difficult from a technological point of view".[6]

The possibility of accidents caused by AVs which are controlled by non-human drivers raises interesting and novel aspects with respect to the issue of product liability. In particular, unlike other and less significant product innovations, in the case AVs a new type of injurer emerges: the AV's operating system.[7] As computer systems cannot be held legally responsible, however, the question arises who should be liable for the harm arising from accidents caused by AVs? As one main feature (and alleged benefit) of AVs is that its driver retains less control and hence becomes more like a passenger rather than the operator of the car, many argue that there should be less scope for legal responsibility of drivers, and more scope for car manufacturers and automotive software developers.[8] Despite the current legal uncertainty to which extent producers of AVs and their operating systems will ultimately be held liable, they are expected to face substantially higher liability costs than they currently do with conventional cars.[9]

One crucial question is to which degree these higher liability risks affect firms' innovation activities and other business policies, such as R&D investments (e.g. to improve the safety of AVs), the timing of market

[5]See e.g. `https://www.nytimes.com/2018/03/31/business/tesla-crash-autopilot-musk.html` and `https://www.nytimes.com/2018/03/19/technology/uber-driverless-fatality.html`.

[6]See `https://www.faz.net/-ijt-9gpve`.

[7]Marchant and Lindor (2012, p.1325) discuss autopilots in airplanes as one example with similar features.

[8]Galasso and Luo (2018a) discuss some general legal challenges for tort law in the age of AI. Moreover, Douma and Palodichuk (2012), Colonna (2012), Marchant and Lindor (2012), Schellekens (2015) and Gless, Silverman, and Weigend (2016) provide legal discussions in the context of AVs in the United States. Moreover, the European Parliament has adopted a resolution containing suggestions how to establish legal rules for tort cases with AI tortfeasors (European Parliament, 2017). Finally, Germany's "Ethics Commission for Autonomous and Connected Driving" recommends that apart from car owners, also the producers and operators of autonomous cars and its supportive technological systems should be considered as potentially liable parties (Federal Ministry of Transport and Digital Infrastructure, 2017).

[9]Currently (i.e. without AVs on the streets), in the vast majority of accidents the damage is apportioned between the involved drivers and/or their insurance companies. For example, according to the 2008 *National Motor Vehicle Crash Causation Survey* (NHTSA), more than 93 percent of the analyzed crashes were classified as being caused by erroneous behavior of the respective drivers, and only the remaining 7 percent were due to vehicle problems or adverse weather conditions. Car manufacturers currently face a risk of product liability, which might apply when a car model exhibits systematic technical defects. Moreover, Viscusi (2012) argues that there is a tendency of courts and juries to be tougher on injurers in liability cases where "novel risks" are involved. Given the high degree of product novelty of AVs compared to conventional cars, this could further contribute to the expected liability cost of firms engaged in the development and production of AVs.

introduction of AVs and pricing. Together these decisions determine (i) when AVs are eventually launched (if at all) and with which safety level, (ii) how quickly they penetrate the market, and (iii) how long it takes until the long-term benefit of a higher safety of AVs compared to conventional cars can be expected to materialize. For example, legal scholars such as Marchant and Lindor (2012, pp. 1336), Colonna (2012, p.84) and Schellekens (2015) argue that too stringent liability for producers of AVs might considerably delay these processes.

More generally, from a theoretical perspective, the effect of stringent liability on the incentives to innovate seems ambiguous: one the one hand, it might increase the incentive to develop innovations which improve product safety. But on the other hand, it might have a "chilling effect" on innovative activities, for example in the form of choosing to not develop new products (Viscusi, 1991; Parchomovsky and Stein, 2008; Viscusi, 2012), even if they are potentially superior than existing ones.[10] From an empirical point of view, both effects seem relevant (see e.g. McGuire, 1988; Viscusi and Moore, 1993; Galasso and Luo, 2017, 2018b). Hence, in the light of the potentially large benefits from AVs in the long-term, the stringency of liability for firms involved in the development, production and operation of AVs and other smart products seems a highly relevant policy issue.

Framework and results We consider a dynamic model of product innovation, where there already exists a market for an established product (the *old* product), and where a monopolistic innovator can introduce a new one (the *smart* product). Consumers differ with respect to their valuation for the smart product. Both products are prone to accidents, which are socially harmful. At each point in time, the innovator decides on (i) how much to invest into the safety level of the smart product which reduces its accident rate, and (ii) the price at which it is sold, once it has been launched.

We analyze how the dynamically optimal decisions and the resulting state dynamics depend on the liability regime, which apportions the damage accruing from accidents with the smart product between the innovator and the consumers. More precisely, we are interested in how liability affects (i) the safety level of the smart product in the long run (steady state), (ii) the optimal dynamic paths for safety investments as well as output and pricing choices, thereby also determining when and with which safety level the smart product is launched, (iii) the development of the smart product's safety level after market introduction and (iv) overall welfare. In the main part of this article, we examine these issues in a framework of strict (partial) liability for the innovator. The stringency of liability is determined by the fraction of the damage from AV accidents accruing to the innovator.

[10] This was also a crucial argument in the decision of the U.S. Supreme Court in *Riegel vs. Meditronic* in 2008.

We first provide an analytical characterization of the effect of the stringency of liability on the steady states of the model. The dynamic properties of the optimal paths are then examined using numerical methods, where we calibrate the model using data from the US automobile market. Considering the impact of liability on the optimal behavior of the innovator, our analysis reveals two basic trade-offs for policy makers and legislators: The first trade-off concerns the existence and optimality of an *active* steady state, in which the innovator indeed accumulates a safety stock for the smart product, and eventually introduces it to the market: As liability becomes more stringent, the safety level of the smart product in an active steady state increases. However, too stringent liability might forestall the development of the smart product altogether. In particular, the innovator might prefer to induce a *passive* steady state, in which there is no investment into the safety stock of the smart product, and where it is never launched. Overall, these results are consistent with empirical findings that the effect of liability on innovation activities is ex ante unclear and might go in either direction (Viscusi and Moore, 1991, 1993; Galasso and Luo, 2017, 2018b).

The second trade-off concerns the short- and long term effects of liability: More stringent liability has the long-run benefit of inducing a higher steady state safety level of the smart product. However, the analysis of the investment dynamics reveals that it also leads to a delay in the accumulation of the safety stock in early periods. Intuitively, under more stringent liability it takes longer until the smart product is launched, quantities are smaller, and hence the market penetration occurs more slowly. This reduces the innovator's incentives for safety investments. As a result, the units brought to the market in the short- and intermediate term exhibit a lower safety level compared to those that would have been produced under less stringent liability. For the context of AVs, our analysis suggests that these short-term costs outweigh the long-term benefit from more stringent liability. As a result, under strict (partial) liability it would be socially optimal to *not* impose additional liability costs on the producers of AVs, compared to the status quo with conventional cars only.

We also analyze two alternative liability rules: First, we consider *residual manufacturer liability* as analyzed by Hay and Spier (2005), and our findings are qualitatively very similar to the baseline setting. Second, we analyze *negligence-based* liability, where the innovator is only liable if the AV's safety level is below a given negligence standard. We find that such a liability rule ameliorates the issue of delay in safety stock accumulation, and thereby also reduces quantity distortions relative to the social optimum. Consequently, from a welfare perspective, an appropriately designed negligence-based rule might outperform a rule with strict (partial) liability. Finally, we also consider a policy of direct safety regulation of AVs in the sense that they can only be launched once they satisfy a given minimum safety standard. We show that the induced investment patterns and market dynamics are very similar to the case of negligence-based liability.

Relation to literature Our article is related to a sizeable literature in law & economics on how (product) liability affects the incentive of firms to improve the safety of their products, either by making existing products safer or by developing new products which exhibit higher safety levels. For example, in the survey by McGuire (1988) conducted among firms, about a third of respondents answered that the threat of liability has led them to improve the safety of their products, but an equal number of respondents said that it has lead them to not introduce new products. In their rather critical general discussion of product liability, Polinsky and Shavell (2010) argue that incentives to invest into product safety can also be stipulated through other (and arguably less expensive) channels, such as market forces (as consumers anticipate that they will have to bear the consequences of the accident risk, which reduces their willingness to pay for the product) or direct safety regulation. These arguments are supported by our analysis in the sense that the consumer side is indeed an important driver of investments into the safety of AVs, even when product liability is weak.

As for the impact of liability on innovation incentives, the evidence is mixed. For example, in their seminal empirical study in the U.S. manufacturing sector Viscusi and Moore (1993) find that the relationship between liability and the intensity of firms' innovative activities is positive (negative) when the stringency of liability is high (low), with the average effect being positive.[11] More recently, Galasso and Luo (2017) empirically analyze the impact of tort reforms in the health sector, which have reduced the stringency of liability. They find that this has led to a decrease of innovative activities of upstream suppliers. By contrast, Galasso and Luo (2018b) consider vertical relationships in the medical implant industry. They find a strong negative effect of higher liability of upstream suppliers on their innovative activities in that sector. Overall, the existing empirical evidence suggests that whether or not liability has a chilling effect on innovation activities crucially depends on the industry under consideration and how stringent product liability actually is. This view is supported also in our analysis using a calibration based on the U.S. automobile market. Viscusi and Moore (1993), Galasso and Luo (2017) and Galasso and Luo (2018b) also develop static models in which they capture the main effects of the empirical analysis. By contrast, our dynamic framework allows for a more detailed analysis of the interplay of the various effects in the short- und long-term, and of the timing of innovations.

Daughety and Reinganum (1995) consider a setting of incomplete information, where a monopolistic firm may use the product price as a signaling device for product safety (which is unobserved by consumers). They analyze the impact of liability on the equilibrium properties. Their model is static, but the timing of events is similar in the sense that safety investments have an impact on subsequent market decisions. In addition, their finding that direct safety regulation might outperform liability rules in terms of social

[11] See also Viscusi (1991), Viscusi and Moore (1991), and Viscusi (2012).

welfare, also occurs in our framework, although the underlying mechanisms differ.

In addition, two recent articles have addressed the issue of liability in the specific context of AVs: Shavell (2019) analyzes optimal liability rules in a static setting with AVs only. Our approach can hence be seen as complementary, as we focus on the dynamic process of development and market penetration of AVs, and the (potentially long) transition period in which autonomous and conventional vehicles coexist. A hybrid setting with both types of vehicles is also considered in Friedman and Talley (2019). In a static model with multilateral care, in which not only producers can contribute to the reduction of accident risks, they analyze the efficiency properties of several negligence-based rules. However, they do not consider investment or innovation incentives, which are crucial aspects of our framework.

A further related strand of literature are dynamic models of innovation. As the safety standard of an AV is a key quality parameter, the innovator's investment in building up the safety stock is closely related to (quality-enhancing) product innovation investments. Although dynamic models of optimal innovation investments have mainly focused on process innovations, several contributions (e.g. Lambertini and Mantovani, 2009; Dawid et al., 2015) have characterized optimal product innovation investments of monopolists in different market settings. With respect to the explicit consideration of the interplay between innovation investments and the timing of market introduction, our article is related to Hinloopen, Smrkolj, and Wagener (2013), where the focus is, however, on process innovation. Moreover, none of these articles considers the impact of liability. Our numerical analysis uses collocation methods for solving the Hamilton-Jacobi-Bellman equation characterizing the optimal investment behavior of the innovator. In that respect our article is related to contributions such as Doraszelski (2003), Dawid et al. (2015) and Dawid, Keoula, and Kort (2017) where such methods have been employed for the numerical analysis of dynamic investment problems.

Last, but not least, we contribute to a current academic and non-academic legal debate about how existing laws and other legal procedures need to be adapted in response to the new legal challenges arising in the digital economy, in particular in the context of AVs (see e.g. Marchant and Lindor, 2012; Colonna, 2012; Douma and Palodichuk, 2012; Schellekens, 2015; Smith, 2017; Wagner, 2018). To the best of our knowledge, our article is the first one to analyze these issues from a theoretical perspective in a dynamic model framework.

The remainder of the article is organized as follows: Section 2 lays out the model framework. Section 3 analyzes the different steady states of the model, and how the steady state optimally induced by the innovator depends on the stringency of product liability. Section 4 analyzes in more detail how product liability affects the dynamic paths leading into the respective steady state. In doing so, we proceed numerically and calibrate the model to the US automobile market. In Section 5 we consider negligence-based liability rules and direct safety regulation. Section 6 explores the policy implications of our findings and discusses further model extensions. The Appendix contains all proofs, details about our numerical method, and additional numerical results.

2 Model

We consider a continuous-time framework in which a monopolistic firm (the "innovator") can introduce a smart product to the market. At each point in time t, there is a unit mass of consumers. Each consumer i has unit demand and is characterized by her gross valuation v_i obtained from the smart product in every period in which it is functional. The valuations v_i are distributed uniformly on $[\underline{v}, \bar{v}]$ with $\underline{v}, \bar{v} > 0$, so that the resulting cumulative distribution function is $F(v_i) = \frac{v_i - \underline{v}}{\bar{v} - \underline{v}}$. Alternatively, consumers can purchase a standard or "old" product which yields the same value v_o to each consumer in every period in which it is functional, where $\underline{v} < v_o < \bar{v}$. The unit cost of production of both the old and the smart product is $c \geq 0$ (there are no fixed costs of production).[12] The old product is sold at a competitive market at price $p_o = c$ in all periods, whereas the price of the smart product in period t is denoted by $p(t)$.

The lifetime of both products is distributed exponentially with a constant failure rate $\rho > 0$. In the course of a product's lifetime accidents may occur, each causing a damage $D > 0$. For the old product, an accident occurs at a given constant rate $\alpha_o > 0$. The accident rate for the smart productis $\alpha(S)$, where S is the innovator's knowledge stock determining the safety of the smart product at the time when it is sold (explained in more detail below). We assume that a higher safety stock reduces the accident rate $\alpha(S)$ of the smart product, with a decreasing marginal effect (i.e. $\alpha'(S) < 0$ and $\alpha''(S) > 0$). Moreover $\lim_{S \to \infty} \alpha(S)$ is denoted by $\alpha^{min} \in [0, \alpha_o)$, i.e. AVs can become safer than conventional cars if the innovator builds up a sufficiently high safety stock, but accidents might not be fully avoided even under enormous safety investments.

As discussed in the Introduction, the scope of product liability is expected to increase upon the introduction of smart products. Consequently, in our model, producers are not liable for accidents with the old product, and all liability risk is borne by the consumers. By contrast, for accidents with a smart product, a fraction $\beta \cdot D$ is borne by the innovator, so that the parameter $\beta \in [0, 1]$ measures the strength of the liability risk for producers of smart products.[13] We refer to this form of liability as *strict (partial) liability.*

[12] As the focus of our analysis is on the accident rate of the smart product and the induced liability risk (compared to the old product), we abstract from differences in production costs between the two products.

[13] Of course, also the producers of conventional cars are currently facing product liability (at least in the U.S.). As our focus is on the effect of additional liability risks due to accidents being caused by the operating system of AVs, we implicitly subsume all other liability risks (e.g. design defects of the brake system or the gas tank) in the production costs of firms. Accordingly, we associate a legal regime where the innovator faces the same liability risk as the producers of the conventional car with $\beta = 0$.

In our model, we capture the empirical fact that, at the margin, consumers often do not fully internalize their liability share, for example, due to insurance policies with deductibles, or under-insurance (or even no insurance at all) in combination with wealth constraints.[14] Formally, denoting by y the damage not borne by the producers (that is, $y = D$ for accidents with the old and $y = (1-\beta)\cdot D$ for accidents with the smart product), we assume that consumers cover $k(y)$ where $k(\cdot)$ satisfies $k' > 0$, $k'' < 0$, $k(0) = 0$, and $k'(0) = 1$.[15] Intuitively, whereas consumers strongly (or even fully) internalize small damages, the degree of internalization decreases as the damage becomes larger. As a result, the expected liability of each consumer per unit of time is $\alpha_o \cdot k(D)$ and $\alpha(S)\cdot k((1-\beta)D)$ for the old and smart product, respectively. Hence, for either product only a part of the total damage is internalized by the producers and consumers. This gives rise to a negative externality for each unit sold, which in case of the smart product is decreasing in β.

For a given price p for the smart product in a given period, consumer i's expected discounted value from purchasing the old respectively the smart product is

$$\begin{aligned} U_o &= \int_0^\infty e^{-(r+\rho)t}\left[v_o - \alpha_o k(D)\right] dt - c \\ &= \frac{v_o - \alpha_o k(D) - (r+\rho)c}{r+\rho} \qquad (1) \\ U(p,S;v_i) &= \int_0^\infty e^{-(r+\rho)t}\left[v_i - \alpha(S)k((1-\beta)D)\right] dt - p \\ &= \frac{v_i - \alpha(S)k((1-\beta)D) - (r+\rho)p}{r+\rho}, \qquad (2) \end{aligned}$$

where $r > 0$ denotes the common discount rate of all market participants. We assume $v_o > \alpha_o k(D) + (r+\rho)c$ and hence $U_o > 0$, so that all consumers purchase either one of the two products. Consumer i with valuation v_i prefers the smart product if and only if $U(p,S;v_i) \geq U_o$. The innovator has an incentive to sell positive

All qualitative findings of our analysis would remain unchanged when considering instead a strictly positive (but not too large) lower bound on β.

[14]Uninsured driving is an empirically relevant phenomenon. For example, whereas there is mandatory insurance in the virtually all states of the U.S., a recent study of the Insurance Research Institute (IRC) estimates that in 2015, 13 percent of all motorists were uninsured (`https://www.insurance-research.org/sites/default/files/downloads/UMNR1005.pdf`). The high empirical relevance of *judgement-proofness* in tort cases is extensively discussed by Gilles (2006). As a result, the personal liability of many individuals is often much smaller than the damage caused in the course of their negligent behavior. For example, when the consumer has an insurance contract with a deductible, her marginal liability effect is one for damages below the amount of the deductible, and zero above it. Whereas larger damages might result in higher insurance premiums, all we need is that the expected liability cost faced by consumers increases under-proportionally in the underlying (remaining) damage amount.

[15]As will become clear below, when the residual damage not covered by the innovator is completely borne by consumers (i.e. $k((1-\beta)D) = (1-\beta)D$), then under linear demand, any shift of liability between the innovator and the consumers would be offset one-to-one by a respective price change. This feature also arises, for example, in the setting of Daughety and Reinganum (1995) with complete information, where the firm's profit depends only on the total social damage from an accident, but not on how it is divided between the firm and the consumers.

quantities of the smart product if and only if there are some consumers who would purchase it when it is priced at marginal costs. This corresponds to the condition $U\left(c+\frac{\beta\alpha(S)D}{r+\rho}, S; \bar{v}\right) \geq U_o$. We denote the lowest safety stock at which this condition holds by $\underline{S}(\beta)$. The corresponding accident rate is denoted by $\underline{\alpha}(\beta) = \alpha(\underline{S}(\beta))$. The value of $\underline{\alpha}(\beta)$ is determined by $U\left(c+\frac{\beta\underline{\alpha}D}{r+\rho}, \underline{S}; \bar{v}\right) = U_o$ and therefore $\underline{S}(\beta)$ is characterized by

$$\alpha(\underline{S}(\beta)) = \frac{\bar{v} - v_o + \alpha_o k(D)}{\beta D + k((1-\beta)D)} > 0, \tag{3}$$

where we assume $\alpha(0) > \underline{\alpha}(\beta) > \alpha_{min}$ for all $\beta \in [0,1]$.[16] Hence, the innovator sells positive quantities of the smart product whenever $S > \underline{S}(\beta)$. In this case demand is determined by the mass of consumers whose valuation for the smart product is above that of the indifferent consumer with valuation $v^{ind}(p,S)$. The (linear) demand function $q(p,S)$ for the smart product can be derived from $q(p,S) = 1 - F(v^{ind}(p,S))$ as

$$q(p,S) = \min\left[1, \frac{1}{\bar{v} - \underline{v}}\Big(A - (r+\rho)(p-c) - \alpha(S)k((1-\beta)D)\Big)\right], \tag{4}$$

where $A = \bar{v} - v_o + \alpha_o k(D)$. The second term in (4) gives the (flow utility equivalent of the) price difference between the two products, and the last term captures the effect of consumers' liability risk from the smart product.

At each point in time t, the innovator chooses an investment level $x(t) \geq 0$ which improves the safety stock of the smart product. This safety stock evolves according to

$$\dot{S}(t) = x(t) - \delta S(t) \tag{5}$$

where $\delta > 0$ measures the stock depreciation.[17] As we are interested in the innovator's ex ante decision whether or not to invest in a wholly new technology (the smart product), we assume $S(0) = 0$. As is common in dynamic investment models, the cost of investment is quadratic and given by $h(x) = \theta x\left(1 + \frac{\eta}{2}x\right)$, where $\theta > 0$ measures the marginal costs at $x = 0$, and $\eta > 0$ is a parameter capturing convexity.

The innovator chooses the trajectories $p(t)$ and $x(t)$ to maximize its expected discounted profit, which is given by

$$\widetilde{\Pi} = \int_0^\infty e^{-rt}\left[q(p(t), S(t)) \cdot \left(p(t) - c - \int_t^\infty e^{-(r+\rho)(\tau-t)}\alpha(S(t))\beta D \, d\tau\right) - h(x(t))\right] dt \tag{6}$$

subject to the state dynamics of the safety stock as described in (5). Thereby, the second integral term in

[16] Due to the strict monotonicity of the function $\alpha(S)$, this implies that (3) has a unique strictly positive (and finite) solution $\underline{S}(\beta) > 0$.

[17] That a positive investment is needed to maintain the stock can be interpreted as a reduced-form modeling of the fact that the complexity of the smart product increases over time.

(6) denotes the expected (discounted) liability cost generated by each unit of the AV sold at date t.

As the price choice has no intertemporal effect, in each point in time t, it can be determined by standard monopoly pricing.[18] This leads to the following expressions for the monopoly price and quantity for $S > \underline{S}(\beta)$:

$$p^*(S) = c + \frac{1}{2}\left[\frac{A + \alpha(S) \cdot (\beta D - k((1-\beta)D)}{r+\rho}\right], \quad (7)$$

$$q^*(S) = \frac{A - \alpha(S)(\beta D + k((1-\beta)D))}{2(\bar{v} - \underline{v})}. \quad (8)$$

Note that whereas the impact of the safety stock S on the monopoly quantity is clearly positive (i.e. q^* is increasing and concave in S), the effect on the price depends on the sign of the expression $\beta D - k((1-\beta)D)$ and is hence ambiguous: Intuitively, a safer product increases the consumers' willingness to pay (demand effect), and it also lowers the innovator's cost for each unit of the smart product sold (liability effect). If the innovator's share of liability is high (i.e. for β large), the liability effects dominates, so that $p^*(S)$ is decreasing in S. By contrast, for β sufficiently low, the demand effect is stronger than the liability effect and $p^*(S)$ is increasing in S.[19] Finally, when product liability becomes more stringent, this leads to a higher price of the smart product and to a lower quantity. Intuitively, a higher β reduces the amount of non-internalized accident costs by the innovator and consumers. For further reference we denote the valuation of the indifferent consumer under optimal pricing as $v^*(S) = v^{ind}(p^*(S), S)$.

Substituting the optimal price $p^*(S)$ and the optimal quantity $q^*(S)$ into (6) yields the reduced-form objective function

$$\Pi = \int_0^\infty e^{-rt}\left[\pi^*(S) - h(x)\right] dt \quad (9)$$

where $\pi^*(S) = q^*(S) \cdot \left(p^*(S) - c - \frac{\alpha(S)\beta D}{r+\rho}\right)$ denotes the expected instantaneous profit which is generated by the (optimal) sales in a given period for a given safety stock S. The intertemporal optimization problem for the innovator is then to maximize (9) w.r.t $x(\cdot)$ under the state dynamics (5) and the constraint $x(t) \geq 0$ for all t. We denote by $V(\tilde{S})$ the value function of the problem, which corresponds to the value of (9) for the initial safety stock $S(0) = \tilde{S}$ under the optimal investment behavior of the innovator. As $\pi^*(S) - h(x)$ does not explicitly depend on time and there is an infinite time horizon, also for any $t > 0$ with $S(t) = \tilde{S}$ the future discounted profit of the innovator under optimal behavior is $V(\tilde{S})$, i.e. the value function is stationary.

[18] Note, however, that the optimal (monopoly) price is not uniform, but will differ across periods, as the safety stock S also varies over time. We abstract from learning effects of the form that the marginal cost of production in a given period is decreasing in the cumulative output in earlier periods, see the discussion in Section 6.

[19] Similarly, such an ambiguous relationship between the optimal price and accident risk also arises in Daughety and Reinganum (1995) for the case of complete information.

3 Optimal safety stock accumulation: active and passive steady states

In this section we characterize the intertemporally optimal investment path of the innovator, combining Optimal Control and Dynamic Programming methods. Using the Maximum Principle (see e.g. Grass et al., 2008) we directly obtain the following necessary optimality conditions:

Lemma 1. *Let $x^*(t)$ be a solution of the innovator's intertemporal optimization problem and $S^*(t)$ the corresponding state trajectory. Then there exists a piece-wise differentiable costate trajectory $\lambda(t)$ such that*

$$x^*(t) = \max\left[\frac{\lambda(t)-\theta}{\theta\eta}, 0\right], \tag{10}$$

$S^*(t)$ evolves according to the state equation given in (5), and $\lambda(t)$ satisfies the costate equation

$$\dot{\lambda}(t) = (r+\delta)\lambda(t) + \alpha'(S^*(t)) \cdot q^*(S^*(t))\frac{\beta D + k((1-\beta)D)}{r+\rho}. \tag{11}$$

Furthermore, the following transversality condition holds:

$$\lim_{t\to\infty} e^{-rt}\left[\pi^*(S^*(t)) - h(x^*(t)) + \lambda(t)(x^*(t) - \delta S^*(t))\right] = 0$$

The optimal investment path x^* is the typical one for a dynamic investment problem with quadratic investment costs. The innovator only invests if the value of the co-state λ is larger than the marginal investment costs at $x = 0$, where $h'(0) = \theta$. Above that level, investment grows linearly with λ. As the co-state λ corresponds to the derivative of the innovator's value function with respect to the safety stock S (see e.g. Grass et al., 2008), this investment rule reflects the trade-off between the marginal return and the marginal cost of investment. The level of investment decreases as the cost function becomes more convex (i.e. as $\theta\eta$ increases). For a given initial condition $S(0)$ the optimal investment path can be written as $x^*(t) = \phi^*(S^*(t))$, where $\phi^*(.)$ is the inter-temporally optimal investment strategy.

In a next step, we characterize potential steady states of the system under optimal investment. We refer to a safety stock S^{ss} as a steady state if it is a fixed point of the state dynamics (5) under the investment strategy $x = \phi^*(S)$.[20] In our analysis we use the following terminology throughout:

Definition 1.

*i) A steady state is called **active** if it exhibits strictly positive levels of both safety stock and production of the smart product (i.e. $S^{ss} > 0$ and $q^*(S^{ss}) > 0$).*

[20]Note that not all fixed points of the canonical system (5) and (11) correspond to steady states in that sense. However, for every steady state there exists a suitable co-state such that the state/co-state tuple is a fixed point of the canonical system. Hence, the set of fixed points of the canonical system provides candidates for steady states.

Figure 1: Isoclines of the canonical system for different values of θ.

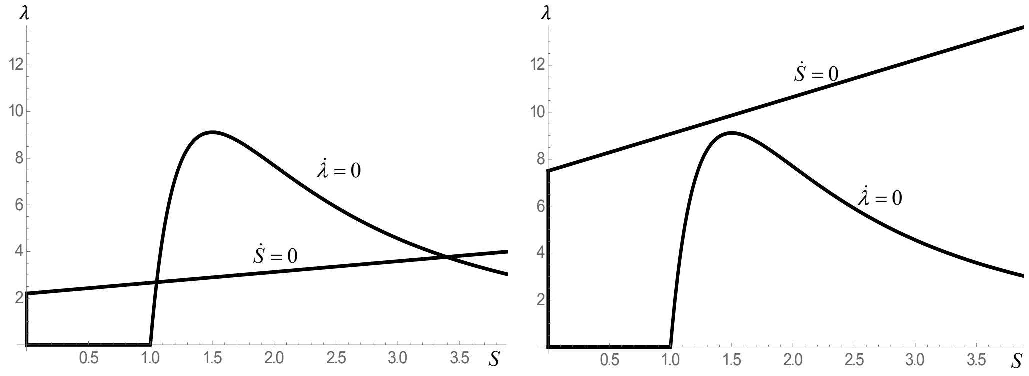

(a) θ low (b) θ high

ii) A steady state is called ***passive*** *if these two levels are zero (i.e.* $S^{ss} = q^*(S^{ss}) = 0$*).*

We only consider these two types of potential steady states, because the two other cases $S^{ss} > 0, q^*(S^{ss}) = 0$ or $S^{ss} = 0, q^*(S^{ss}) > 0$ are irrelevant: The former case can never be optimal because the firm does not gain from a positive safety stock if it does not produce the smart product at all. The latter case is ruled out by the strictly positive minimum safety stock $\underline{S}(\beta) > 0$ required for production to take place, i.e. $q^*(S^{ss}) = 0$ must hold whenever $S^{ss} = 0$.

The candidates for a steady state of the problem can be identified by considering the intersection of the two isoclines $g_\lambda(S)$ and $g_S(S)$ of the canonical system (5) and (11), i.e. setting $\dot{\lambda} = 0$ and $\dot{S} = 0$:

$$g_\lambda(S) \;:=\; -K_1\alpha'(S)q^*(S), \tag{12}$$

$$g_S(S) \;:=\; \begin{cases} \theta\eta\delta S + \theta & \text{for } S > 0 \\ [0,\theta] & \text{for } S = 0, \end{cases} \tag{13}$$

where $K_1 = \frac{\beta D + k((1-\beta)D)}{(r+\rho)(r+\delta)}$.

First, as a strictly positive safety stock $\underline{S}(\beta) > 0$ is required before the smart product is launched, we have $q^*(0) = 0$, and hence $g_\lambda(0) = 0$. This implies that the two isoclines intersect at $S = 0$. Hence, the point $S = 0$ is always a candidate for a (passive) steady state in the sense that under the expectation that $S = 0$ prevails in the long run, it is indeed optimal for the innovator to invest $x = 0$ (and to produce $q = 0$) in any point in time. Clearly, this does not necessarily imply that staying at $S = 0$ is the globally optimal trajectory for the innovator. Furthermore, for $S > \underline{S}(\beta)$ we have $g_\lambda(S) > 0$, as $q^*(S) > 0$, where $\lim_{S\to\infty} g_\lambda(S) = 0$.[21] Taking into account that $g_\lambda(S)$ is independent from θ and that $\lim_{\theta\to 0} g_S(S) = 0$ for all $S > 0$, it follows

[21] The last equality follows because due to $\alpha'(S) < 0$ and $\alpha(S) \geq \alpha^{min}$ for all $S > \underline{S}(\beta)$ we must have $\lim_{S\to\infty} \alpha'(S) = 0$.

that for sufficiently small values of θ there also exist (at least two) candidates for a steady state with $S > 0$. In Figure 1 we illustrate this observation by showing the isoclines of the canonical system for two values of θ. Whereas for a large θ the only candidate for a steady state is the one with $S^{ss} = 0$, for small θ two additional fixed points of the canonical system with positive safety stock exist. This implies that an active steady state can only exist if θ is not too large.

Before continuing with the analysis of steady states, we first establish that under optimal investment the safety stock evolves in a monotonous way:

Lemma 2. *Every optimal path of the innovator's intertemporal optimization problem converges in a (weakly) monotone way to a finite steady state. There exists a finite upper bound $\bar{S}(\theta)$ with $\bar{S}'(\theta) \leq 0$ such that any steady state is in $[0, \bar{S}(\theta)]$ for all $\beta \in [0, 1]$.*

In light of our assumption that the initial safety stock is given by $S(0) = 0$, in what follows we focus on the innovator's optimal behavior under these initial conditions. The following proposition characterizes the conditions for the existence of active and passive steady states and their properties:

Proposition 1. *Assume $S(0) = 0$. There exist thresholds $\underline{\theta} < \bar{\theta}$ with the following properties.*

(i) For $\theta \leq \underline{\theta}$ the optimal trajectory converges to an active steady state with $S^{ss}(\beta; \theta) > \underline{S}(\beta)$ where the innovator produces $q^*(S^{ss}) > 0$ for all $\beta \in [0, 1]$.

(ii) For $\theta \geq \bar{\theta}$ the innovator implements a passive steady state by optimally choosing $x^*(t) = 0$ and $q^*(0) = 0$ for all t and for all $\beta \in [0, 1]$.

(iii) For $\theta \in (\underline{\theta}, \bar{\theta})$ there exists a threshold $\hat{\beta}(\theta) \in (0, 1)$ such that for $\beta > \hat{\beta}$ the innovator implements the passive steady state by optimally choosing $x^*(t) = 0$ for all t. For $\beta < \hat{\beta}$ the optimal trajectory converges to an active steady state with $S^{ss}(\beta; \theta) > 0$ where the innovator produces $q^*(S^{ss}) > 0$. For $\beta = \hat{\beta}$ the innovator is indifferent between the optimal path leading to the active steady state with $S^{ss}(\hat{\beta}; \theta)$, and staying at the passive one.

The first two parts of the proposition show that for sufficiently high or low values of the cost parameter θ, whether an active or a passive steady state is reached is independent of the innovator's liability share (β). For S sufficiently large, the innovator's profit from the smart product is always positive. Hence, in the absence of any investment costs, the firm would always invest in building up a safety stock, independent of the liability regime. By continuity, there exists a threshold $\underline{\theta}$ such that this also holds for sufficiently small values of $\theta \leq \underline{\theta}$. Furthermore, as discussed above, a sufficiently large value of the cost parameter θ implies that $S = 0$ is the only candidate for a steady state for all $\beta \in [0,1]$. Hence, for $\theta \geq \bar{\theta}$ it is not optimal for the firm to invest in the smart product.

The most interesting scenario occurs for intermediate values of θ, as described in part (iii): In the absence of liability (i.e. $\beta = 0$), the active steady state is stable and the path leading from $S(0) = 0$ to this stable steady state generates a strictly positive discounted payoff stream. This is larger than the payoff of zero generated at the passive steady state $S^{ss} = 0$. As β increases, the value of the path leading to the active steady state becomes smaller. Hence, for $\theta \in (\underline{\theta}, \bar{\theta})$ there exists a threshold $\hat{\beta} \in (0,1)$ such that for this value the optimal path leading from $S(0) = 0$ to the active steady state yields a discounted payoff stream of zero. Therefore, only for $\beta \leq \hat{\beta}$ is it optimal for the innovator to pursue the path to the active steady state.[22] By contrast, for all $\beta > \hat{\beta}$, the expected liability cost is excessively high such that the innovator optimally forgoes the development of the smart product.

4 Dynamic patterns of product safety and market evolution: The countervailing effects of liability

In Proposition 1, we have characterized the steady state properties of our model (i.e. the long run attractor from $S(0) = 0$), and the impact of the stringency of liability. In this section, we focus on several crucial aspects of the emerging dynamics. Due to the non-linear nature of the firm's investment problem, it is infeasible to obtain a closed-form solution for the optimal investment strategy. Consequently, we perform a numerical analysis of a calibrated version of the model. After explaining our calibration approach (Section 4.1), we show how the innovator's liability share β affects the active steady state (Section 4.2). We then discuss the dynamic patterns of optimal investment and the corresponding market dynamics (Section 4.3), and how β affects the timing of AV market introduction and the evolution of AV safety (Section 4.4). Finally, we study welfare implications (Section 4.5).

4.1 Calibration for US automobile market

To obtain an empirically meaningful parametrization of the model, we focus on the US automobile market and consider the development of Autonomous Vehicles (AV, the smart product) as an alternative to

[22] Considering the problem from the perspective of the state space, for any $\theta \in (\underline{\theta}, \bar{\theta})$ and $\beta > \hat{\beta}$ there is a positive threshold $S^{SK} > 0$ such that the optimal trajectory for any $S(0) < S^{SK}$ leads to the passive steady state $S^{ss} = 0$, whereas for $S(0) > S^{SK}$ the optimal path leads to the active steady state S^{ss}. Such a threshold S^{SK} is referred to as a Skiba point (see e.g. Skiba, 1978; Hinloopen, Smrkolj, and Wagener, 2013). In this interpretation the value $\hat{\beta}$ is defined as the point where the Skiba point coincides with $S = 0$.

conventional cars (the old product).[23] For most of our model parameters we were able to obtain appropriate empirical targets (reference year 2017). For the remaining parameters, where such targets cannot be obtained, we have chosen values leading to reasonable outcomes for the timing of AV market introduction.

As for the discount rate, we follow Blincoe et al. (2015) and use $r = .03$. Moreover, Bento, Roth, and Zuo (2018) find a current average lifetime of cars in the US of 15.6 years, leading to a failure rate $\rho = .065$, which we assume to apply for both types of car.

As for accident rates, for conventional cars we set $\alpha_o = .074$. This is obtained from US data from the Insurance Information Institute, reporting a total of 7.4 claims per year per 100 insured cars.[24] For the AV, we use the functional form $\alpha(S) = \bar{\alpha}/(m + S)$, which satisfies all properties for $\alpha(S)$ as assumed in Section 2, and where we set $\bar{\alpha} = 0.12$ and $m = 0.05$.

For the damage in the course of an accident, we use $D = 39.2$ which is calculated as follows. According to a report by the U.S. National Highway Traffic Safety Adminstration (Blincoe et al., 2015), the total social costs due to car accidents in 2010 were \$836 billion with a total of 24 million damaged vehicles. Taking also into account price inflation between 2010 and 2017 (12.4 % in total), we estimate the average damage per damaged vehicle per accident in 2017 as $(836.000 \cdot 1.124)/24 = 39.152$. As for the share of the expected damage borne by the consumers, we use $k(z) = \frac{\tau \cdot z}{\tau + z}$ which satisfies all the properties as assumed in Section 2, where we set $\tau = 50.000$.

According to data from Kelley Blue Book, the average transaction price for cars in 2017 was 34.782\$ and, hence, $p_o = 34.8 = c$.[25] As for consumer valuations, we could not find any direct evidence for the US market. Instead, we use data reported by Brenkers and Verboven (2006) in their empirical study of several EU automobile markets, and we find that for the case of Germany, on average the consumer surplus equals the price of the car.[26] In the numerical analysis, we assume that the same property also holds for the US. Based on this, the consumer surplus can be derived as the difference between the willingness to pay (per period of usage) net of expected liability costs and the sales price (i.e. $\frac{v_o - \alpha_o k(D)}{r+\rho} - p_o$), which leads

[23] Empirical studies of the automobile markets in the U.S. and in Europe have been carried out for example by Goldberg and Verboven (2001), Brenkers and Verboven (2006) and Kagawa, Tasaki, and Moriguchi (2006).

[24] See `https://www.iii.org/fact-statistic/facts-statistics-auto-insurance`.

[25] Kelley Blue Book is a leading source of price information for the US car market. They issue monthly press releases with information on transaction prices for the different manufacturers and models. Our value of 34.782\$ is obtained from taking the average of the prices reported for the 12 months of the year 2017. That also $c = 34.8$ holds follows from our model assumption that conventional cars are sold at marginal costs.

[26] Brenkers and Verboven (2006) report that a loss in consumer surplus of 444 Mio Euro corresponds to 1.7% of the average annual consumer surplus generated in the market for new cars in Germany between 1970-1999 (see their Table 6). Hence, the average annual consumer surplus during that time period can be calculated as 26.1 Billion Euro. Dividing this number by the average annual car sales in Germany of 2.48 Mio during that period (obtained from their Table 1), we obtain a consumer surplus per sold car of 10.531 Euro. This number almost coincides with the average sales price of 10.520 Euro in Germany during the same period (also obtained from their Table 1).

to $v_o = 8.2$. With respect to consumer valuations for the AV, we do not consider those consumers, whose expected discounted value derived from an AV is negative even if the AV had an accident rate of zero and would be sold at the price of the conventional car. This implies $\underline{v} = p_o(r + \rho) = 3.3$. To approximate the market size for AVs, we follow the seminal work on the diffusion of innovations by Rogers (2003), who argues that 16% of the consumers can be classified as either *innovators* or *early adopters*, which are the most eager ones to switch to a new product. Hence, in our numerical analysis, we assume that 16% of the mass of the (uniform) distribution of valuations on the interval $[\underline{v}, \bar{v}]$ is above $v_o (= 8.2)$. This leads to $\bar{v} = 9.2$.

Finally, with respect to the safety stock S of the AV, we assume that it depreciates with a rate $\delta = 0.07$. Moreover, as for the cost of investment into that stock, $h(x) = \theta x(1 + \frac{\eta}{2}x)$, we use $\theta = 2.2$ and $\eta = 3$. This parameter choice ensures that there indeed exist a liability parameter $\hat{\beta} \in (0, 1)$, such that the innovator prefers the optimal path leading to the active steady state for $\beta \leq \hat{\beta}$, and staying at the passive steady state for $\beta > \hat{\beta}$.[27]

In order to determine the optimal paths for our calibrated model, we numerically determine the value generated by the optimal path from $S(0) = 0$ leading to the candidate for an active steady state (if it exists), and compare it to the value of zero from staying at the passive steady state. The optimal path leading to the candidate for the active steady state and the associated value are calculated by solving the Hamilton-Jacobi-Bellman (HJB) equation of a variation of the innovator's dynamic optimization problem, in which the additional constraint is added that the state converges to the active steady state. The function $\tilde{V}(S)$ solving this HJB equation is evaluated at the initial state $S(0) = 0$ and we obtain the value $V(0) = \max[0, \tilde{V}(0)]$ at this state. The scenarios for which positive investment and convergence to an active steady state are optimal for the innovator are therefore characterized by the condition $\tilde{V}(0) \geq 0$. As an analytical solution for the HJB equation is not available, we obtain an (approximate) solution to the equation using a collocation method. The formulation of the HJB equation and a short description of our method is given in Appendix B.

4.2 Active and passive steady states of safety stock accumulation

We start with analyzing the impact of the stringency of product liability (β) on the existence and properties of active and passive steady states as characterized in Proposition 1 above. As illustrated in Figure 2, when the innovator's liability share (β) is sufficiently low (i.e. when $\beta \leq \hat{\beta}$), an *active* steady state emerges in which the innovator accumulates a positive safety stock $S^{ss}(\beta) > 0$. Moreover, the steady state level of the

[27] Put differently, our specification ensures that the isoclines (12) and (13) intersect twice as in panel (a) of Figure 1, and that we are in part (iii) of Proposition 1.

Figure 2: AV safety stock in *active* and *passive* steady states under strict (partial) liability

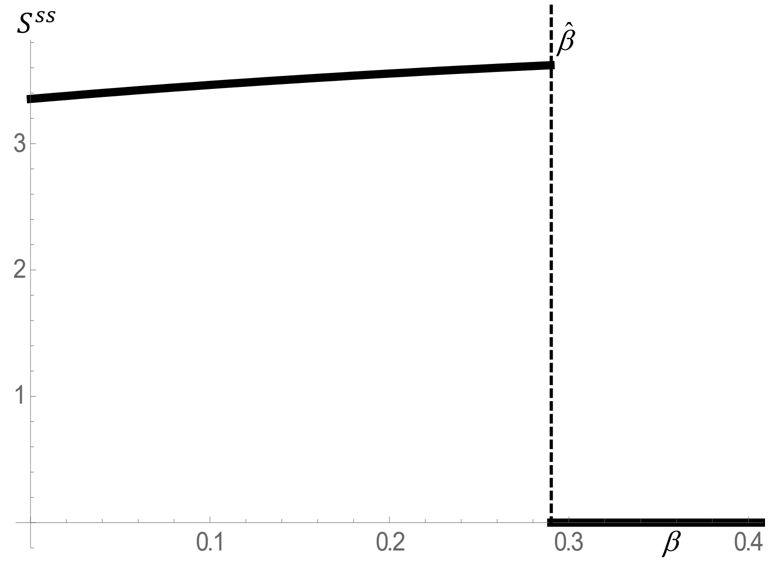

safety stock is increasing in β. By contrast, for high values of $\beta > \hat{\beta}$, the *passive* steady states emerges, where the innovator (with a starting point of $S = 0$) will optimally not accumulate any safety stock over time (i.e. $S^{ss}(\beta) \equiv 0$ for all $\beta > \hat{\beta}$). For our calibration, the critical value above which the active steady state ceases to be optimal is $\hat{\beta} = 0.29$. Hence, this property illustrates one trade-off when shifting more liability to the producers of smart products: Given that these products will eventually be launched at all, more liability indeed increases the level of product safety in the long-run, but it might also discourage the activity altogether. Overall, these results show that the relationship between liability and product innovation can in principle be positive or negative, and that it depends (among others) on the characteristics of industry under consideration and the level of liability actually chosen. This is in line with the empirical findings of McGuire (1988); Viscusi and Moore (1993); Galasso and Luo (2017, 2018b) as discussed above.

4.3 Optimal dynamics

From now on, we focus on liability regimes $\beta \leq \hat{\beta}$, such that optimal investment by the innovator leads to convergence to the active steady state. In the following, we analyze in more detail the dynamic properties of the key variables of interest, and how these trajectories are affected by the stringency of liability. The results are illustrated in Figure 3, where we depict trajectories for $\beta = 0, \beta = 0.15$ and $\beta = \hat{\beta} = 0.29$. Lemma 2 directly implies that the safety stock grows monotonously over time under any optimal investment path leading to an active steady state. Hence, understanding how the key variables depend on S also allows to understand their behavior over time.

Figure 3: AV market dynamics under optimal investment under strict (partial) liability

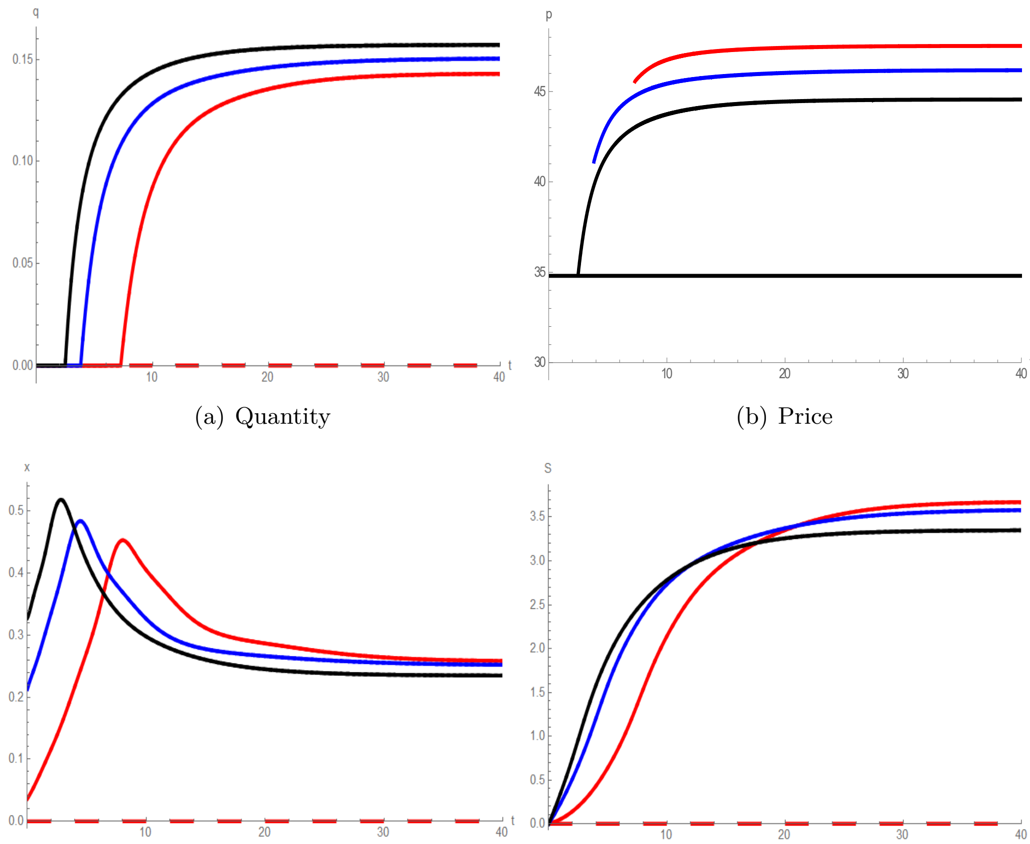

(a) Quantity

(b) Price

(c) Investment

(d) Safety stock accumulation

The solid lines of different colors show the optimal paths for different values of firm liability: $\beta = 0$ (black), $\beta = 0.15$ (blue), $\beta = 0.29$ (red). For the threshold value $\hat{\beta} = 0.29$ both the path converging to the active steady state (solid red line) and the one staying in the passive steady state (dashed red line) are optimal. In panel (b), the three depicted paths show the optimal price for the AV starting from the respective period for which the optimal quantity $q^*(S)$ becomes strictly positive. The horizontal black line indicates the price of the conventional car ($p_o = 34.8$).

4.3.1 AV production quantities and prices

Production quantities As can be seen from Figure 3 (a), the AV is not immediately launched. Rather, there is an initial time interval, during which the innovator does not produce and only invests to improve the AVs' safety (this issue will be discussed on more detail in Section 4.4 below). Once production does take place, the (instantaneous) quantity is indeed increasing and concave over time for all three depicted liability regimes (see the discussion of Eqn. (8) above). Intuitively, as the safety stock increases over time (see panel (e)), the AV becomes safer which, ceteris paribus, decreases the innovator's unit costs and increases consumers' net benefit from it. Both effects lead to higher sales. Moreover, more stringent liability does not affect the dynamics of output decisions qualitatively. It does, however, induce a later date of market introduction of the AV and lower production quantities afterwards, and hence a slower market penetration.

Prices Figure 3(b) shows that the price for the AV optimally chosen by the innovator increases over time as the safety stock increases. Recall from the discussion of Eqn. (7) above, that the effect of S on the optimal price of the AV $p^*(S)$ is a priori ambiguous, and that it is increasing in S for β sufficiently small, in which case the demand effect dominates the liability effect. In our calibration, it turns out that this is indeed the case in the relevant parameter range $\beta \in [0, \hat{\beta}]$ where the innovator optimally induces the active steady state. Moreover, as product liability become more stringent, this leads to higher prices.

4.3.2 Investment pattern and accumulation of safety stock

Figure 3(c) reveals a hump-shaped investment pattern over time. Again, this process is affected by the stringency of liability: For low values of β, initial investments are high, with a peak being reached already in early periods, and then exhibiting a relatively strong decline. By contrast, for larger values of β, the initial investment outlays are lower, the peak is reached at a later date, and the decline is less severe. As a result, the more stringent liability, the lower (higher) the investment incentives in the short-run (intermediate and long run). In the remainder of this subsection, we investigate the drivers for these results in more detail.

Basic effects determining optimal investment Several effects are at work which jointly determine the dynamic investment pattern. To gain an intuition, recall from Lemma 1 that the optimal investment is linearly increasing with the co-state λ (see Eq. (10)). To illustrate these effects, in what follows we approximate the co-state by the derivative of the innovator's (discounted future) market profits with respect to the safety stock S. Formally, we define $\tilde{\Pi}^*(\tilde{S}) = \int_0^\infty e^{-rt}\pi^*(S(t))dt$ with $S(t)$ following the path induced by optimal investment from $S(0) = \tilde{S}$ and approximate $\lambda(t)$ by $\frac{d\tilde{\Pi}^*(\tilde{S}(t))}{d\tilde{S}}$.[28] Using this approximation, the optimal investment level is linearly increasing in the term $\frac{d\tilde{\Pi}^*}{d\tilde{S}}$, and we consider this term to gain an understanding of the factors determining optimal investment. Taking into account the envelope theorem, it can be written as follows:

$$\begin{aligned}\frac{d\tilde{\Pi}^*(\tilde{S})}{d\tilde{S}} &= \int_0^\infty e^{-rt}\left[\frac{\beta D}{r+\rho}q^*(S(t)) + \left(p^*(S(t)) - c - \frac{\alpha\beta D}{r+\rho}\right)\frac{k((1-\beta)D)}{\bar{v}-\underline{v}}\right]\cdot\left|\frac{\partial\alpha(S(t))}{\partial\tilde{S}}\right| dt \\ &\approx \Big[\underbrace{\frac{\beta D}{r+\rho}Q^*(\tilde{S})}_{\substack{\text{Profit Margin Effect}\\ \text{(PME)}}} + \underbrace{\left(P^*(\tilde{S}) - \frac{c}{r+\delta}\right)\frac{k((1-\beta)D)}{\bar{v}-\underline{v}}}_{\substack{\text{Quantity Effect}\\ \text{(QE)}}}\Big]\cdot\underbrace{\left|\frac{\partial\alpha(\tilde{S})}{\partial\tilde{S}}\right|}_{\substack{\text{Safety Effect}\\ \text{(SE)}}},\end{aligned} \tag{14}$$

where $Q^*(\tilde{S}) = \int_0^\infty e^{-(r+\delta)t}q^*(S(t))dt$ and $P^*(\tilde{S}) = \int_0^\infty e^{-(r+\delta)t}\left(p^*(S(t)) - \frac{\alpha(S(t))\beta D}{r+\rho}\right)dt$ denote aggregated discounted future sales quantities and net prices (i.e. prices minus liability costs) for an initial safety stock $\tilde{S}$. Note that to obtain the second line we have used the approximation $\frac{\partial\alpha(S(t))}{\partial\tilde{S}} \approx e^{-\delta t}\frac{\partial\alpha(\tilde{S})}{\partial\tilde{S}}$ based on the fact that the safety stock depreciates at a rate of δ.

Intuitively, the multiplicative third term captures the direct effect of an increase S, namely a reduction in the accident rate of the AV (SE). The first two terms capture the effect of a decrease in the accident rate on the profit margin (PME) and the quantity of AV sold (QE), respectively: As for the PME, for a given aggregate quantity Q^*, a lower $\alpha(S)$ increases the profit margin of the innovator for each unit sold now and in the future, as it leads to a decrease of liability costs. As for the QE, the demand function (4) shifts upwards by $\frac{k((1-\beta)D)}{\bar{v}-\underline{v}}$ as α decreases. To obtain the total marginal profit of this quantity increase, it is multiplied with the discounted sum of the marginal profits over time.[29]

Investment incentives over time To understand the dynamics of the optimal investment over time, we analyze how the different effects in Eq. (14) vary in S (i.e. we now look at $\frac{\partial^2\tilde{\Pi}^*}{\partial\tilde{S}^2}$). The following three (countervailing) channels drive the optimal investment pattern for the time *after* the AV has been launched (i.e. when $q^* > 0$). First, the PME becomes stronger because the discounted sum of future sales, $Q^*(\tilde{S})$, increases with $\tilde{S}$. Hence, this channel leads to an increasing investment incentive over time. Second, also the

[28] To clarify our approximation, recall first that the co-state λ corresponds to the derivative of the value function with respect to S. The value function is given by $V(\tilde{S}) = \int_0^\infty e^{-rt}(\pi^*(S(t) - h(x^*(t))dt$ with $S(t)$ following the path induced by the optimal investment $x^*(t)$) from $S(0) = \tilde{S}$. Hence, our approximation $\lambda = \frac{\partial V(\tilde{S})}{\partial\tilde{S}} \approx \frac{\partial\Pi^*(\tilde{S})}{\partial\tilde{S}}$ holds under the simplifying assumption that $\frac{\partial}{\partial\tilde{S}}\int_0^\infty e^{-rt}h(x^*(t))dt \approx 0$, i.e. when abstracting from the fact that an increase of S might also have an impact on the future investment paths and the associated discounted costs.

[29] The change of the optimal price induced by an decrease in α and the resulting quantity change do not affect the maximized profit because of the envelope theorem.

QE becomes stronger, because an increase in S leads to both a higher price (see Figure 3(b)) and a lower (expected) liability cost per unit sold (as α decreases with S) in all future periods. Hence the term $P^*(\tilde{S})$ increases over time and also this channel generates growing investment incentives. Third, the SE generates a decreasing investment incentive over time, because the marginal effect of a higher safety stock on the accident rate decreases over time (recall that $\alpha'(S) < 0$ and $\alpha''(S) > 0$). This term is multiplicative such that it scales the size of the two previous channels. For the time period *before* the market introduction of the AV (i.e. when $q^* = 0$), a fourth channel is at work. During this time, the innovator seeks to accumulate, at minimum cost, the level of safety stock $\underline{S}(\beta)$ with which the AV is launched. Note that the discounted profit generated after the introduction of the AV does not depend directly on the time of the market introduction, but only on the value of S. For a given date of market introduction, in the absence of discounting it would be optimal to spread the investment outlays evenly over the time period before the AV is launched (this due to the convexity of the investment cost function $h(x)$). However, with discounting, the value of this fixed future profit becomes larger as the date of market introduction comes closer. This explains the initially increasing part of the investment trajectories before the AV is launched. The interplay of these four effects generates the observed hump-shaped pattern.

Impact of liability on investment pattern Equation (14) is also useful to develop an intuition of how the optimal investment trajectories are affected by the stringency of liability (β). Note first that, ceteris paribus, an increase in β does not affect the SE. With respect to the PME and QE, we obtain the following derivative with respect to β:

$$
\begin{aligned}
\frac{\partial [PME + QE])}{\partial \beta} &= \underbrace{\frac{D}{r+\rho} Q^*(\tilde{S})}_{>0} + \underbrace{\frac{\beta D}{r+\rho} \frac{\partial Q^*(\tilde{S})}{\partial \beta}}_{<0} + \underbrace{\frac{\partial P^*(\tilde{S})}{\partial \beta} \frac{k((1-\beta)D)}{\bar{v} - \underline{v}}}_{<0} \\
&\quad + \underbrace{\left(P^*(\tilde{S}) - \frac{c}{r+\delta} \right) \frac{-Dk'((1-\beta)D)}{\bar{v} - \underline{v}}}_{<0} . \qquad (15)
\end{aligned}
$$

Intuitively, the first of the four terms captures that a higher β fortifies the impact of changes in α on the profit margin (i.e. the coefficient of Q^* increases with β). This leads the innovator to respond with higher investment, and the effect is proportional to quantity. As for the second term, in light of (8), there is a direct negative effect of β on the accumulated discounted future quantities ($\frac{\partial Q^*(\tilde{S})}{\partial \beta} < 0$). This reduces the size of the PME and dampens investment incentives. Note that this argument abstracts from the indirect effect that a change in β also affects the optimal path of $S(t)$, which also influences the accumulated discounted quantities and prices. Similarly, with respect to the third term, abstracting from this indirect effect and using (7) we obtain

$$\frac{\partial P^*(\tilde{S})}{\partial \beta} = \frac{D(k'((1-\beta)D)-1)}{2(r+\rho)} \int_0^\infty e^{-(r+\delta)t} \alpha(S(t))dt < 0. \tag{16}$$

This expression is negative as a higher β leads to higher unit costs, which is not fully offset by the increase in demand resulting from the lower remaining liability share $(1-\beta)D$ borne by consumers (this is due to $k'(\cdot) < 1$). As a result, an increase in β induces a lower profit margin (price minus liability costs) for the innovator, which lowers investment incentives. Finally, as for the last term in (15), shifting liability from the consumers to the innovator, makes demand react less strongly to an increase of the safety stock, thereby reducing investment incentives.

The interplay of these four effects as captured in (15) explains why the qualitative implications of a more stringent product liability differs between the short-term and the long-term (see panel (c) of Figure 3). Note that only the first term in (15) generates a positive effect of β on investment incentives, which is proportional to the discounted stream of future sales (Q^*). In early periods, Q^* is small and therefore this positive effect is dominated by the three negative effects. This explains why initially investments are lower for larger values of β. Eventually, in later periods Q^* becomes sufficiently large such that the first positive effect dominates, and we obtain a positive relationship between β and safety stock investment.

Finally, it is noteworthy that in the absence of liability ($\beta = 0$), the innovator's investment incentives are not zero, even in the short-run. The reason is that in this case, a substantial liability risk would be borne by the consumers. This would lower their willingness to pay for AVs, thereby shrinking the market for these vehicles. To countervail this effect, the innovator does optimally invest in product safety, even when she does not face a direct liability risk. These are exactly the "market forces" stressed by Polinsky and Shavell (2010) in their critical discussion of product liability.

Accumulation of safety stock The dynamic patterns of investment incentives determine how the stringency of product liability affects the process of safety stock accumulation. In line with our observation that in the long run an increase of β induces higher investment, Figure 3(d) shows that the safety stock ultimately reached in the active state steady is increasing in the innovator's liability share β. However, panel (e) also conveys that, as β increases, the accumulation of the safety stock occurs more slowly in early periods, due to the lower initial investment incentives. For example, during the first 8 years, S is largest for $\beta = 0$ followed by $\beta = 0.15$ and $\beta = 0.29$ and it takes 20 years until the safety stock is highest under the most stringent form of liability. This effect is illustrated in more detail in Figure 4. It shows that the time elapsing until the AV eventually reaches the same safety standard (accident rate) as the conventional car, is increasing and convex in the stringency of liability. That is, as long as β is relatively small, making the liability more

Figure 4: Delay in AV safety stock accumulation under strict (partial) liability

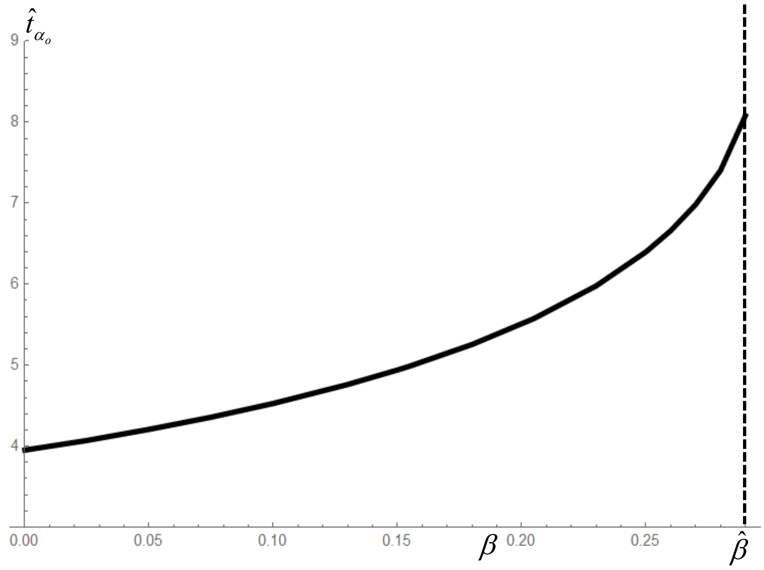

The vertical axis measures the time ($\hat{t}_{\alpha_o}$) until the accident rate of the AV equals that of the conventional car, i.e. $\alpha(S(\hat{t}_{\alpha_o})) = \alpha_o = 0.074$.

stringent only causes a moderate delay in the accumulation process compared to the benchmark $\beta = 0$. But the delay becomes considerable as β moves closer towards the threshold level $\beta = 0.29$, for which the active steady state ceases to be optimal for the innovator.

In summary, even when confining attention to liability regimes which are not too stringent in the sense of preventing the introduction of the smart product altogether, our analysis points to a second trade-off associated with shifting more liability to the producers of smart products: One the one hand, this leads to a higher safety level of these products in the long run. But on the other hand, it delays the process of safety stock accumulation, so that it takes longer until a given safety standard is reached. As a result, under more stringent liability, the units being sold in the short- and intermediate run are less safe than they would have been under a less stringent regime.

4.4 Timing of AV market introduction, and safety level of first generation

In this section we investigate in more detail the impact of liability on the date of market introduction of the AV (denoted by t^{en}) and the corresponding safety level $S(t^{en})$ of this first AV generation. Recall that the innovator will optimally launch the AV when it has just reached the minimum safety standard such that demand becomes positive, i.e. when $S(t^{en}) = \underline{S}(\beta)$ holds (see Eq. (3) above). Moreover, a more stringent liability leads to a higher the safety stock of the first generation (i.e. $\underline{S}'(\beta) > 0$, which follows from $\alpha' < 0$ and the properties of $k(\cdot)$). Furthermore, as discussed in the previous section, in early periods an increase

Figure 5: Date of AV market introduction and corresponding safety standard

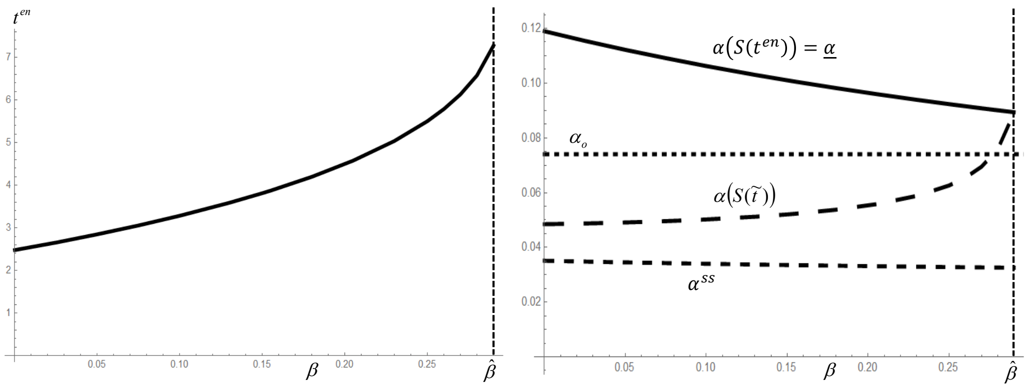

(a) Date of AV market introduction (b) AV accident rate

Panel (a): $t^{en}(\beta)$ indicates the point in time at which the AV is launched (i.e. where the optimal instantaneous quantity becomes strictly positive) for a given β. Panel (b): accident rate of AV i) in the first generation ($\alpha(S(t^{en})) = \underline{\alpha}(\beta)$), ii) at a given point in time $\tilde{t}$ ($\alpha(S(\tilde{t}))$) and iii) in the active steady state ($\alpha^{ss}(\beta)$). The dotted line shows the accident rate of the conventional car ($\alpha_o = 0.074$).

of β has a negative effect on investment such that it takes longer until any given level of safety stock S is reached. Due to these two effects, more stringent liability leads to a later market introduction of the AV (see Figure 5(a)).

Figure 5(b) shows that, whereas β has a beneficial effect on the AV accident rate at the time of market introduction and in the long run (steady-state), there is also a negative effect in the short run. If we consider a given point in time $t = \tilde{t}$ not too far in the future, then the accident rate at $\tilde{t}$ is increasing in β.[30] Intuitively, when β is low, the AV is launched earlier and the innovator's investment incentives increase once it sells positive quantities. In particular, the difference between $\alpha(S(t^{en}))$ and $\alpha(S(\tilde{t}))$ shows by how much the accident rate is reduced due to the earlier market introduction for a low β compared to the case of $\beta = \hat{\beta}$. Put differently, if a policy maker aims at maximizing AV safety at a given point in time not too far in the future, then lenient product liability ($\beta = 0$) is optimal. However, the policy maker should be aware that under such a regime, AVs are introduced to the market early and the first product generations exhibit a safety level that is substantially lower than first generation AVs under more stringent liability. Figure 5(b) also reveals that for a large range of β the accident rate at $t = \tilde{t}$, is already below that of the conventional car. Last, but not least, the figure also illustrates that for our calibration, the long run AV accident rate ($\alpha^{ss}(\beta)$) is below that of the conventional car (α_o) for all β. This is consistent with the alleged long-term benefit of AVs compared to conventional cars.

[30] For the purpose of illustrating this effect, in the figure we have set $\tilde{t} = 7.29$, which is the date of market introduction of the AV for $\beta = \hat{\beta}$.

4.5 Welfare

Having studied the implications of an increase in β on the timing of market introduction of AVs and the evolution of the safety level, we now consider the total welfare effects of product liability. The welfare generated by each vehicle is given by the expected lifetime utility for the buyer minus the sum of production costs and total expected damages.[31] As we are interested in the welfare implications of AVs, in what follows we consider the additional welfare generated by their introduction relative to the status quo where they are not available.[32] This leads the following measure ΔW for the (instantaneous) welfare effect of AVs:

$$\Delta W = \int_{v^*(S)}^{\bar{v}} \left(\frac{v - \alpha(S)D - (v_o - \alpha_o D)}{r + \rho} \right) \frac{1}{\bar{v} - \underline{v}} dv - h(x), \tag{17}$$

where the first term measures the additional welfare generated by those consumers who prefer to purchase the AV instead of the conventional car. Recall that the threshold $v^*(S) = \bar{v} - q^*(S)(\bar{v} - \underline{v})$ is determined by the valuation of the consumer who is just indifferent between purchasing either type of car, where $q^*(S)$ is the optimal quantity of the innovator as given in (8) above.

From a social welfare perspective, a consumer should purchase the AV whenever this generates a positive net effect on welfare. This leads to a valuation threshold $v^{FB}(S) = v_o - \alpha_o D + \alpha(S)D$ and hence to the first-best quantity

$$q^{FB}(S) = 1 - F(v^{FB}(S)) = \frac{\bar{v} - v_o + (\alpha_o - \alpha(S))D}{\bar{v} - \underline{v}}. \tag{18}$$

Taking the difference yields

$$q^{FB}(S) - q^*(S) = \underbrace{\frac{\bar{v} - v_o + (\alpha_o - \alpha(S))D}{2(\bar{v} - \underline{v})}}_{\text{monopoly distortion}} + \underbrace{\frac{\alpha_o(D - k(D)) - \alpha(S)\left((1-\beta)D - k((1-\beta)D)\right)}{2(\bar{v} - \underline{v})}}_{\text{externality distortion}}. \tag{19}$$

This reveals that at a given point in time t, there are two sources which distort the quantity of AVs sold relative to the efficiency benchmark: First, there is the standard downwards distortion due to the monopoly power of the innovator (*monopoly distortion*). This distortion is independent of the liability parameter β.

Second, as discussed in Section 2, both the AV and the conventional car generate externalities which are due to the incomplete internalization of the damage from accidents (*externality distortion*). Depending on which of these two externalities is larger, this creates an upwards or downwards distortion of the AV

[31] As for the damage resulting from accidents, note that simply adding up the liability shares of the respective consumer and the innovator (i.e. $\beta D + k((1-\beta)D)$) would leave unaccounted for the remaining share, which has to be borne by a third party such as, for example, another driver involved in the accident or an insurance company. To avoid this, our welfare measure captures the total social damage D.

[32] Hence, we consider the difference in welfare between the cases with and without AVs, respectively. As in the latter case, the resulting welfare is independent of S and β, this does not have any qualitative effect on the results.

quantity relative to the first best. In particular, when all liability costs of AVs are borne by the innovator (i.e. $\beta = 1$, so that the last term in (19) is zero), the externality distortion only arises through the conventional car. In this case, both the monopoly and the externality distortion go in the same direction, and the AV optimal quantity $q^*(S)$ is inefficiently small. By contrast, when consumers are liable for all AV accidents ($\beta = 0$), then the externality distortion becomes negative if AVs are less safe than conventional cars (i.e. if $\alpha(S) > \alpha_o$). If this effect outweighs the (positive) monopoly distortion, then $q^*(S)$ is too high compared to the efficiency benchmark, i.e. $v^*(S) < v^{FB}(S)$. In such a case launching the AV might have negative welfare implications even if S is sufficiently large such that $v^{FB}(S) \leq \bar{v}$. In particular, AV consumers with $v \in [v^*(S), v^{FB}(S))$ reduce welfare, whereas those with $v \in (v^{FB}(S), \bar{v}]$ increase it. We define S^{wel} as the minimal level of the safety stock such that, from a static perspective, the provision of AVs enhances welfare (i.e. ΔW as given in (17), net of investment costs is non-negative). For our parameterization, we obtain $S^{wel} = 1.19$ and $\alpha(S^{wel}) = 0.096$. These values will be a useful benchmark for the analysis of negligence standards in Section 5 below.

An overall assessment of the welfare effects of product liability also needs to take into account the dynamic aspects arising through the induced investment incentives into product safety. To analyze these dynamic aspects, we proceed in two steps: First, we compare the optimal investment path of the innovator with the welfare-maximizing path for a *given* level of β.[33] Second, we consider the welfare effect of a change of β under the innovator's optimal path. Throughout, we fix the AV quantity at $q^*(S)$ which allows for a clean comparison of the individually and socially optimal investment paths.

Figure 6: Socially optimal investment and quantity paths

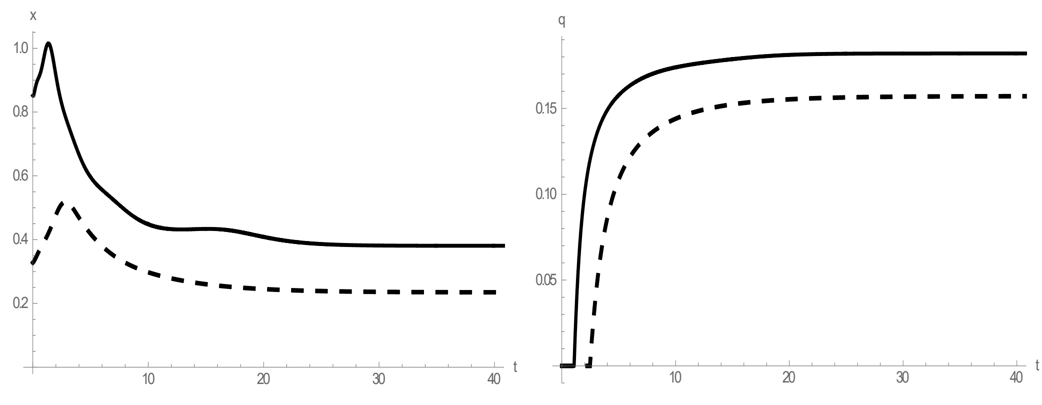

(a) Investment (b) Quantity

Dynamics of investment (panel (a)) and quantity (panel (b)) under the welfare maximizing path (solid line) and the optimal path of the innovator (dashed line) under strict (partial) liability for $\beta = 0$.

[33]Formally, we consider the objective function $\int_0^\infty e^{-rt} \Delta W(S(t), x(t))dt$ subject to the state dynamics (5), where $\Delta W(S(t), x(t))$ is given by (17).

Figure 7: The welfare effect of liability in the active steady state

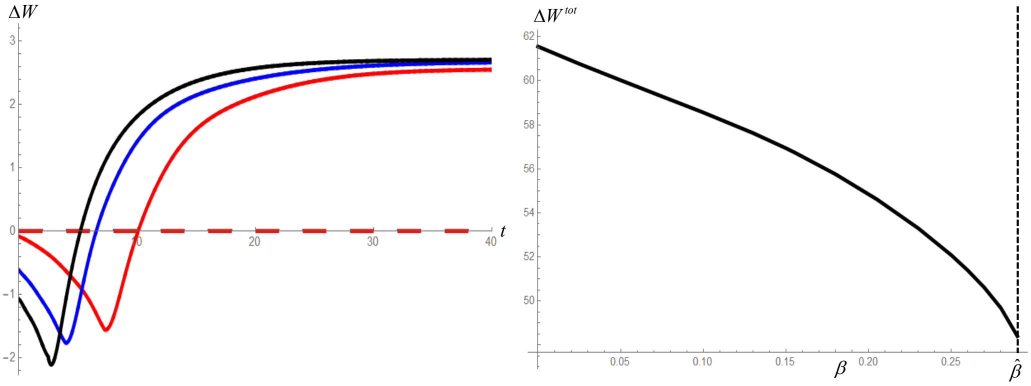

(a) Welfare dynamics

(b) Total welfare effect

Panel (a): Dynamics of instantaneous welfare effect for $\beta = 0$ (black), $\beta = 0.15$ (blue) and $\beta = \hat{\beta} = 0.29$ (red). Panel (b): Total discounted welfare effect ΔW^{tot} for $\beta < \hat{\beta}$.

As for the first step, Figure 6 illustrates that the socially optimal path exhibits substantially larger safety stock investments compared to the innovator's optimal path under $\beta = 0$. As the safety stock at market introduction is $\underline{S}(0)$ under both paths, the higher investment leads to earlier market introduction of AVs under the socially optimal path. Furthermore, at each point in time after market introduction, the quantity is higher under the socially optimal investment level. Overall, this comparison suggests that from a welfare perspective, it is desirable to speed up investments into the safety stock of AVs and to increase their sales volume.

As for the second step, the results for the dynamic analysis are illustrated in Figure 7. Panel (a) shows that instantaneous welfare exhibits a U-shaped pattern over time. Intuitively, it is initially negative due to the investment outlays before the AV is launched, and then eventually becomes positive afterwards. Moreover, more stringent liability leads to a more balanced pattern (i.e. to lower initial welfare losses, but for a longer period of time, and lower gains in later periods). In particular, the long-run welfare level decreases with β, although a higher value of β has the beneficial effect of inducing a lower long-run accident rate. Intuitively, an increase of β leads to lower AV sales in the long-run, and the resulting welfare loss dominates the positive welfare effect of a better safety.

Figure 7(b) depicts the difference in total discounted welfare ΔW^{tot} as a function of the stringency of liability. As long as the active steady state is induced at all (i.e. for $\beta \leq \hat{\beta} = 0.29$), the resulting welfare difference is positive, and decreasing in β. Intuitively, the initially high welfare losses induced by higher investment under a low value of β are outweighed by the intermediate and long term gains due to earlier market introduction and a larger sales volume of AVs. Our finding of an overall negative net effect of an increase of β on welfare suggests that under rule of strict (partial) liability, producers of AVs should not be held more liable than producers of conventional cars.

To check the robustness of this result, we have also considered *residual manufacturer liability*, an alternative form of product liability as analyzed by Hay and Spier (2005). Under this rule consumers are always liable up to their available funds, whereas the manufacturers cover (a fraction of) the remainder. Hay and Spier (2005) show that residual manufacturer liability can be optimal from a welfare perspective in situations where consumers cannot cover the entire damage. In our framework, residual manufacturer liability can be captured by assuming that in case of an accident, a consumer pays $k(D)$ and the innovator pays $\beta(D - k(D))$. Carrying out an analysis analogous to the one presented above, we find that setting $\beta = 0$ is socially optimal also under residual manufacturer liability.[34]

In summary, one main conclusion that can be drawn from our analysis is that, even if producers of AVs are not directly liable for accidental harm caused by these vehicles, substantial investment incentives are generated by the fact that low product safety reduces demand. In this sense our findings corroborate the argument put forward by Polinsky and Shavell (2010), that safety investments can be stipulated by market forces even in the absence of liability.

5 Negligence-based liability

In this section, we consider a model variant where the liability of the innovator is negligence-based. More precisely, we assume that a policy maker sets a safety standard (or due care level) S^{neg} such that the innovator faces no (full) liability if the safety stock of a sold AV satisfies (does not satisfy) this standard. Formally, for every accident of an AV with safety level S, the liability cost of the innovator is 0 for $S \geq S^{neg}$ and D for $S < S^{neg}$. The safety level at which it is optimal for the innovator to launch the AV, depending on the negligence standard S^{neg}, is then given by

$$\underline{S}^{neg} = \begin{cases} \underline{S}(0) & \text{if} \quad S^{neg} < \underline{S}(0), \\ S^{neg} & \text{if} \quad \underline{S}(0) \leq S^{neg} \leq \underline{S}(1), \\ \underline{S}(1) & \text{if} \quad \underline{S}(1) < S^{neg}, \end{cases} \tag{20}$$

where $\underline{S}(\beta)$ is given by Eq. (3). Intuitively, in the first case of (20) the negligence standard is so lenient that the innovator has no incentive to introduce the AV with a safety stock violating the negligence standard, even when not facing liability at all. Hence, such a standard does not influence the innovator's behavior.

[34]Also the dynamic patterns of investment, quantities and AV safety under residual manufacturer liability are qualitatively identical to those obtained for the baseline setting. The details of the analysis are available from the authors upon request.

This scenario corresponds to the case $\beta = 0$ in our analysis of Section 4.

In the second case of (20), the negligence standard induces the innovator to delay the market introduction of the AV as long as it would be fully liable, but starts to produce once the standard is met. This is optimal for the innovator as for any $S(t) < S^{neg} \leq \underline{S}(1)$ it faces full liability ($\beta = 1$), and by the definition of $\underline{S}(1)$, selling AVs with such a safety level would generate a negative expected profit. Furthermore, for $S(t) \geq S^{neg}$ the innovator faces no liability and due to $\underline{S}(0) \leq S^{neg}$ the resulting profit is non-negative. Hence, it is optimal to introduce the AV exactly when its safety level $S(t)$ reaches the negligence standard S^{neg}.

Finally, in the third case of (20) the negligence-standard is so high, that it is optimal for the innovator to launch the AV even before the negligence standard is reached, i.e. the innovator is prepared assume full liability in case of accidents.

Figure 8 illustrates the impact of S^{neg} on the dynamic patterns of the two optimal controls (i.e. investment in safety stock and output) in our calibrated model. With respect to optimal investment, consider first the black line in panel (a) which corresponds to the case of a lenient negligence standard ($S^{neg} \leq \underline{S}(0) = 0.95$). As this type of liability rule effectively does not constrain the innovator, the investment pattern is the same as in the analysis of strict (partial) liability in Section 4 for $\beta = 0$ (see Figure 3(c) above).

As the negligence standard becomes stricter (i.e. as S^{neg} increases such that we have $S^{neg} > \underline{S}(0)$), the optimal dynamic investment pattern changes and exhibits a jump.[35] Intuitively, this jump occurs at the point in time at which the negligence standard S^{neg} is reached. From that point onwards, the investment pattern closely resembles the one characterized in Section 4 for $\beta = 0$.[36] By contrast, before the threshold S^{neg} is reached, negligence-based liability gives rise to a prolonged period of increasing investment compared to the case without negligence above (see Figure 3(d) above), and this period lasts the longer, the stricter the standard. The reason is that the innovator has an incentive to reach the negligence standard S^{neg} "as soon as possible" as this leads to a downward jump in its liability costs. Hence, one advantage of introducing negligence-based liability compared to t strict (partial) liability is that it accelerates the process of safety stock accumulation, rather than slowing it down.

With respect to optimal output choices illustrated in Figure 8(b), by the same argument as above, there is no effect when the negligence standard is sufficiently weak such that it does not affect behavior. Furthermore, after the negligence standard has been reached, the output pattern is determined by the optimal quantity choice for $\beta = 0$. However, for negligence standards $S^{neg} > \underline{S}(0)$ the market introduction of the AV is

[35]As detailed at the end of Appendix B, the possibility of such jumps require slight modifications of the numerical methods used.

[36]Note that the actual investment dynamics depends on the level of the safety stock $\underline{S}^{neg}$, which varies with S^{neg}.

Figure 8: Impact of negligence-based liability on optimal dynamic AV investment and output patterns

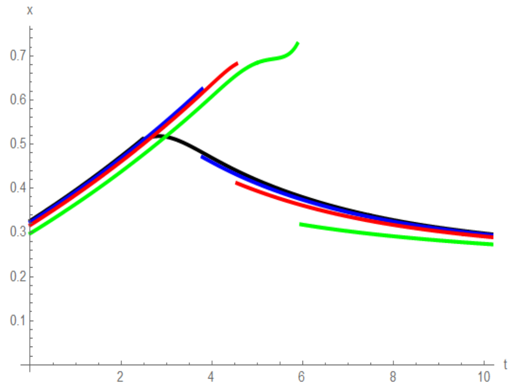

(a) Investment

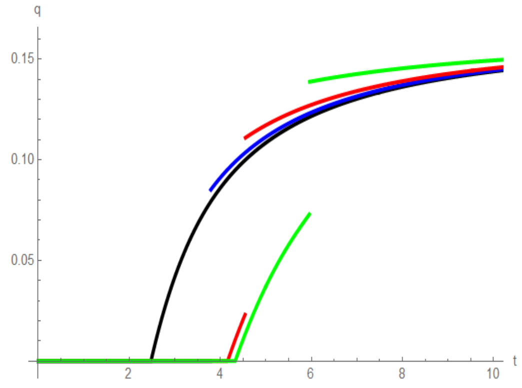

(b) Quantity

The lines of different colors show the optimal paths for different values of the negligence standards: $S^{neg} = \underline{S}(0) = 0.95$ (black), $S^{neg} = 1.57$ (blue), $S^{neg} = 1.94$ (red) and $S^{neg} = 2.61$ (green), giving rise to AV accident rates of $\alpha(0.95) = 0.12$, $\alpha(1.57) = \alpha_o = 0.074$, $\alpha(1.94) = 0.06$ and $\alpha(2.6) = 0.045$.

delayed. For $S^{neg} \in (\underline{S}(0), \underline{S}(1))$ the innovator waits with launching the AV until S^{neg} is reached, so that is exempt from liability. In this case the quantity directly jumps to a strictly positive value at the time of market introduction. Under our parametrization, $\underline{S}(1) = 1.73$ and therefore the blue line ($S^{neg} = 1.57$) in Figure 8(b) corresponds to this scenario. Alternatively, if $S^{neg} > \underline{S}(1)$ the innovator launches the AV with a continuously increasing quantity before the negligence standard is met, thereby accepting full liability. The quantity then again exhibits an upward jump as $S(t)$ reaches S^{neg}. In Figure 8(b), this case is illustrated by the red and green lines ($S^{neg} = 1.94$ and $S^{neg} = 2.61$, respectively).

Next, consider the effect of negligence-based liability on the date of market introduction of AVs and accident rates as depicted the upper part of Figure 9. Panel (a) reveals that a stricter negligence standard leads to later market introduction of AVs, which is similar to the effect of an increase of β under strict liability(compare Figure 5 (a)).

Moreover, in Figure 9(b) we observe that regardless of the negligence standard S^{neg}, the active steady state is given by $S^{ss}(0)$ and hence the long-run accident rate is $\alpha^{ss}(0)$. This is due to the fact that S^{neg} no longer has any impact once the safety stock $S(t) \geq S^{neg}$ holds.[37] Hence, tightening the negligence standard has no positive long-run effect on AV safety, which is qualitatively different from the effect of an increase in β under strict (partial) liability (see Section 4 above). However, the accident rate in the short- and intermediate term (as captured by $\alpha(S(\tilde{t}))$ at a given time $\tilde{t}$) becomes *lower* as the negligence standard becomes stricter. Again, this is qualitatively different compared to strict liability. This can be seen by

[37] Of course, this no longer holds for negligence standard $S^{neg} > S^{ss}(0)$. As in our setting $\hat{\beta} < 1$, it follows that $S^{ss}(1)$ can never be an active steady state. Therefore, for $S^{neg} > S^{ss}(0)$ the only candidate for an active steady state has a long-run safety stock equal to S^{neg}.

Figure 9: Negligence-based liability: Timing of AV market introduction, speed of safety stock accumulation and welfare

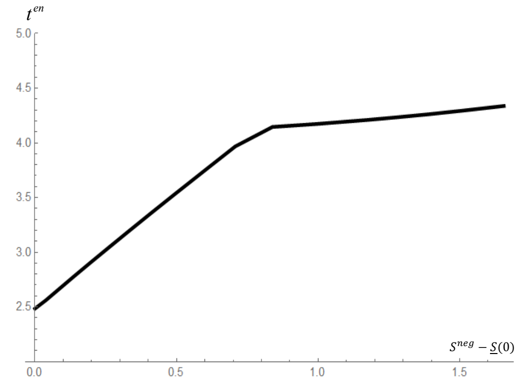

(a) Date of AV market introduction

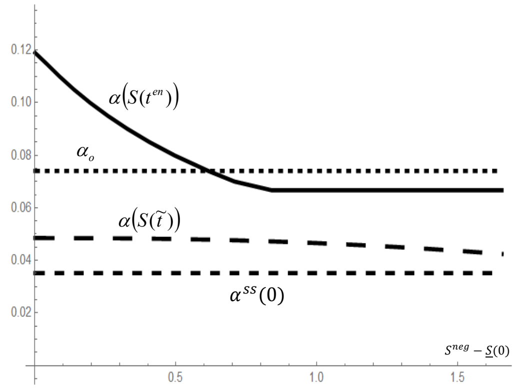

(b) AV accident rate

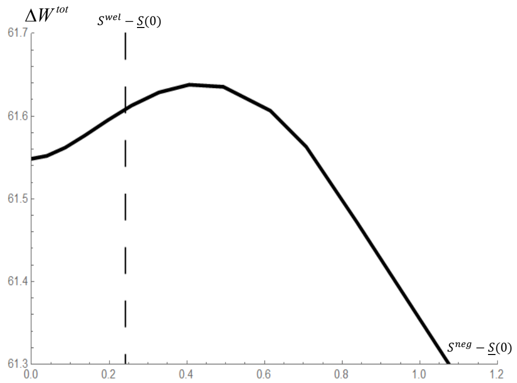

(c) Welfare

In panel (a), $t^{en}(\beta)$ indicates the point in time at which the AV is launched. Panel (b) shows the accident rate of AVs i) in the first generation ($\alpha(S(t^{en}))$), ii) at a given point in time $\tilde{t}$ ($\alpha(S(\tilde{t}))$) and iii) in the active steady state, which for all considered values of S^{neg} is given by $\alpha^{ss}(0)$. The dotted line shows the accident rate of the conventional car ($\alpha_o = 0.074$). Panel (c) indicates the total discounted welfare effect of negligence-based liability. In all panels, the horizontal axis measures the difference between the negligence standard S^{neg} and the safety stock upon market introduction in the absence of liability ($\underline{S}(0) = 0.95$).

comparing the function $\alpha(S(\tilde{t}))$ which is increasing in Figure 5(b), whereas it is decreasing in Figure 9(b). Intuitively, this is due to our finding that a tighter negligence standard boosts investment before market introduction. By contrast, as shown in Section 4, increasing β under strict liability has exactly the opposite effect and dampens investments.

From a welfare perspective tightening the negligence standard has two countervailing effects. On the positive side, investments become higher in early periods and hence get closer to the socially optimal level. This leads to a faster accumulation of the safety stock. On the negative side, a higher negligence standard induces a later market introduction of AVs and thereby, at least for large values of S^{neg}, prevents firms from selling AVs when this would be socially beneficial. In our calibrated model the positive effect prevails if S^{neg} is chosen appropriately. As can be seen in Figure 9(c), the (total discounted) welfare effect of AVs exhibits a hump-shaped pattern as the negligence standard becomes stricter, and it is maximized for $S^{neg} = 1.356$. This value is not only above the safety standard $\underline{S}(0) = 0.95$, at which the innovator would introduce the AVs without any liability, but also above $S^{wel} = 1.192$ at which AVs should be introduced from a static welfare perspective (see Section 4.5). This observation has two implications. First, a properly designed negligence-based liability rule can be welfare-enhancing compared to the optimal strict (partial) liability rule (i.e. setting $\beta = 0$). Comparing Figure 9(c) with Figure 7(b) shows that a negligence rule has quite different properties. Second, due to the intertemporal implications, the negligence standard which maximizes total discounted welfare is higher than the one for which AV should be launched from a static point of view.

Direct safety regulation An alternative instrument for inducing high safety standards at market introduction of AVs is direct safety regulation, where the regulator imposes a minimum safety standard which a new product must fulfill before it can be launched. The implications of such regulation closely resemble those of the negligence rule considered above. In particular, as shown in Appendix C, all findings obtained for the negligence rule qualitatively carry over to a scenario with direct safety regulation, apart from the obvious observation that under safety regulation the quantity has to be zero as long as the safety standard is not met. In particular, we find that tightening the safety standard increases early safety investments, and that combining direct safety regulation with the optimal strict (partial) liability ($\beta = 0$) outperforms the sole use of such a liability rule in terms of total discounted welfare.[38]

[38] Daughety and Reinganum (1995) highlight a different channel through which safety regulation might be welfare improving compared to the sole use of liability. In their signaling framework the price set by the firm serves as a signaling device for unobserved product safety. They argue that, by credibly excluding low realizations from the set of relevant safety levels, such regulation can improve overall efficiency attained in equilibrium. Our results also complement the findings from previous studies analyzing the joint use of safety regulation and liability rules (see e.g. Shavell, 1984a,b; Kolstad, Ulen, and Johnson, 1990;

6 Discussion

In this paper, we set up a dynamic model of product innovation and product safety to analyze the short- and long-term effects of product liability on the market evolution and safety features of smart products, in particular autonomous vehicles (AVs). Because of non-human tortfeasors (i.e. the product's operating system, which is based on algorithms), this issue constitutes novel challenges for liability law.

Our analysis informs legislators and policy makers about crucial trade-offs associated with product liability in the context of such innovations. In particular, under strict (partial) liability there is a long-term benefit, in the form a higher steady-state product safety, of shifting more liability to producers of smart products. However, such a shift also bears the risk of either leading to a slower accumulation of product safety, later market introduction and slower market penetration of smart products, or foregoing such innovations altogether. In our model calibration for the U.S. automobile market we find that these downsides are so significant that extending the scope of producer liabilitycompared to the status quo with conventional cars only, would lead to welfare losses. Instead, an appropriately designed negligence-based liability can actually speed up the development of safe smart products and hence be welfare-improving. Similar results are obtained for direct safety regulation.

Our model exhibits a number of simplifying features, which were helpful in fleshing out the main trade-offs. In future work our framework could also be extended in number of further directions. First, in our calibration negligence-based liability outperforms strict liability by strengthening safety investments in early stages. This suggests that complementing liability rules with policies that foster investments more directly, e.g. through subsidies, could lead to even more efficient outcomes. Such subsidies would reduce the innovator's marginal investment cost (θ), and would ceteris paribus lead to higher levels of investment in each period. However, as such subsidies will have to be financed by the consumers (e.g. through taxation), this might lower their demand for consumption goods. It would be interesting to explore the overall effects of such a subsidy scheme in our model framework.

Second, our model does not yet capture a potential externalities which seems particularly relevant in the context of smart products such as AVs. It arises from the fact that the accident rate of an AV might not only depend on its own safety stock S, but also on the number of AVs on the streets. As AVs can communicate with other AVs but not with conventional cars, a higher share of AVs improves the overall "connectedness" among vehicles, which could lead to fewer accidents.

Third, there might exist learning effects on the production side implying that the marginal cost of

Schmitz, 2000). This literature shows that combining the two instruments in an optimal way can lead to strictly higher welfare compared to relying on one instrument only.

production in a given period decrease in the total output produced prior to this period. This would give the innovator an incentive to depart from monopoly pricing in each period. From a technical point of view, it would make the analysis of the model more challenging, as this would add a second state variable to the dynamic framework.

References

Bento, A., K. Roth, and Y. Zuo (2018): "Vehicle Lifetime Trends and Scrappage Behavior in the US Used Car Market," *The Energy Journal*, 39.

Blincoe, L., T. Miller, E. Zaloshnja, and B. Lawrence (2015): "The Economic and Societal Impact of Motor Vehicle Crashes, 2010.(Revised)," *US Department of Transportation, National Highway Traffic Safety Administration, Report No. DOT HS 812 013.*

Brenkers, R. and F. Verboven (2006): "Liberalizing a Distribution System: the European Car Market," *Journal of the European Economic Association*, 4, 216–251.

Colonna, K. (2012): "Autonomous Cars and Tort Liability," *Journal of Law, Technology & the Internet*, 4, 81–130.

Daughety, A. F. and J. F. Reinganum (1995): "Product safety: Liability, R&D, and Signaling," *American Economic Review*, 1187–1206.

Dawid, H., M. Keoula, M. Kopel, and P. Kort (2015): "Product Innovation Incentives by an Incumbent Firm: A Dynamic Analysis," *Journal of Economic Behavior and Organization*, 117, 411–438.

Dawid, H., M. Keoula, and P. Kort (2017): "Numerical Analysis of Markov-Perfect Equilibria With Multiple Stable Steady States: A Duopoly Application With Innovative Firms," *Dynamic Games and Applications*, 7, 555–577.

Doraszelski, U. (2003): "An R&D Race With Knowledge Accumulation," *RAND Journal of Economics*, Spring, 20–42.

Douma, F. and S. Palodichuk (2012): "Criminal Liability Issues Created by Autonomous Vehicles," *Santa Clara Law Review*, 52, 1157.

European Parliament (2017): "Civil Law Rules on Robotics," *2015/2103(INL) - 16/02/2017.*

Federal Ministry of Transport and Digital Infrastructure (2017): "Ethik-Kommission Automatisiertes und Vernetztes Fahren: Bericht Juni 2017," *https://www.bmvi.de/goto?id=348344*.

Friedman, E. and E. Talley (2019): "Automatorts: How Should Accident Law Adapt to Autonomous Vehicles? Lessons From Law & Economics," *Columbia Univeristy, mimeo.*

Galasso, A. and H. Luo (2017): "Tort Reform and Innovation," *The Journal of Law and Economics*, 60, 385–412.

——— (2018a): "Punishing Robots: Issues in the Economics of Tort Liability and Innovation in Artificial Intelligence," in *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press.

——— (2018b): "When Does Product Liability Risk Chill Innovation? Evidence from Medical Implants," *NBER Working Paper No. 19-002*.

Gilles, S. (2006): "The Judgment-Proof Society," *Washington & Lee Law Review*, 63, 603–715.

Gless, S., E. Silverman, and T. Weigend (2016): "If Robots Cause Harm, Who is to Blame? Self-Driving Cars and Criminal Liability," *New Criminal Law Review*, 19, 412–436.

Goldberg, P. and F. Verboven (2001): "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, 68, 811–848.

Grass, D., J. Caulkins, G. Feichtinger, G. Tragler, and D. Behrens (2008): *Optimal Control of Nonlinear Processes*, Springer, Berlin.

Hartl, R. (1987): "A Simple Proof of the Monotonicity of the State Trajectories in Autonomous Control Problems," *Journal of Economic Theory*, 41, 211–215.

Hay, B. and K. Spier (2005): "Manufacturer Liability for Harms Caused by Consumers to Others," *American Economic Review*, 95, 1700–1711.

Hinloopen, J., G. Smrkolj, and F. Wagener (2013): "From Mind to Market: A Global, Dynamic Analysis of R&D," *Journal of Economic Dynamics and Control*, 37, 2729–2754.

Judd, K. (1998): *Numerical Methods in Economics*, MIT Press, Cambridge.

Kagawa, S., T. Tasaki, and Y. Moriguchi (2006): "The Environmental and Economic Consequences of Product Lifetime Extension: Empirical Analysis for Automobile Use," *Ecological Economics*, 58, 108–118.

Kolstad, C., T. Ulen, and G. Johnson (1990): "Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements?" *American Economic Review*, 888–901.

Lambertini, L. and A. Mantovani (2009): "Process and Product Innovation by a Multiproduct Monopolist: A Dynamic Approach," *International Journal of Industrial Organization*, 27, 508–518.

Marchant, G. and R. Lindor (2012): "The Coming Collision Between Autonomous Vehicles and the Liability System," *Santa Clara Law Review*, 52, 1321–1340.

McGuire, E. (1988): "The Impact of Product Liability," *The Conference Board*, Report 908.

Michel, P. (1982): "On the Transversality Condition in Infinite Horizon Optimal Problems," *Econometrica*, 975–985.

Miranda, M. and P. Fackler (2002): *Applied Computational Economics and Finance*, MIT Press, Cambridge.

OECD (2017): "Key Issues for Digital Transformation in the G20," *www.oecd.org/G20/key-issues-for-digital-transformation-in-the-G20.pdf.*

Parchomovsky, G. and A. Stein (2008): "Torts and Innovation," *Michigan Law Review*, 107, 286–315.

Polinsky, M. and S. Shavell (2010): "The Uneasy Case for Product Liability," *Harvard Law Review*, 123, 1437–1492.

Porter, M. and J. Heppelmann (2015): "How Smart, Connected Products are Transforming Companies," *Harvard Business Review*, 93, 96–114.

Rogers, E. M. (2003): *The Diffusion of Innovation. (5th ed.)*, Free Press, New York.

Sassone, A., M. Grosso, M. Poncino, and E. Macii (2016): "Smart electronic Systems: An Overview," in *Smart Systems Integration and Simulation*, ed. by N. Bombieri, M. Poncino, and G. Pravadelli, Springer, 5–21.

Schellekens, M. (2015): "Self-Driving Cars and the Chilling Effect of Liability Law," *Computer Law & Security Review*, 31, 506–517.

Schmitz, P. W. (2000): "On the Joint Use of Liability and Safety Regulation," *International Review of Law and Economics*, 20, 371–382.

Shavell, S. (1984a): "Liability for Harm Versus Regulation of Safety," *The Journal of Legal Studies*, 13, 357–374.

——— (1984b): "A Model of the Optimal Use of Liability and Safety Regulation," *The Rand Journal of Economics*, 15, 271–280.

——— (2019): "On the Redesign of Accident Liability for the World of Autonomous Vehicles," *Harvard University, mimeo.*

Skiba, A. (1978): "Optimal Growth With a Convex-Concave Production Function," *Econometrica*, 527–539.

Smith, B. (2017): "Automated Driving and Product Liability," *Michigan State Law Review*, 1, 1–74.

UN Conference on Trade and Development (2017): *Information Economy Report 2017: Digitalization, Trade and Development*, United Nations Publication, UNCTAD/IER/2017.

Viscusi, K. and M. Moore (1993): "Product Liability, Research and Development, and Innovation," *Journal of Political Economy*, 101, 161–184.

Viscusi, W. K. (1991): *Reforming Products Liability*, Harvard University Press.

——— (2012): "Does Product Liability Make Us Safer," *Regulation*, 35, 24.

Viscusi, W. K. and M. Moore (1991): "An Industrial Profile of the Links Between Product Liability and Innovation," in *The Liability Maze: The Impact of Liability Law on Safety and Innovation*, ed. by P. Huber and R. Litan, Brookings Institution Press, 81–119.

Wagner, G. (2018): "Robot Liability," *SSRN Working Paper No. 3198764*.

Appendix (For Online Publication)

A Proofs

A.1 Proof of Lemma 1

We define the current-value Hamiltonian function

$$H(x, S, \lambda) = [\pi^*(S) - h(x)] + \lambda(x - \delta S), \tag{21}$$

where λ is the costate for the safety stock S. Applying Pontryagin's maximum principle (see e.g. Grass et al., 2008), we derive the optimal control from the first order condition $\frac{dH}{dx} = 0$. Taking into account the non-negativity constraint of the control, this yields (10). The costate equation follows from

$$\begin{aligned} \dot{\lambda} &= r\lambda - \frac{\partial H}{\partial S} \\ &= r\lambda - \frac{\partial \pi}{\partial q^*} \cdot \frac{\partial q^*}{\partial S} - \frac{\partial \pi}{\partial S} + \delta\lambda \\ &= (r + \delta)\lambda - \frac{\partial \pi}{\partial S} \end{aligned}$$

where $\pi = \Big[q^*(S) \cdot \Big(p(q^*(S), S) - c - \frac{\alpha(S)\beta D}{r+\rho}\Big) - h(x)\Big]$ is the integrand of (9). With respect to the term $\frac{\partial \pi}{\partial S}$, it follows from the envelope theorem that $\frac{\partial \pi}{\partial q^*} = 0$, so that only the direct effect of S on π needs to be considered. That is, we have $\frac{\partial \pi}{\partial S} = q^*(S) \cdot \Big(\frac{\partial p(q,S)}{\partial S} - \frac{\alpha'(S) \cdot \beta D}{r+\rho}\Big)$, where $p(q, S)$ is the demand function derived from (4) and given by

$$p(q, S) = c + \frac{A - (\bar{v} - \underline{v}) \cdot q - \alpha(S) \cdot k((1 - \beta)D)}{r + \rho}.$$

Taking this into account when calculating the derivate of $\pi(s)$ with respect to S yields the expression in the lemma. The transversality condition follows from Michel (1982). ■

A.2 Proof of Lemma 2

As this is a time-autonomous infinite horizon optimization problem with a one-dimensional state space, standard results (see e.g. Hartl, 1987) imply that every optimal trajectory has to be (weakly) monotone. To see that no optimal path can induce divergence of S to infinity, note that $\pi^*(S)$ is bounded from above. For any monotonously increasing path we must have $x > \delta S$. Furthermore there exists a lower bound $\tilde{S}(\beta, \theta)$ such that $h(\delta S) > \pi^*(S)$ for all $S > \tilde{S}$. As $h(\delta S)$ is shifted upwards for increasing θ, whereas $\pi^*(S)$ is independent of θ, if follows that $\frac{\partial \tilde{S}}{\partial \theta} \leq 0$ for all β. Defining $\bar{S}(\theta) = \max_{\beta \in [0,1]} \tilde{S}(\beta, \theta)$, we obtain that $\bar{S}' \leq 0$. Therefore for any path diverging to infinity there is a $T > 0$ such that $S(t) > \bar{S}(\theta)$ for $t > T$. Hence,

$\pi^*(S(t)) - h(x(t)) < 0$ for all $t > T$, which implies that such a path can never be optimal. Together these insights imply that every optimal path monotonously converges to some steady state in $[0, \bar{S}(\theta)]$. ■

A.3 Proof of Proposition 1

(i): In order to prove this claim, we show that for any β it holds that $\lim_{\theta\to\infty} V(0;\beta,\theta) = 0$. Considering the canonical system derived from the Maximum Principle, any steady state $S^{ss} > 0$ has to satisfy $g_\lambda(S^{ss}) = g_S(S^{ss})$, where g_λ and g_S are given by (12) and (13). The function g_λ is finite, continuous in S and does not depend on θ. Fix some small but positive θ_{min}. The function $g_S(S)$ increases monotonously in θ for all $S \in [0, \bar{S}(\theta^{min})]$ with $\lim_{\theta\to\infty} g_S(S) = \infty$. Therefore for any given $\beta \in [0,1]$ there exists a finite $\theta^{max}(\beta) = \inf\left\{\theta \geq \theta^{min} \mid g_S(S) > g_\lambda(S)\ \forall S \in [0, \bar{S}(\theta_{min})]\right\}$. By definition for each $\theta > \theta^{max}(\beta) \geq \theta^{min}$ we have $g_S(S) > g_\lambda(S)$ for all $S \in [0, \bar{S}(\theta)]$. Hence, there is no candidate for a steady state in $(0, \bar{S}(\theta)]$. As shown in Lemma 2, any optimal path entering $(0, \bar{S}(\theta)]$ must monotonously converge to a steady state in this interval. This implies that the optimal path starting from $S(0) = 0$ cannot leave this state and therefore $V(0;\beta,\theta) = 0$ holds for sufficiently large θ.

Turning to the main claim of part (i) of the proposition, consider some $\beta \in [0,1]$ and denote by $\epsilon < \delta \underline{S}(\beta)$ a small positive number. Then consider the intertemporal optimization problem of the innovator with the additional constraint $x(t) \geq \epsilon\ \forall t \geq 0$. We denote this as the $\epsilon-$constrained problem of the innovator. Standard arguments show that the problem has an optimal solution and we denote by $V_\epsilon(0;\beta,\theta)$ the value function of this problem for the initial state $S(0) = 0$. A decrease in θ shifts the function $h(x)$ downwards, and as investment is positive along the optimal path in the $\epsilon-$constrained problem, it follows that $V_\epsilon(0;\beta,\theta)$ is a strictly decreasing function of θ. Furthermore, we show that $V_\epsilon(0;\beta,\theta) > 0$ for sufficiently small θ. To see this, note that $\lim_{S\to\infty} \pi^*(S) > 0$ and hence $\pi^*(\hat{S}) > 0$ for sufficiently large $\hat{S}$. Now consider a path $\{\hat{x}(t)\}$ with the properties that the induced state dynamics $S(t)$ reaches $\hat{S}$ for $t = T$ and $\hat{x}(t) = \delta\hat{S}\ \forall t \geq T$. For the discounted net present value of this path, $\hat{V}$, we obtain

$$\begin{aligned}
\hat{V} &\geq -\int_0^T e^{-rt} h(\hat{x}(t))dt + \frac{e^{-rT}}{r}\left(\pi^*(\hat{S}) - h(\delta\hat{S})\right) \\
&= -\theta\left[\int_0^T e^{-rt}\left(\hat{x}(t) + \frac{\eta}{2}\hat{x}(t)^2\right)dt + \frac{e^{-rT}}{r}\left(\delta\hat{S} + \frac{\eta}{2}(\delta\hat{S})^2\right)\right] + \frac{e^{-rT}}{r}\pi^*(\hat{S})
\end{aligned}$$

The square bracket does not depend on θ and therefore the negative term goes to zero for $\theta \to 0$. Moreover, $\pi^*(\hat{S}) > 0$, which implies that $\hat{V} > 0$ for sufficiently small θ. As $V_\epsilon(0;\beta,\theta) \geq \hat{V}$, we obtain $\lim_{\theta\to 0} V_\epsilon(0;\beta,\theta) > 0$. Considering now large values of θ, as shown above $\lim_{\theta\to\infty} V(0;\beta,\theta) = 0$, which implies $\lim_{\theta\to\infty} V_\epsilon(0;\beta,\theta) \leq \lim_{\theta\to\infty} V(0;\beta,\theta) = 0$. Taking into account the monotonicity and using the intermediate value theorem we obtain that there exists a unique $\hat{\theta}_\epsilon(\beta)$ such that $V_\epsilon(0;\beta,\hat{\theta}_\epsilon(\beta)) = 0$. Clearly any optimal path giving rise to this value must converge monotonously to some steady state and as $\pi^*(S) = 0$ for all $S \leq \underline{S}(\beta)$ and $h(x(t)) > 0$ along any path in the $\epsilon-$constrained problem, this steady-state, which we denote by $S_\epsilon^{ss}(\beta)$, must satisfy $S_\epsilon^{ss}(\beta) > \underline{S}(\beta)$ for all $\epsilon > 0$. Taking the limit $\epsilon \to 0$ it follows that for $\theta = \hat{\theta}(\beta) := \lim_{\epsilon\to 0} \hat{\theta}_\epsilon(\beta)$ the value of the unconstrained problem of the innovator is zero and that this value can be obtained both by the path $x(t) = 0$ $\forall t$ and by a path with positive investment, which induces convergence of the state S to the steady state $S^{ss}(\beta) = \lim_{\epsilon\to\infty} S_\epsilon^{ss}(\beta) > \underline{S}(\beta)$.[39]

Furthermore, we show that $\hat{\theta}(\beta)$ is strictly decreasing with respect to β. Consider an arbitrary $\beta \in (0,1)$. The value function of the problem for $\theta = \hat{\theta}(\beta)$ by definition is $V(0;\beta,\hat{\theta}(\beta)) = 0$. As shown above, this value is generated by an optimal path for which we have $q(t) > 0$, $\forall t > T$ for some large T. Using the envelope theorem, we obtain from (9)

$$\frac{\partial \pi^*(S(t))}{\partial \beta} = -q^*(S(t))\frac{\alpha(S(t))D}{r+\rho} < 0$$

for all t where $q(S(t)) > 0$ and zero otherwise. Hence, $V(0;\tilde{\beta},\hat{\theta}(\beta)) > 0$ for all $\tilde{\beta} < \beta$. Similar arguments as used above show that $V(0;\beta,\theta)$ is a strictly decreasing function of θ if the optimal path induces $q(S(t)) > 0$ in a time interval with positive measure. It follows that $\hat{\theta}(\tilde{\beta}) > \hat{\theta}(\beta)$, which establishes the strict monotonicity of $\hat{\theta}(\beta)$ with respect to β. Defining $\underline{\theta} = \min_{\beta\in[0,1]} \hat{\theta}(\beta)$ we immediately obtain the claim of (i) with $\underline{\theta} = \hat{\theta}(1)$.

(ii): Defining $\bar{\theta} = \hat{\theta}(0)$, it follows from the arguments given in part (i) that for $\theta = \bar{\theta}$ and $\beta = 0$ the innovator has two optimal paths giving the value $V(0;0,\bar{\theta}) = 0$, one of them converging to a positive steady state and one staying at $S = 0$. For any positive β only the path $S(t) = 0$, $\forall t$ is optimal. This implies directly that for any $\theta > \bar{\theta}$ any path from $S(0) = 0$ which converges to a positive steady state generates a negative value regardless of $\beta \in [0,1]$.

(iii): Follows directly from the strict monotonicity of $\hat{\theta}(\beta)$ with respect to β and the definitions $\underline{\theta} = \hat{\theta}(1)$ and $\bar{\theta} = \hat{\theta}(0)$. ■

[39] It should be noted that by the same arguments as used above, the value of any path monotonously converging from $S(0) = 0$ to a steady state at $\underline{S}(\beta)$ has to be negative. Therefore $S^{ss}(\beta) > \underline{S}(\beta)$ must hold.

B Details of the Numerical Method

In this section, we briefly outline the numerical method used in the analysis.

Main analysis To determine the innovator's optimal investment path from the initial condition $S(0) = 0$ we first check whether for a given parameter setting there is a candidate for an active steady state. If this is the case, we then rely on the HJB equation associated to the innovator's dynamic optimization problem in order to determine the optimal path leading to this candidate and the associated value of the innovator's objective function (9).[40]

To identify candidates for an active steady state we numerically solve the equation $g_\lambda(S) = g_S(S)$ with g_λ and g_S given by (12) and (13). As discussed in Section 3, this equation (apart from the non-generic case of tangency between the two isoclines) has at least two solutions if it has any. The smallest solution is always repelling in the state/co-state space and therefore we always consider the second smallest solution of this equation as the location of the candidate for an active steady state.[41] We denote this candidate by $\tilde{S}^{ss}$.

The value of an optimal path to $\tilde{S}^{ss}$ is determined using the HJB equation. As the problem is time-autonomous and has an infinite horizon, we can consider stationary investment and value functions, which do not explicitly depend on time. Denoting by $\tilde{V}(S)$ the value function associated to the optimal path leading to $\tilde{S}^{ss}$, the HJB equation is given by

$$r\tilde{V}(S) = \begin{cases} \max_{x\geq 0}\left[-h(x) + \frac{\partial \tilde{V}(S)}{\partial S}(x - \delta S)\right] & S \in [0, \underline{S}(\beta)), \\ \max_{x\geq 0}\left[q^*(S)\cdot\left(p^*(S) - c - \frac{\alpha(S)\beta D}{r+\rho}\right) - h(x) + \frac{\partial \tilde{V}(S)}{\partial S}(x - \delta S)\right] & S \geq \underline{S}(\beta). \end{cases} \tag{22}$$

Maximization of the right hand side of the HJB equation yields (see also (10))

$$x^*(S) = \max\left[\frac{1}{\theta\eta}\frac{\partial \tilde{V}(S)}{\partial S} - \frac{1}{\eta}, 0\right], \tag{23}$$

and inserting this expression into (22) yields a non-linear first order differential equation in $\tilde{V}$. The fact the we consider paths (at least locally) leading to $\tilde{S}^{ss}$ is incorporated by the condition

$$\tilde{V}(\tilde{S}^{ss}) = \frac{1}{r}\left(q^*(\tilde{S}^{ss})\cdot\left(p^*(\tilde{S}^{ss}) - c - \frac{\alpha(\tilde{S}^{ss})\beta D}{r+\rho}\right) - h\left(\delta\tilde{S}^{ss}\right)\right). \tag{24}$$

We numerically solve the HJB equation on the state space $[0, \bar{S}]$ with $\bar{S} > \tilde{S}^{ss}$ relying on a collocation method using Chebychev polynomials. To this end we generate a set of n Chebychev nodes $\mathcal{N}$ in $[0, \bar{S}]$ (see

[40] See also Miranda and Fackler (2002) or Dawid et al. (2015) for more details of the application of collocation for the solution of HJB equations.

[41] In all scenarios covered in this article the equation $g_\lambda(S) = g_S(S)$ has at most two solutions.

e.g. Judd (1998) for the definition of Chebychev nodes and Chebychev polynomials).

Our goal is to calculate a polynomial approximation of $\tilde{V}(S)$ which (approximately) satisfies (22) on the set of interpolation nodes $\mathcal{N}$. The set of basis functions for the polynomial approximation is determined as $\mathcal{B} = \{B_j(S), j = 1, .., n\}$ with

$$B_j(S) = T_{j-1}\left(-1 + \frac{2S}{\bar{S}}\right),$$

where $T_j(x)$ denotes the j-the Chebychev polynomial. As Chebychev polynomials are defined on $[-1, 1]$, the state variables have to be transformed accordingly.

For a given value of β the function $\tilde{V}$ is then approximated by

$$\tilde{V}(S) \approx \hat{V}(S;\beta) = \sum_{j=1}^{n} C_j(\beta)B_j(S), \;\; S \in [0, \bar{S}], \tag{25}$$

where $C(\beta) = \{C_j(\beta)\}$ with $j = 1, \cdots, n$ is the set of n coefficients to be determined. To calculate these coefficients we set up a system of non-linear equations derived from the condition that $\hat{V}$ satisfies the HJB equation (22) on the set of interpolation nodes $\mathcal{N}$. This system consists of n equations with n unknowns (i.e. the coefficients $C_j(\beta)$) and is solved by a standard numerical procedure for solving systems of non-linear equations.

Once a solution vector $C(\beta) \in \mathcal{R}^n$ is obtained the accuracy of the solution is checked in two ways. First, the absolute difference between the left hand side and right hand side of (22) for $\tilde{V}(S) = \hat{V}(S;\beta)$ relative to the value of $\hat{V}(S)$ is calculated on the state space $[0, \bar{S}]$. In all numerical results presented in the article this difference value is below $2 \cdot 10^{-3}$, which shows that $\hat{V}(S;\beta)$ almost exactly solves the HJB equation, not only on $\mathcal{N}$ but on the entire state space. Second, the value $\hat{V}(0;\beta)$ has been verified by calculating a discrete time approximation (for small time steps) of the optimal investment path and state dynamics derived from $\hat{V}$ and explicitly calculating the corresponding value of (9). For our numerical results also this test confirms that the obtained value functions and optimal investments obtained from inserting $\hat{V}(S;\beta)$ into (23) are very close approximations of the actual optimal values.

The condition (24) is enforced by initially determining the coefficient vector $C(\beta)$ for $\beta = 0$ and verifying that with this coefficient vector $\tilde{S}^{ss}$ is indeed the fixed point of the state dynamics induced by the feedback function (23) if $\tilde{V} = \hat{V}$, which implies by the HJB equation that (24) holds. Furthermore for $\beta = 0$ we have $\hat{V}(0;0) > 0$ and therefore the path leading to $\tilde{S}^{ss}$ is optimal from the initial condition $S = 0$. The parameter β is then increased stepwise with a small step-size ϵ and the coefficient vector $C(\beta - \epsilon)$ is used as the initial guess in the iterative algorithm solving (22) on $\mathcal{N}$ for β. This continuous adjustment of the value function implies that there is always a positive steady state of the dynamics induced by $\hat{V}$ and the function $\hat{V}$ captures the value of the optimal paths converging to this positive steady state. As β increases the value of $\hat{V}(0;\beta)$ declines and the value of the threshold $\hat{\beta}$ is determined by the condition that it is the smallest value of β such that $\hat{V}(0;\beta+\epsilon) < 0$.

Adjustments for case of negligence-based liability In order to determine the value function and the optimal feedback function for the case of negligence-based liability discussed in Section 5, we have to slightly adjust the numerical algorithm. The reason is that under this liability rule the value of β jumps from $\beta = 1$ to $\beta = 0$ once the safety stock S crosses the negligence standard S^{neg}. As such a jump in the instantaneous objective function of the innovator implies a kink in the value function and such a kink cannot be captured by a polynomial approximation, we have to solve the HJB equation separately on the intervals $[0, S^{neg}]$ and $[S^{neg}, \bar{S}]$. As $\beta = 0$ holds on $[S^{neg}, \bar{S}]$ and we only consider scenarios where the negligence standard is below the steady state without liability, i.e. $S^{neg} < S^{ss}(0)$, the value function on the interval $[S^{neg}, \bar{S}]$ does not depend on S^{neg}. Actually, $\hat{V}$ on this interval coincides with the value function we have calculated in Section 4 for $\beta = 0$, i.e. $\hat{V}(S;0)$. To determine the value function $\tilde{V}$ on $[0, S^{neg}]$ for a given S^{neg} we use the collocation method for this reduced state space to solve the HJB equation (22) for $\beta = 1$ and the boundary condition $\tilde{V}(S^{neg}) = \hat{V}(S^{neg};0)$.

C Safety regulation

In this section we consider the case of direct safety regulation. In particular, the regulator imposes a minimum safety stock S^{min} which a new product must fulfill before it can be launched. A high (low) value of S^{min} indicates a strict (weak) regulation. The numerical procedure for solving the model under safety regulation is analogous to that under negligence-based liability in the sense that the HJB equation is solved separately on the intervals $S \in [0, S^{min})]$ and $[S^{min}, \bar{S}]$. Figure 10 shows the dynamics of investment and quantity under different levels of minimum safety standard. Comparing with Figure 8 obtained under negligence-based liability shows that the dynamic patterns are indeed very similar apart from the obvious observation that under safety regulation the quantity has to be zero as long as $S < S^{min}$. Also with respect to the date of AV market introduction, the speed of safety stock accumulation and total discounted welfare an increase in S^{min} has qualitatively very similar effects as increasing the negligence standard S^{neg}. This can be seen by comparing Figure 11 with Figure 9. In particular, the welfare maximizing minimum safety standard is above both $\underline{S}(0)$ and S^{wel}. Hence, safety regulation can increase welfare relative to the benchmark of no regulation and $\beta = 0$.

Figure 10: Impact of direct AV safety regulation on optimal dynamic investment and output patterns

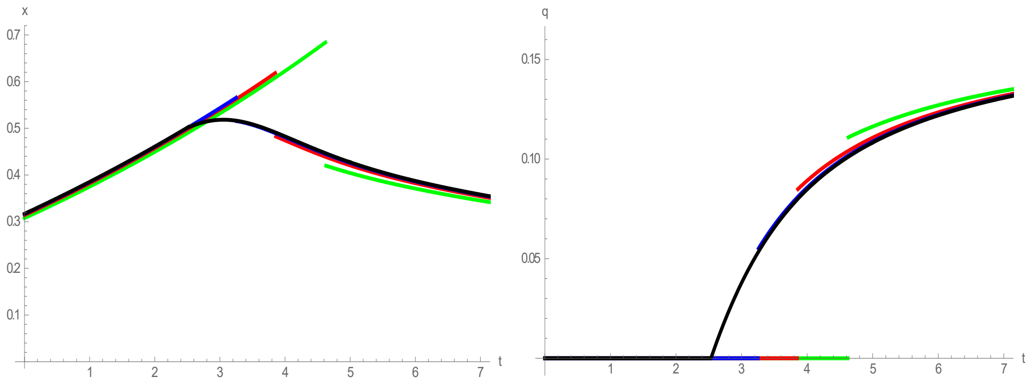

(a) Investment (b) Quantity

The lines of different colors show the optimal paths for different values of the maximum accident rate α^{max} for the AV: $\alpha^{max} = 0.12$ (black), $\alpha^{max} = 0.09$ (blue), $\alpha^{max} = \alpha_0 = 0.074$ (red), $\alpha^{max} = 0.06$ (green).

Figure 11: Impact of direct AV safety regulation on timing of market introduction, speed of safety stock accumulation and welfare

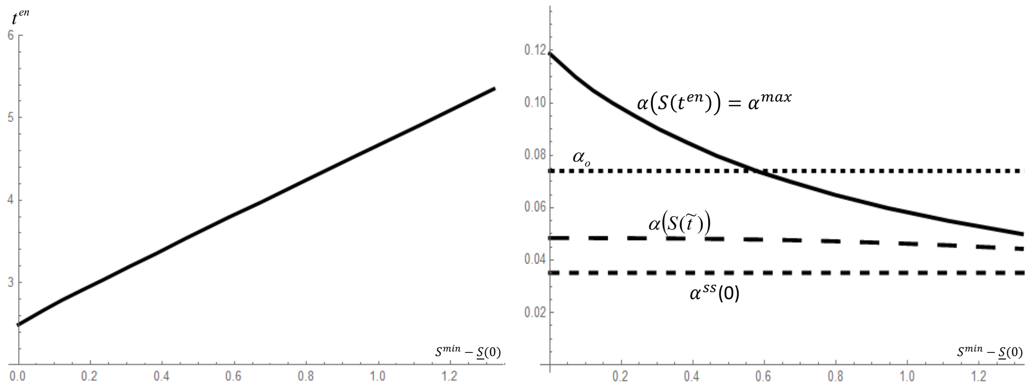

(a) Date of AV market introduction (b) AV accident rate

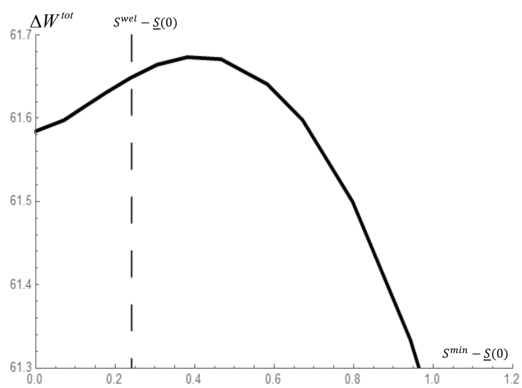

(c) Welfare

Panel (a): $t^{en}(\beta)$ indicates the point in time at which the AV is launched (i.e. where the required safety stock S^{min} is reached. Panel (b): accident rate of AV i) in the first generation ($\alpha(S(t^{en})) = \alpha(S^{min})$), ii) at a given point in time $\tilde{t}$ ($\alpha(S(\tilde{t}))$) and iii) in the active steady state (α^{ss}). The dotted line shows the accident rate of the conventional car ($\alpha_o = 0.074$). Panel (c): the total discounted welfare effect of safety regulation. All panels are drawn for $\beta = 0$ and the horizontal axis measures the difference between the minimum safety stock S^{min} and the safety stock upon market introduction in the absence of liability ($\underline{S}(0) = 0.95$).