Platform liability, product safety and optimal pricing[*]

Herbert Dawid[†] Mohammad Ali Kashefi[‡] Gerd Muehlheusser[§]

September 8, 2025

Preliminary and incomplete: PLEASE DO NOT CITE

Abstract

Should platforms be held legally responsible for harm caused by products or services offered by third-party sellers? We address this question in a novel setting where product safety (the likelihood of accidents) is determined by a third-party seller, while the consumer price is set by a monopolistic platform that pays the seller a commission. Ride-hailing platforms serve as a relevant example. When an accident occurs, the resulting harm is divided between the platform and the seller according to a liability rule (strict partial liability) set ex ante by a welfare-maximizing policymaker. In addition to the liability regime, the seller's incentive to provide high product safety is influenced by overall demand, which is determined by the platform's pricing. The platform's optimal pricing balances maximizing consumer revenue and minimizing expected liability costs, with the latter depending on the seller's chosen safety level. Consequently, the platform may set a price below the monopoly level to boost demand, thereby encouraging the seller to improve product safety. We show that the optimal liability rule is dichotomous: it either exempts the platform from liability or assigns it the maximum share possible while still incentivizing the seller to provide high product safety. In an extension we consider platform externalities, where consumers' valuation for the platform good depends on the number of active sellers, thereby also allowing for an endogenous number of active sellers. Our results are qualitatively very similar to the baseline model.

Keywords: Platform liability, consumer pricing, product safety

JEL classification: K13, L13, L15, L86, D42,

[*]This paper has benefited from comments and suggestions from Xinyu Hua and from seminar audiences at the University of Hong Kong, National Dong Hwa University (Taiwan), and the University of Hamburg. We gratefully acknowledge the hospitality and financial support by the Center of Interdisciplinary Research (ZiF) at Bielefeld University (Research Group "Economic and Legal Challenges in the Advent of Smart Products").

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

[‡]Bielefeld University, Department of Business Administration and Economics, mohammadali.kashefi@uni-bielefeld.de

[§]University of Hamburg, Department of Economics, and IZA, and CESifo, gerd.muehlheusser@uni-hamburg.de

1 Introduction

Motivation In recent years, digital platforms have become central to everyday life, playing key roles in social interaction, media consumption, and shopping for goods and services. In the United States, individuals spend an average of 143 minutes per day on social media. The video platform YouTube, for example, accounts for a global viewing time of 1 billion hours per day (Global Media Insight, 2025). With respect to shopping, e-commerce accounted for over 15.6% of total US retail sales in 2023. Currently, eight out of ten internet users in the EU shop online, and global retail e-commerce sales in 2024 are estimated at $4.1 trillion (Statista, 2025a,b; DeStatis, 2025). Moreover, by the end of 2024, 62% of all units sold on the shopping platform Amazon were provided by third-party sellers (Marketplace Pulse, 2025).

Platforms manage the interactions between their users. A key issue in this context is the question of legal responsibility for harm arising from transactions conducted on the platform, particularly when the content, good, or service causing the harm was *not* provided by the platform itself. Examples include users uploading material on hosting platforms that it is either protected by intellectual property rights (such as movies or songs, see e.g. NME, 2023) or outright illegal (such as terrorist propaganda, see e.g. BBC, 2020; GovInfo, 2017). Moreover, third-party sellers often provide defective products on shopping platforms that cause injuries to buyers (e.g. an exploding battery, see New York Post, 2024), or engage in offensive behavior such as harassment or even rape in the course of personal interaction with customers (see e.g. New York Times, 2025; CNN, 2022; USA Today, 2022).

To what extent should a platform be held liable for such harm? One could argue that it would be inappropriate to hold a platform liable when it was not directly involved in the harmful activity in question. Indeed, traditional legal practice has reflected this lenient approach. For example, platforms were initially shielded from liability under Section 230 of the U.S. Communications Decency Act, while the EU e-Commerce Directive (2000/31/EC) established a safe harbor for hosting services, provided they lacked knowledge of unlawful activities (see e.g., Buiten, de Streel, and Peitz, 2020).[1] Even under the more recent EU Digital Services Act, liability exemptions persist, albeit with more stringent obligations for very large platforms (see e.g. Lefouili and Madio, 2022). Also for the case of shopping platforms, courts have repeatedly ruled that platforms are not "sellers" in the traditional sense and hence are not liable for harm caused by defective goods traded on the platform.[2]

However, this traditional approach has been increasingly challenged. For example, in a landmark case involving shopping platforms, Amazon was found liable for harm caused by a defective (i.e., exploding) battery purchased from a third-party seller. In particular, the California Court of Appeal found Amazon liable as "an integral part of the overall producing and marketing enterprise," thereby reversing the lower court's ruling in favor of Amazon.[3] More recently, in 2024 the U.S. Consumer Product Safety Commission (CPSC, 2024) has established that, in its role of 'distributor', Amazon bears legal responsibility for the recall of defective goods by third-party sellers under Federal Safety Law.

The ongoing shift in legal practice regarding platform liability is accompanied by a growing academic literature in law and economics (discussed in detail below). This literature demonstrates that, even when platforms cannot influence product safety directly, they can nonetheless exert substantial control over the

[1] Historically, regulatory leniency toward platforms was motivated by the initial goal of fostering competition in digital markets, allowing platforms to emerge as viable challengers to incumbent sellers. Policymakers largely believed that holding platforms liable could inhibit innovation, market expansion, and investment in digital infrastructure (see, e.g., Lefouili and Madio, 2022; Regibeau, 2025; Spier and Van Loo, 2025). However, as many platforms have since become strong or even dominant market players, facilitating their growth and prosperity is no longer a compelling regulatory priority.

[2] See e.g. Erie Ins. Co. v. Amazon.com 925 F.3d 135 (4th Cir. 2019).

[3] See Bolger v. Amazon.com, 53 Cal.App. 5th 431 (Cal. Ct. App.2020) and the discussion in Busch (2021).

quality and safety of products and services traded in digital markets—for example, through (costly) monitoring activities that enable the identification and removal of harmful products or sellers from the platform (Hua and Spier, 2025; Yasui, 2022). Moreover, holding platforms liable will also affect the incentives of other parties (e.g. sellers) to invest in product safety (Hua and Spier, 2025; De Chiara et al., 2024; Zennyo, 2023; Yasui, 2022). In their overview, Spier and Van Loo (2025) challenge the view that holdings platforms liable inevitably chills innovation and that platforms can be trusted to self-police harmful conduct. They argue that, in practice, market incentives often fail to prevent harm, making a case for carefully designed liability regimes.

To date, the academic literature on platform liability has primarily studied platform types where the consumer price either plays a minor role only (e.g., social platforms such as Facebook and hosting platforms such as YouTube), or where it is set by the third-party sellers active on the platform (e.g., shopping platforms such as Amazon).[4] However, there are also platform types on which potentially harmful goods or services are offered solely by third-party sellers, while the consumer price is set by the platform. A case in point are ride-hailing platforms such as Uber and Lyft who by now manage millions of trips per day.[5] Both platforms set the consumer fares, but the rides are provided by independent drivers rather than the platform itself. Consumers pay the platform directly, and drivers receive a commission from the platform.[6] Drivers are responsible for ensuring a smooth and safe trip (e.g., through regular vehicle maintenance and being physically fit to drive), yet accidents regularly occur in the course of service provision. For example, a recent survey of rideshare drivers by Shannon et al. (2024) found that one-third of them had already been involved in a work-related crash. Many of the main accident causes – such as cell phone use while driving, fatigue, and navigating unfamiliar roads – are under the driver's control.[7] Nevertheless, although Uber and Lyft do not directly provide the rides themselves, they frequently face liability claims from harmed passengers (see e.g. GAO, 2024). This raises the question of the extent to which these platforms should be held legally responsible for harm suffered by consumers in such accidents.[8]

Framework and results Against this background, we consider a theoretical framework in which a platform cannot directly influence product safety, but can indirectly affect the expected number of accidents through its choice of consumer price. This price determines overall demand for the potentially harmful good, which in turn shapes sellers' incentives regarding product safety. We analyze the interactions of four types of agents: a policymaker, a monopolistic platform, a seller, and consumers. The seller offers a product or service that consumers can purchase through the platform at a price set and collected by the platform. For each unit sold, the seller receives a commission from the platform. The product is potentially harmful and the likelihood of an accident depends on the level of product safety chosen by the seller. The policymaker decides on the liability rule which divides the harm from an accident between the platform and the seller (strict partial liability), whereas harmed consumers are fully made whole in the course of applying

[4]Rather than charging users directly, social platforms primarily generate revenue from selling advertisements shown to users. In addition to managing different groups of users, shopping platforms such as Amazon often also act as sellers themselves, thereby competing with third-party sellers. For a systematic discussion of different platform types and their business models, see, e.g., Belleflamme and Peitz (2021).

[5]For example, Uber reported an average of 34 million trips per day worldwide for the last quarter of 2024, and Lyft reported 26.1 million active riders (defined as taking at least one trip per quarter) in the second quarter of 2025 (Uber, 2024; Lyft, 2025).

[6]The exact distribution of the fare between platforms and drivers is not publicly known and seems to vary over time. For example, Uber's share has recently been estimated to lie between 16-20 percent of the fare, see e.g. UCLA Labor Center (2023). Zhang et al. (2019) argue that Uber has a strong bargaining strong position in their negotiations with drivers.

[7]As a further example, Lyft reported 111 motor vehicle fatalities between 2020 and 2022, an increase of over 14% compared to the period 2017- 2019, see Lyft (2024), "Safety Transparency Report (2020-2022)".

[8]For an overview of current legal practice in different states in the U.S. for accidents involving ride-sharing platforms (see e.g. The Super Lawyer, 2025).

the liability rule.

Providing high product safety is costly for the seller, and the incentive to do so comes not only directly from the liability regime, but also indirectly from the interaction in the market: The lower the price set by the platform, the higher the demand for the product, the higher the seller's expected liability costs (for a given liability share), and hence the stronger its incentive to provide high product safety. This market channel hence establishes a positive relationship between market demand and product safety.

In turn, the platform's optimal pricing strategy is driven by the following trade-off: When setting the monopoly price, the platform maximizes revenue from consumers. However, the relatively low demand may lead the seller to provide only poor product safety, resulting in a higher number of accidents and, consequently, increased liability costs for the platform. Alternatively, the platform may set a price below the monopoly price in order to stimulate demand, thereby inducing the seller to improve product safety and which in turn would reduce the platform's own liability costs. Given this trade-off, we characterize the platform's optimal pricing behavior and analyze how it depends on crucial factors such as the liability rule or the seller's cost of product safety provision. We show that the platform will optimally choose monopoly pricing when its share of liability is either high or low. Intuitively, in the former case, the platform has little incentive to reduce accidents, so the benefits of a deviation from monopoly pricing (lower liability costs due to higher product safety) are outweighed by the corresponding revenue loss. Conversely, when the platform's liability share is high (and the seller's is low), incentivizing the seller to provide high product safety is either infeasible or would require a too large price cut by the platform. By contrast, for intermediate levels of platform liability, the expected liability costs are significant for both the platform and the seller. In this case, the platform optimally sets a price below the monopoly price, thereby inducing the seller to provide high product safety (as long as doing so is not prohibitively costly for the seller).

We then determine the optimal liability rule set by a welfare-maximizing policymaker. The liability rule influences (i) the overall demand for the good (through the platform's pricing decision), and thus allocative efficiency in the market, and (ii) the frequency of accidents for a given quantity (through the level of product safety provided by the seller). We show that the optimal liability rule is dichotomous: it either assigns no liability to the platform (and hence full liability to the seller), or it assigns the maximum liability share to the platform that just leads it to set a price which induces the seller to provide high product safety. Intuitively, for a given level of product safety, the policymaker minimizes allocative inefficiency by selecting the liability rule that results in the lowest platform price. No liability is assigned to the platform if either the cost of providing high product safety are so low that the monopoly price of the platform for $\beta = 0$ is below the price implied by the maximum liability share that makes the platform induce high product safety, or so high that the platform does not induce high product safety regardless of β. For cost values in-between the second option, that is, a positive value of the liability share, is socially optimal.

We also study a model extension allowing for across-group externalities that are typical for many platform goods (Belleflamme and Peitz, 2021). In particular, consider the case in which consumers' willingness to pay for the platform good depends on the number of active sellers. For example, in our context of ride-hailing platforms, a larger number of drivers/vehicles might reduce the average waiting time for consumers, which makes the service more valuable to them. This extension also requires to endogenize the number of active sellers. Each active seller's profit now depends on the overall number of active sellers on the platform, determined by a zero-profit (fixed-point) condition. We show that the main qualitative properties of the baseline model are preserved.

Related literature Our paper is related to the following strands of literature: Firstly, we contribute to the by now vast economic literature studying platforms. In their seminal work, Rochet and Tirole (2003) introduced a framework to analyze the pricing and interactions between distinct user groups within a platform, emphasizing the importance of cross-side network effects. Armstrong (2006) built on this by exploring competition between platforms, offering insights into how platforms can differentiate themselves, and the implications for market structures and welfare. There exists by now a vast literature in this area, studying further issues such as platform design, multi-homing, bundling, matching, exclusivity, dynamics, and regulation (for overviews, see e.g. Belleflamme and Peitz, 2021; Jullien, Pavan, and Rysman, 2021).

Our paper is most closely related to a recent theoretical literature that studies whether platforms should be held responsible for harm suffered by consumers, caused by products they have purchased via the platform (either directly sold by the platform or by third-party sellers).[9] In the framework of Hua and Spier (2025), holding platforms liable incentivizes them to monitor and remove sellers of harmful products, but it may also reduce sellers' incentives to invest in product safety. Hua and Spier (2025) find sharing liability between the platform and sellers becomes optimal, in particular when sellers might be judgment-proof, leading to inefficient entry decisions. In Hua and Spier (2025), there is no consumer pricing, but the transaction fee between the platform and sellers plays a crucial role, while in our model, it is just the reverse, i.e. we focus on the case where the consumer price is set by the platform, taking the transaction fee as given. Hence, the two models apply to different platform types.

In Zennyo (2023), it is the platform's decision how to allocate liability for consumer harm between itself and the sellers of potentially harmful products. However, a regulator can set a minimum level of liability for the platform. Compared to a regime without platform liability, this reduces seller's precaution incentives, but leads more sellers to enter the market. As a result, the effect of the regulation on consumer surplus is ambiguous and depends on the relative importance of these countervailing effects.

In Hua and Spier (2023), the platform directly sells the product to consumers (no third party sellers), thereby also generating revenue from advertisement. It then depends on the size of the resulting network effects whether the consumer price charged by the platform is zero or strictly positive. Hua and Spier (2023) show that, strict liability outperforms negligence when the optimal price is zero, as it induces more consumers to join the platform. By contrast, there is no difference between the two rules when the optimal price is strictly positive.

De Chiara et al. (2024) also study the interplay between a monopolistic platform, sellers and consumers as in our model. In contrast to our paper, there are no pricing decisions. Rather, they explore the role of reputational sanctions to induce platforms and sellers to maintain safety and quality standards. Also in Yasui (2022), sellers' incentives to provide high product safety arise from reputation concerns. Under platform liability, these incentives might be crowed out, as the threat of liability increases the platform's incentive to increase its monitoring intensity. Another strand of literature on platform liability focuses on other important aspects such as preventing infringements of intellectual property rights or the distribution of otherwise illegal content (for recent policy discussions from a law & economics perspective, see e.g. Buiten, de Streel, and Peitz, 2020; Lefouili and Madio, 2022).

Secondly, our paper is related to the literature on product liability (both theoretical and empirical), which studies the role of liability in fostering producers' incentives to improve the safety of existing products or to develop new and safer ones (see e.g. McGuire, 1988; Viscusi and Moore, 1993; Daughety and Reinganum, 1995; Hay and Spier, 2005; Schwartzstein and Shleifer, 2013; Galasso and Luo, 2022; Dawid and Muehlheusser, 2022). This literature has also stressed the role of market forces as a potentially impor-

[9]See Busch (2021) for a comparative legal analysis of platform liability in the U.S. and in the EU.

tant driver of product safety. Intuitively, when sellers' liability is weak or even absent, consumers anticipate that they will have to cover the lion's share of the harm, leading to a low willingness to pay for unsafe products (see e.g. Polinsky and Shavell, 2010; Dawid and Muehlheusser, 2022; Dawid et al., 2024; Hua and Spier, 2025). In our setting, although consumers do not face any liability risk, a seller's incentive to improve product safety also arises "via the market" and increases with the demand for the product.

Finally, our paper is also related to a recent literature studying the issue of liability for accidents involving autonomous systems, such as autonomous vehicles (AVs). From a legal perspective, this raises a number of novel questions (see e.g. Colonna, 2012; Gless, Silverman, and Weigend, 2016; Geistfeld, 2017; Wagner, 2019; Guerra, Parisi, and Pi, 2022a; Buiten, De Streel, and Peitz, 2023; Buiten, 2024; Di, Dawid, and Muehlheusser, 2025). Several recent papers study more specific game-theoretic models of AV liability, focusing on issues such as negligence in a setting of bilateral care (Friedman and Talley, 2019; Di, Chen, and Talley, 2020), the timing of AV market introduction and market penetration (Dawid and Muehlheusser, 2022) or AV adoption incentives (De Chiara et al., 2021; Dawid et al., 2024).[10] Our model could also be adapted to a setting of AV liability, with the (monopolistic) producer of AVs in the role of the platform, interacting with consumers (as buyers and passengers/drivers of AVs) and other traffic participants such as pedestrians, whose level of precaution also affects the likelihood of accidents.

The remainder of this paper is organized as follows: The model framework is presented in Section 2, and the main results are derived in Section 3. In Section 4 we study the model extension with across-group externalities on the side of consumers and an endogenous number of active sellers. Section 5 discusses some (policy) implications and concludes. The Appendix contains the proofs and additional figures.

2 Model framework

General setup We study the interaction between four types of agents (all risk neutral): a policymaker, a platform, a seller and consumers. The seller offers a good or service which consumers can purchase via the platform at a price $p \geq 0$ set and reaped by the platform. In turn, the platform pays the seller a per-unit fee $\alpha > 0$ for providing the good. Consumers differ with respect to their valuation of the product which leads to a demand function $Q(p)$.[11] The good is potentially harmful to consumers in the course of accidents, and the level of harm per accident is fixed and normalized to 1. The likelihood of an accident depends on the level of product safety provided by the seller.[12] The seller can either provide a low level of product safety that comes with a high accident probability $\bar{x} \in (0, 1]$. Alternatively, the seller can provide a high level of product safety, so that the accident probability is reduced to $\underline{x} < \bar{x}$. For example, the seller might (not) perform additional safety checks or provide safety training for staff before the product is made available to consumers or while in use. The seller's cost of implementing accident probability x is $C(x)$, where $C(\underline{x}) = 0$ and $C(\bar{x}) = c > 0$.[13]

A liability rule stipulates how the harm from an accident is divided between the parties involved. In our

[10] All of these papers consider a setting of *mixed traffic*, where AVs interact with human-driven vehicles. Optimal liability rules in a setting with only AVs one the streets are studied, for example, in Shavell (2020), Guerra, Parisi, and Pi (2022b) and Schweizer (2024).

[11] In Section 4, we consider a model extension with many potential sellers that are heterogeneous with respect to their entry cost, deciding upon market entry so that the number of active sellers in equilibrium becomes endogenous.

[12] We restrict attention to the case of unilateral precaution where only the seller, but not the platform, can affect product safety directly. See Section 5 for a discussion.

[13] In many cases the seller might not directly choose the level of product safety; rather this will result from the number or intensity of (costly) precaution measures. Since one would naturally assume that more precaution leads to higher product safety (at least in expectation), for the sake of notational simplicity we can assume without loss of generality that the seller chooses the level of product safety directly. Note that in our model, the effectiveness of the safety investment is independent of the actual demand for the good or service.

context, a natural starting point is the case where consumers (who have no influence on product safety in our framework) initially suffer the damage, but are fully made whole, so that the harm is ultimately divided between the platform and the seller. We denote by β and $(1-\beta)$ the liability shares of the harm accruing to the platform and the seller, respectively, where $\beta \in [0,1]$.

Our model focuses on the policymaker's choice of liability regime (β), the platform's pricing decision (p) and the seller's choice of product safety (x), thereby taking the per-unit fee $\alpha > 0$ paid by the platform to the seller as given.[14] There is no asymmetric information in the model, so that each player's choice is observed by the other players. There is only ex ante uncertainty with respect to the occurrence of accidents. In the following, we discuss the objectives, decisions and payoffs of the different players in more detail.

Platform For a given liability rule β, the platform chooses the consumer price $p \geq 0$ to maximize its expected profit. For each unit sold at price p, the platform obtains a payoff $p-\alpha$, so that the total revenue is $(p-\alpha)\cdot Q(p)$. On the cost side, the platform faces expected liability costs which depend on (i) the total number of units sold ($Q(p)$), (ii) the product safety per unit (x, determined by the seller) , and (iii) the platform's liability share (β, determined by the policymaker). There is no cost of operating the platform. The platform's objective is hence

$$\max_p \ \Pi_x(p;\beta) = (p-\alpha-x\cdot\beta)\cdot Q(p). \tag{1}$$

Seller For a given liability rule β and a given consumer price p, the seller aims to maximize expected profit by choosing the level of product safety determining the accident probability, $x \in \{\underline{x}, \bar{x}\}$. Apart from the eventual cost of product safety provision, $C(x)$, the seller obtains the fee α from the platform for each unit sold, and faces the remaining liability costs in each accident. The seller's objective is hence

$$\max_x \ \pi_x(p;\beta) = (\alpha - x\cdot(1-\beta))\cdot Q(p) - C(x). \tag{2}$$

Consumers Consumers differ with respect to their willingness to pay for the product. In the spirit of Hotelling (1929), they are distributed on a subset of the non-negative real line, and the product offered by the seller is located at the left boundary point 0. All consumers have the same gross valuation $\gamma > 0$ for the product. Moreover, a consumer at location θ_i faces preference (or travel) costs $t\cdot\theta_i$, where the parameter $t > 0$ measures the strength of consumer heterogeneity. Consumer i's net utility is then

$$u_i(p) = \gamma - t\cdot\theta_i - p, \tag{3}$$

Since a consumer is fully compensated in case of an accident, their utility directly depends neither on the level of product safety (x) nor on the liability rule (β). If follows from (3) that only consumers with $\theta_i \in [0, \gamma/t]$ might end up buying the product if the price is sufficiently low, and all other consumers are irrelevant for the further analysis. We assume that a mass γ/t of consumers is uniformly distributed in this relevant interval $[0,\gamma/t]$ (which, as shown below, gives rise to a demand function that is linear in p).

[14]This assumption is less restrictive than it might appear. For example, α might reflect the relative bargaining power of the platform and the seller when sharing the surplus created. All we need is that the value of α allows both the platform and the seller to earn non-negative profits (see Assumption 1 below). Moreover, the modeling approach ensures that the platform is not the residual claimant, in which case it would fully internalize the seller's incentives regarding product safety.

Policymaker The policymaker chooses the liability rule β to maximize expected social welfare:

$$\max_{\beta} \; W_x(Q(p)) = \int_0^{Q(p)} (\gamma - t \cdot \theta)\, d\theta - C(x) - x \cdot Q(p). \tag{4}$$

The three terms capture the consumer surplus in the market, the cost of product safety, and the expected harm from accidents, respectively. Note that the liability rule β does not directly affect social welfare, as it only specifies the division of consumers' harm from accidents between the platform and the seller. However, it will do so indirectly by affecting the platform's pricing decision and hence demand $Q(p)$, as well as the seller's choice of product safety and the resulting accident probability (x).

The following assumption ensures that the subsequent analysis focuses on economically meaningful settings.

Assumption 1.

(i) The cost of high product safety satisfies $c < c^{max} := \frac{(\bar{x}-\underline{x})\gamma}{t}$.

(ii) The consumers' gross valuation γ *fulfills* $\gamma > 2\bar{x}$.

(iii) The per-unit fee α *satisfies* $\alpha \in [\bar{x}, \gamma - \bar{x})$.

The first condition implies that the cost of high product safety is lower than the maximal potential gain in terms of a lower accident risk, which is $\frac{(\bar{x}-\underline{x})\gamma}{t}$.[15] The other two conditions guarantee that for any given liability rule $\beta \in [0, 1]$ and accident probability $x \in \{\underline{x}, \bar{x}\}$ both the platform and the seller expect non-negative profits under optimal platform pricing.[16]

The sequence of events is as follows: At stage 1, the policymaker decides on the liability rule (β). At stage 2, the platform decides on the consumer price (p). At stage 3, the seller decides on the level of product safety determining the per-unit accident probability (x). At stage 4, for a given price and level of product safety, each consumer decides whether not to purchase the product. Afterwards, accidents occur as random events, and consumers are compensated for their harm by the platform and the seller according to the liability rule chosen at stage 1.

3 Analysis

In this section, we determine the subgame perfect equilibrium outcome of the game, using backward induction.

3.1 Stage 4: Optimal consumer behavior

Consider first optimal consumer behavior at stage 4 for a given price p. Using (3), all consumers with sufficiently small values of θ obtain a strictly positive net utility from the product and will purchase it. The indifferent consumer θ^* is given by the condition $\gamma - t \cdot \theta^* - p = 0$, and all consumers with $\theta_i \in [0, \theta^*]$

[15] The gain from providing high instead of low product safety is maximal when all consumers (a mass γ/t) purchase the good, and each of them enjoys the reduction of the accident probability from $\bar{x}$ to $\underline{x}$.

[16] To see this, note that the maximal marginal cost of the seller (i.e., the marginal liability cost under low product safety and $\beta = 0$) is given by $\bar{x}$, so that $\alpha \geq \bar{x}$ ensures non-negative profits under the optimal choice of product safety. Considering the platform, we have for the marginal costs that $\alpha + \beta x < \gamma - \bar{x} + \bar{x} = \gamma$. Therefore, the platform is a monopolist with marginal costs strictly below the reservation price γ for any $\beta \in [0, 1]$ and accident probability $x \in \{\underline{x}, \bar{x}\}$, which implies that it can generate a positive profit under optimal pricing.

will purchase the good. This leads to a demand function

$$Q(p) = \theta^* = \frac{\gamma - p}{t}, \tag{5}$$

that is linearly decreasing in the price p, increasing in the gross valuation γ and decreasing in consumers' taste sensitivity t. Since consumers are fully made whole in case of an accident, demand is independent of the level of product safety (x) and the liability rule (β).

3.2 Stage 3: The seller's optimal choice of product safety

Consider next the seller's optimal choice of product safety at stage 3 for a given price p and liability rule β. From (2), the seller can either provide low product safety at no cost, leading to high accident probability $\bar{x}$ and profit $\pi_{\bar{x}}(p)$. Alternatively, it can provide high product safety at cost $c > 0$, thereby reducing the accident probability to $\underline{x} < \bar{x}$, which leads to profit $\pi_{\underline{x}}(p)$. It follows that $\pi_{\underline{x}}(p) \geq \pi_{\bar{x}}(p)$ if and only if

$$(\alpha - \underline{x} \cdot (1 - \beta)) \cdot Q(p) - c \geq (\alpha - \bar{x} \cdot (1 - \beta)) \cdot Q(p). \tag{6}$$

The incentive condition (6) implicitly defines a maximum consumer price $\bar{p}(\beta)$ such that the seller is willing to provide high product safety if and only if $p \leq \bar{p}(\beta)$. Substituting the demand function (5), the threshold price $\bar{p}(\beta)$ is given by

$$\bar{p}(\beta) = \gamma - \frac{t \cdot c}{(\bar{x} - \underline{x}) \cdot (1 - \beta)}. \tag{7}$$

Denoting the seller's optimal decision by $x^*(p, \beta)$, we hence have $x^*(p, \beta) = \underline{x}$ for $p \leq \bar{p}(\beta)$, and $x^*(p, \beta) = \bar{x}$ otherwise. Intuitively, the seller's incentive to provide high product safety is not only provided by the liability rule β, but also by the market: A high demand for the product (induced by a low price) leads to higher expected liability costs for the seller, which in turn increases the seller's incentive to invest in product safety. The market interaction between the platform, consumers and the seller hence leads to a positive relationship between consumer demand and product safety.

Importantly, $\bar{p}(\beta)$ decreases with β: A higher liability share β for the platform lowers the seller's remaining share $1 - \beta$, so that a larger demand (and hence a lower price) is required to induce the seller to provide high product safety. Furthermore, $\bar{p}(\beta)$ decreases with the cost of providing a safe product (c): If c increases, then a higher quantity and therefore a lower price is needed to make it optimal for the seller to provide high product safety.

As a next step we investigate under which conditions the seller can be induced to provide high product safety *at all*:

Lemma 1. *There exists a liability threshold $\bar{\beta} = 1 - \frac{t \cdot c}{\gamma(\bar{x} - \underline{x})} > 0$ such that the seller will never optimally provide high product safety for $\beta > \bar{\beta}$. That is, if $\beta > \bar{\beta}$, then $x^*(p, \beta) \equiv \bar{x}$ for all $p \geq 0$.*

The threshold $\bar{\beta}$ is derived from $\bar{p}(\bar{\beta}) = 0$. Intuitively, for $\beta > \bar{\beta}$, the seller's liability share $(1 - \beta)$ is so low that even under maximum demand, the seller will not find it optimal to provide high product safety. Moreover, $\bar{\beta} > 0$ holds as long as the cost of providing high product safety is not excessive (i.e. for $c < c^{max}$ as discussed above).

3.3 Stage 2: Optimal platform pricing

We now study optimal pricing of the platform under a given liability rule β, anticipating consumer choice and the seller's incentives regarding product safety. The platform trades off the revenue it obtains from selling the product to consumers and the expected liability costs. These costs also depend on the level of product safety chosen by the seller, which in turn is driven by consumer demand.

The platform can set the monopoly price, thereby extracting the maximum revenue from consumers. However, at such a high price, the resulting consumer demand might not suffice to induce the seller to provide high product safety which, in turn, would lead to high expected liability costs for the platform. Alternatively, the platform could set a price below the monopoly price, thereby increasing demand and fostering the seller's incentive to provide high product safety. In this case, while the platform faces a revenue loss by deviating from monopoly pricing, also its expected liability costs are lower. With this basic intuition in mind, we will next formally characterize optimal platform pricing, and show how it is affected by the liability rule β and other crucial model parameters.

Denote by $p_x^m(\beta)$ the monopoly price for a given level of product safety $x \in \{\underline{x}, \bar{x}\}$ and a given liability rule $\beta \in [0,1]$. Substituting the demand function (5) into the platform's profit function (1) and maximizing w.r.t p yields

$$p_x^m(\beta) = \frac{1}{2}(\gamma + \alpha + \beta \cdot x) \quad \forall x \in \{\underline{x}, \bar{x}\}. \tag{8}$$

Note that $p_x^m(\beta)$ increases with β. Moreover, $p_{\underline{x}}^m(0) = p_{\bar{x}}^m(0)$ and $p_{\underline{x}}^m(\beta) < p_{\bar{x}}^m(\beta)$ for all $\beta > 0$. Intuitively, when the platform faces no liability ($\beta = 0$), then it does not care about product safety. Only for $\beta > 0$ will the platform take product safety into account. Moreover, for any $\beta > 0$, the monopoly price is higher when the "marginal" (liability) costs $\beta \cdot \bar{x}$ is high due to low product safety.

In a next step we investigate under which conditions the platform will optimally depart from monopoly pricing. Recall from Lemma 1 that for $\beta > \bar{\beta}$, the seller can never be induced to provide high product safety. In this case, the platform has obviously no incentive to deviate from the monopoly price $p_{\bar{x}}^m(\beta)$. However, for $\beta \leq \bar{\beta}$, a high product safety can in principle be induced, and the question is whether the seller would provide it under the monopoly price $p_{\underline{x}}^m(\beta)$, i.e. if $p_{\underline{x}}^m(\beta) \leq \bar{p}(\beta)$. If this is the case, there is again no need for the platform to deviate from monopoly pricing – this time at $p_{\underline{x}}^m(\beta)$– as it gets high product safety "for free". However, if $p_{\underline{x}}^m(\beta) > \bar{p}(\beta)$, the platform faces a trade-off between extracting the maximum revenue from consumers and reducing the liability cost from a higher product safety, and in this case a deviation from monopoly pricing might be optimal.

It is instructive to decompose the effect on the platform's profit when it deviates from monopoly pricing to $\bar{p}(\beta)$, the highest price for which the seller would provide high product safety.[17] On the one hand, such a deviation imposes a revenue loss that, for given product safety x, can be expressed as

$$\begin{aligned}
\Pi_x^L(\beta) \quad &:= \quad \Pi_x(\bar{p}(\beta)) - \Pi_x(p_{\bar{x}}^m(\beta)) = -\left(\frac{1}{t}\right) \cdot (p_{\bar{x}}^m(\beta) - \bar{p}(\beta))^2 \\
&= \quad -\left(\frac{1}{4t}\right) \cdot \left(-\gamma + \alpha + \bar{x} \cdot \beta + \frac{2t \cdot c}{(\bar{x} - \underline{x}) \cdot (1 - \beta)}\right)^2 < 0.
\end{aligned} \tag{9}$$

The last expression emerges when substituting for $\bar{p}(\beta)$ and $p_x^m(\beta)$ from (7) and (8), respectively. Note that $\Pi_x^L(\beta)$ is strictly negative for all $\beta \in [0, \bar{\beta}]$, strictly concave in β, and reaches its maximum of 0 only if β is such that $\bar{p}(\beta) = p_x^m(\beta)$. In particular the absolute value of the loss increases with β whenever $\bar{p}(\beta) < p_x^m(\beta)$. This follows from the observation that $\bar{p}(\beta)$ decreases with β, while $p_x^m(\beta)$ increases with β.

[17] Note that only a deviation to $\bar{p}(\beta)$ needs to be considered. For a given x, the platform's profit is strictly monotone increasing in p for all $p < p_x^m(\beta)$. Therefore, for the case $p_{\underline{x}}^m(\beta) > \bar{p}(\beta)$ considered here, if the platform wants to induce high product safety ($\underline{x}$), it will optimally deviate downward from the monopoly price as little as necessary to just ensure this (i.e. it will move to $\bar{p}(\beta)$).

On the other hand, by choosing $\bar{p}(\beta)$ the platform benefits from lower expected liability costs due to higher product safety. This potential gain is

$$\Pi^G(\beta) \quad := \quad \Pi_{\underline{x}}(\bar{p}(\beta)) - \Pi_{\bar{x}}(\bar{p}(\beta)) = \left(\frac{\beta}{1-\beta}\right) \cdot c, \tag{10}$$

it satisfies $\Pi^G(0) = 0$, and increases with the platform's liability share β. Intuitively, when the platform does not face any liability ($\beta = 0$), it does not gain from inducing higher product safety. But as β increases this leads to higher expected liability cost for the platform, which in turn increases the benefit from high product safety. Importantly, however, the platform reaps this gain only when the price deviation to $\bar{p}(\beta)$ induces the seller to actually *increase* the product safety provided; it is zero when high product safety already emerges under monopoly pricing.

In summary, the change in the platform's profit when charging $\bar{p}(\beta)$ instead of the (incentive compatible) monopoly price[18] is given by

$$\Delta\Pi(\beta) = \begin{cases} \Pi^L_{\underline{x}}(\beta) & \text{for } \ \bar{p}(\beta) \geq p^m_{\underline{x}}(\beta), \\ \Pi^L_{\bar{x}}(\beta) + \Pi^G(\beta) & \text{for } \ \bar{p}(\beta) < p^m_{\underline{x}}(\beta). \end{cases} \tag{11}$$

Since $\bar{p}(\beta)$ decreases with β, while $p^m_{\underline{x}}(\beta)$ increases, a necessary condition for $p^m_{\underline{x}}(\beta)$ to induce high safety for some β is that it does to for $\beta = 0$. The next results characterizes for which values of c high safety is induced by monopoly pricing for $\beta = 0$.

Lemma 2. *There exists a threshold for the cost of providing high product safety, $\bar{c} := \frac{(\gamma-\alpha)\cdot(\bar{x}-\underline{x})}{2t}$, implicitly defined by the condition $\bar{p}(0) = p^m_x(0)$, such that $\bar{p}(0) > (<) p^m_x(0)$ if and only if $c < (>) \bar{c}$.*

Figure 1: Threshold price $\bar{p}(\beta)$ and monopoly prices $p^m_{\underline{x}}(\beta)$ and $p^m_{\bar{x}}(\beta)$ depending on liability rule β for different cost of providing high product safety c

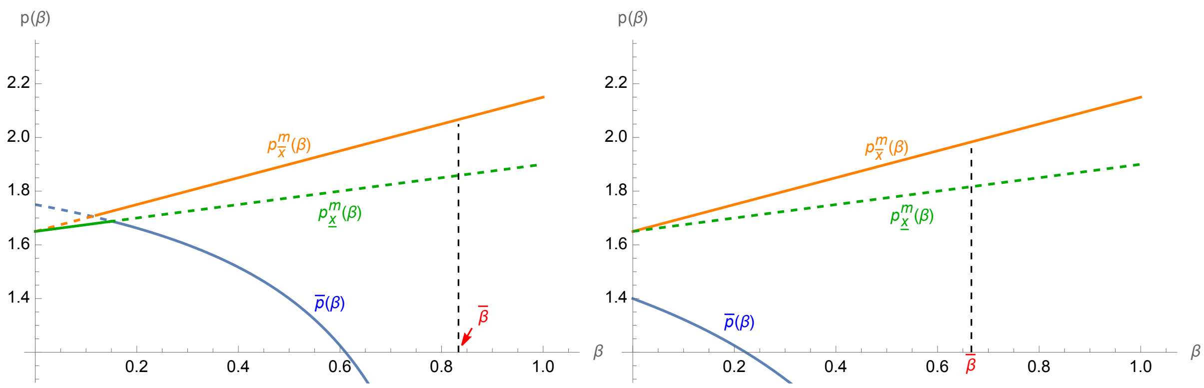

(a) Low cost: $c < \bar{c}$ (b) High cost: $c > \bar{c}$

Numerical specification for both panels: $t = 0.7$, $\gamma = 2.1$, $\alpha = 1.2$, $\bar{x} = 1$, $\underline{x} = 0.5$, leading to $\bar{c} = 0.32$. Moreover, $c = 0.25 < \bar{c}$ and $c = 0.5 > \bar{c}$ in the left and right panel, respectively. In both panels, a dotted segment of $p^m_{\underline{x}}(\beta)$ or $p^m_{\bar{x}}(\beta)$ indicates that this price corresponds to a level of product safety that is inconsistent with the seller's incentive condition (6).

The lemma is illustrated in Figure 1. Recall from (7) and (8) that the maximum price for which the seller provides high product safety, $\bar{p}(\beta)$, decreases with β, while the monopoly price $p^m_x(\beta)$ increase with β for all $x \in \{\underline{x}, \bar{x}\}$, where $p^m_{\underline{x}}(0) = p^m_{\bar{x}}(0)$.

[18] There is an interval of β values where both $p^m_{\underline{x}}(\beta)$ and $p^m_{\bar{x}}(\beta)$ are incentive compatible (see Figure 1). In this case the platform always prefers to choose $p^m_{\underline{x}}(\beta)$.

Panel (a) depicts the case $c < \bar{c}$ so that $\bar{p}(0) > p^m_{\bar{x}}(0)$ and for β sufficiently small, the seller would provide high product safety under the monopoly price $p^m_{\underline{x}}(\beta)$. Intuitively, when the seller's liability share is high (i.e. when β is small) and the cost of high product safety c is sufficiently small, then the seller will optimally provide it even for the relative small consumer demand induced by the platform setting the monopoly price (i.e., $\Pi^G(\beta) = 0$ in this case).

By contrast, when $c > \bar{c}$ as depicted in panel (b) of Figure 1, $\bar{p}(0) < p^m_{\bar{x}}(0)$ holds. In this case, for any given $\beta \in [0, 1]$, under monopoly pricing the platform cannot induce the seller to provide high product safety. For the sake of visual clarity, in both panels of Figure 1, a dotted segment of $p^m_{\underline{x}}(\beta)$ or $p^m_{\bar{x}}(\beta)$ indicates that this price corresponds to a level of product safety that is inconsistent with the seller's incentive condition (6). For example, in the left panel, $p^m_{\bar{x}}(\beta)$ can never be optimal for small β as the seller optimally provides high product safety under this price, while $p^m_{\underline{x}}(\beta)$ is only optimal under low safety. By contrast, in the right panel, $p^m_{\underline{x}}(\beta)$ is infeasible throughout as the seller prefers low product safety for all values of β.

With respect to optimal platform pricing, denoted by $p^*(\beta)$, we first establish some features that apply to both of the two cases delineated in Lemma 2 and illustrated in Figure 1. The specific features of each case are then addressed in Propositions 1 and 2 below.

Lemma 3. *(i) For any $c \neq \bar{c}$, monopoly pricing is optimal for the platform for sufficiently small $\beta > 0$. In particular, we have $p^*(\beta) = p^m_{\underline{x}}(\beta)$ for $c < \bar{c}$ and $p^*(\beta) = p^m_{\bar{x}}(\beta)$ for $c > \bar{c}$ and sufficiently small $\beta > 0$. For $c = \bar{c}$, we have $p^*(\beta) = \bar{p}(\beta)$ for sufficiently small $\beta > 0$. (ii) For $\beta \geq \bar{\beta}$, the platform optimally chooses $p^*(\beta) = p^m_{\bar{x}}(\beta)$.*

As for part (i), as long as $c \neq \bar{c}$, the platform will optimally deviate from monopoly pricing and choose $\bar{p}(\beta)$ only if $\Delta\Pi(\beta) \geq 0$. Since $\Pi^G(0) = 0$, such a deviation never pays off for $\beta = 0$ as there is no gain but only a loss. As a result, for $\beta = 0$ the platform always prefers monopoly pricing, either at $p^m_{\underline{x}}(\beta)$ (when $c \leq \bar{c}$) or at $p^m_{\bar{x}}(\beta)$ (when $c > \bar{c}$). By continuity, these arguments remain valid for β sufficiently small.

The case $c = \bar{c}$ is non-generic in the sense that for $\beta = 0$, $p^m_{\bar{x}}(0) = \bar{p}(0)$ holds, and so there is no loss for the platform when choosing $\bar{p}(\beta)$, thereby ensuring high product safety. Hence, $\Pi^L_{\underline{x}}(0) = \Pi^G(0) = 0$ for $c = \bar{c}$. Furthermore, the marginal loss from deviating to $\bar{p}(0)$ is also zero, while the slope with respect to β of the gain function is positive. Therefore, for $\beta > 0$ but sufficiently small, the platform strictly prefers to stick to $\bar{p}(\beta)$ as choosing the monopoly price $p^m_{\bar{x}}(\beta)$ instead would entail a loss in product safety that outweighs the gain in revenue (see Figure 2(a) for an illustration). As for part (ii), recall from Lemma 1 that the seller cannot be induced to provide high product safety for $\beta > \bar{\beta}$. As a result, $\Pi^G(0) = 0$ and so the platform has no incentive to deviate from monopoly pricing.

We now consider the more specific features of optimal pricing for each of the two cases delineated in Lemma 2, starting with the one in which the seller's cost of providing high product safety is low (i.e. $c < \bar{c}$, see left panel of Figure 1):

Proposition 1. *When the seller's cost of providing high product safety is low ($c < \bar{c}$), the platform's optimal price is*

$$p^*(\beta) = \begin{cases} p^m_{\underline{x}}(\beta) & \text{for } \beta \in [0, \tilde{\beta}], \\ \bar{p}(\beta) & \text{for } \beta \in [\tilde{\beta}, \hat{\beta}], \\ p^m_{\bar{x}}(\beta) & \text{for } \beta \in [\hat{\beta}, 1], \end{cases} \tag{12}$$

where $0 \leq \tilde{\beta} \leq \hat{\beta} \leq \bar{\beta}$.

The proposition is illustrated in Figure 2(a). The cost of providing high product safety is sufficiently low such that the seller would provide it for $\beta = 0$, despite the relatively low demand resulting under monopoly pricing. As was shown in Lemma 3 above, in this case there is no reason for the platform to deviate from the corresponding monopoly price $p^m_{\underline{x}}(\beta)$ for $\beta = 0$, and also as long as β remains sufficiently small (indicated by the threshold $\tilde{\beta}$).

As β increases slightly beyond $\tilde{\beta}$, charging the monopoly price $p^m_{\underline{x}}(\beta)$ would cause a downward jump in the platform's profit as the seller would now provide low product safety only (a first-order effect). Instead, pricing at $\bar{p}(\beta)$ still induces a high level of product safety at only a slightly lower price (a second-order effect). By continuity, $\bar{p}(\beta)$ remains optimal for a whole range of values of β sufficiently close to $\tilde{\beta}$ (indicated by the threshold $\hat{\beta}$). However, as β increases beyond $\hat{\beta}$, the loss from deviating to $\bar{p}$ eventually becomes too large and the platform's optimally switches back to monopoly pricing, this time at $p^m_{\bar{x}}(\beta)$. Finally, for $\beta \geq \bar{\beta}$, high product safety can no longer be induced (see Lemma 1). Hence, the monopoly price $p^m_{\bar{x}}(\beta)$ is optimal for the platform also in this range.

One implication of the proposition is that that the imposition of platform liability does not necessarily lead to more product safety, but it might actually increase it. To see this, note that high product safety is only provided along the equilibrium path for $\beta \leq \hat{\beta}$. For all $\beta > \hat{\beta}$, the optimal platform price will jump upwards to $p^m_{\bar{x}}(\beta)$, and the resulting reduction in demand makes it optimal for the seller to provide low product safety only.

Figure 2: Optimal platform pricing (in bold), depending on liability rule (β) and seller's cost of providing high product safety (c)

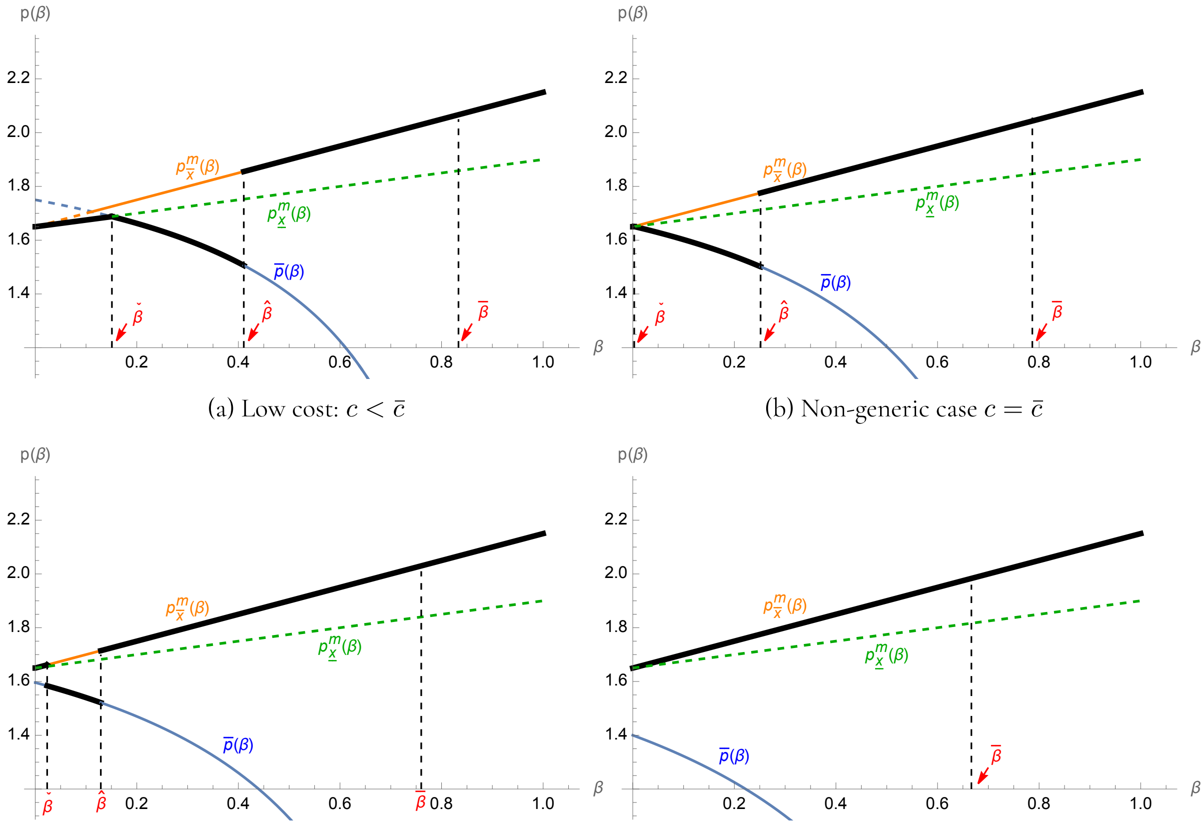

(a) Low cost: $c < \bar{c}$

(b) Non-generic case $c = \bar{c}$

(c) High cost: $c - \bar{c}$ small

(d) Very high cost: $c > \bar{\bar{c}}$

Numerical specification for all panels: $t = 0.7$, $\gamma = 2.1$, $\alpha = 1.2$, $\bar{x} = 1$, $\underline{x} = 0.5$, leading to threshold $\bar{c} = 0.32$. The four panels differ with respect to the value of the costs of providing high product safety, c, relative to the threshold $\bar{c} = 0.32$. Panel (a): $c = 0.25 < \bar{c}$, leading to $\check{\beta} = 0.151$, $\hat{\beta} = 0.41$, and $\bar{\beta} = 0.83$. Panel (b): non-generic case $c = 0.32 = \bar{c}$, leading to $\check{\beta} = 0$, $\hat{\beta} = 0.248$, and $\bar{\beta} = 0.786$. Panel (c): $c = 0.36 > \bar{c}$, leading to $\check{\beta} = 0.0235$, $\hat{\beta} = 0.129$, and $\bar{\beta} = 0.76$. Panel (d): $c = 0.5 > \bar{c}$, leading to $\bar{\beta} = 0.67$ (note that $\check{\beta}$ and $\hat{\beta}$ do not exist in this case). In all panels, a dotted segment of $p^m_{\underline{x}}(\beta)$ or $p^m_{\bar{x}}(\beta)$ indicates that this price corresponds to a level of product safety that is inconsistent with the seller's incentive condition (6).

Consider next the case in which the seller's cost of providing high product safety is high (i.e. $c > \bar{c}$, see right panel of Figure 1):

Proposition 2. *When seller's the cost of providing high product safety is high ($c \geq \bar{c}$), then either $p^*(\beta) = p^m_{\underline{x}}(\beta)$ for all $\beta \in [0,1]$, or there exist an interval of β values for which the platform optimally chooses $\bar{p}(\beta)$. More precisely, in the latter case there exist thresholds $0 \leq \check{\beta} < \hat{\beta} \leq \bar{\beta}$ such that*

$$p^*(\beta) = \begin{cases} p^m_{\underline{x}}(\beta) & \text{for } \beta \in [0, \check{\beta}], \\ \bar{p}(\beta) & \text{for } \beta \in [\check{\beta}, \hat{\beta}], \\ p^m_{\underline{x}}(\beta) & \text{for } \beta \in [\hat{\beta}, 1]. \end{cases} \tag{13}$$

For sufficiently large values of c the optimal price is given by $p^*(\beta) = p^m_{\underline{x}}(\beta)$ for all $\beta \in [0,1]$. For $c = \bar{c}$ we have $0 = \check{\beta} < \hat{\beta}$ and for $c - \bar{c}$ sufficiently small the interval of β values such that $p^*(\beta) = \bar{p}(\beta)$ is non-empty. Furthermore, there exists $\bar{\bar{c}} \in (\bar{c}, c^{max})$ such that the interval $[\check{\beta}, \hat{\beta}]$ is empty for all $c > \bar{\bar{c}}$.

The proposition is illustrated in Figure 2(b) – (d). To gain an intuition, recall first that for $c > \bar{c}$, the seller will not provide high product safety under monopoly pricing by the platform, so that the relevant monopoly price is $p^m_{\underline{x}}(\beta)$. As shown in Lemma 3, for β small, the platform does not care so much about product safety since the burden of liability is low, so that a deviation from the monopoly price is not profitable. Similarly, for β large, inducing high product safety by the seller either becomes very costly for the platform (requiring a huge price cut) or even unfeasible (for $\beta > \bar{\beta}$). Again, in these cases there is no deviation from the monopoly price $p^m_{\underline{x}}(\beta)$. If at all, a deviation to $\bar{p}(\beta)$ arises for intermediate values of β. This depends on how far the cost of high product safety (c), exceeds the threshold $\bar{c}$. When c is sufficiently close to $\bar{c}$, there exists an interval $[\check{\beta}, \hat{\beta}]$ for which it is beneficial for the platform to induce high product safety by choosing $\bar{p}(\beta)$. This case is illustrated in Figure 2(c). By contrast, when c is too large, such an interval does not exist, and the platform optimally prices at $p^m_{\underline{x}}(\beta)$ throughout (see Figure 2(d)).

3.4 Stage 1: Optimal liability rule

Propositions 1 and 2 have characterized how the liability rule affects the platform's optimal pricing, $p^*(\beta)$, and the seller's optimal choice of product safety, $x^*(p^*(\beta), \beta)$. Factoring these optimal choices into the policymaker's objective (4) yields the objective function

$$\max_{\beta} \ W = \int_0^{Q(p^*(\beta))} (\gamma - t \cdot \theta)\, d\theta - C(x^*(p^*(\beta), \beta)) - x^*(p^*(\beta), \beta) \cdot Q(p^*(\beta)). \tag{14}$$

There are two key factors for the policymaker, quantity and product safety. First, with respect to quantity, since the platform is a monopolist, there is the usual downward distortion compared to the efficiency benchmark under perfect competition. When the platform resorts to monopoly pricing, the size of this quantity distortion increases with β (recall from Eq. (8) that $p^m_x(\beta)$ increases with β). By contrast, when the platform departs from monopoly pricing and chooses the price $\bar{p}(\beta)$ instead, the quantity distortion is smaller than under the respective respective monopoly price. The reason is that the platform optimally chooses $\bar{p}(\beta)$ only if $\bar{p}(\beta) < p^m_x(\beta)$ for a given β (see Propositions 1 and 2). Moreover, the quantity distortion under $\bar{p}(\beta)$ decreases with β (recall from Eqn. (7) that $\bar{p}(\beta)$ itself decreases with β, so that the quantity $Q(p)$ increases with β in this case).

Second, the level of product safety provided by the seller (possibly at a cost c) affects social welfare via the expected harm from accidents. When the platform prices at $\bar{p}(\beta)$, then by definition of $\bar{p}(\beta)$, high product safety is always induced. By contrast, under monopoly pricing, high product safety is only provided under $p^m_{\underline{x}}(\beta)$, but not under $p^m_{\bar{x}}(\beta)$. As shown above, apart from the cost of providing high product safety (c), it then depends on the liability rule β, which of the two monopoly prices is consistent with the seller's optimal choice of product safety, and which price is ultimately optimal for the platform.

Recall from Propositions 1 and 2 that for $\beta = \hat{\beta} > 0$ the platform is indifferent between setting $p^*(\beta) = \bar{p}(\beta)$ (thereby inducing the seller to provide high product safety) and $p^*(\beta) = p^m_{\bar{x}}(\beta)$ (thereby inducing the seller to provide low product safety). Denoting the optimal liability rule by β^*, we obtain the following result:

Proposition 3.

(i) The policymaker optimally either does not impose any liability on the platform, or she imposes the maximum liability share for which the platform would optimally choose the price $\bar{p}(\beta)$, i.e. $\beta^* \in \{0, \hat{\beta}\}$.

(ii) When the seller's cost of providing high product safety is low ($c < \bar{c}$), then $\beta^* = 0$ for $p^m_{\underline{x}}(0) < \bar{p}(\hat{\beta})$ and $\beta^* = \hat{\beta}$ otherwise.

(iii) When the seller's cost of providing high product safety is high, but not too much so (i.e. $c \geq \bar{c}$, but $c - \bar{c}$ sufficiently small), then $\beta^* = \hat{\beta}$.

(iv) When the cost of providing high product safety is very high (i.e. $c > \bar{\bar{c}}$), then $\beta^* = 0$.

To gain an intuition, Figure 2 is again useful. Because of the quantity distortion due to the platform's market power, the policymaker seeks to induce any given level of product safety $x \in \{\underline{x}, \bar{x}\}$ at the lowest possible price.

As for part (i), recall from Propositions 1 and 2 that for a given liability rule β, the platform optimally either chooses one of the two monopoly prices or the price $\bar{p}(\beta)$. Since both $p^m_{\bar{x}}(\beta)$ and $p^m_{\underline{x}}(\beta)$ increase with β for all $x \in \{\underline{x}, \bar{x}\}$, so does the quantity distortion. Therefore, under monopoly pricing by the platform, setting $\beta = 0$ would be optimal for the policymaker (i.e. no platform liability). Moreover, $\bar{p}(\beta)$ decreases with β and so does the quantity distortion. In this case, the policymaker would hence optimally choose $\beta = \hat{\beta}$, i.e. the largest value of β for which the platform would choose $\bar{p}(\beta)$. Whether $\beta = 0$ or $\beta = \hat{\beta}$ is ultimately optimal for the policymaker depends on whether high effort is induced by the monopoly price for $\beta = 0$, and if this is the case, whether $p^m_{\bar{x}}(0)$ is smaller or larger than $\bar{p}(\hat{\beta})$. In turn, this depends on the seller's cost of providing high product safety, c, as well as on the other model parameters. When $c < \bar{c}$ as considered in part (ii), we know from Proposition 1 that under both $\beta = 0$ and $\beta = \hat{\beta}$ the platform chooses a price that induces high effort and the relevant comparison is between $p^m_{\underline{x}}(0)$ and $\bar{p}(\hat{\beta})$. For the case depicted in Figure 2(a), it turns out that $\hat{\beta}$ is optimal for the policymaker, but there also exist parameterizations for which $\beta^* = 0$ holds.

In the remaining parts of Proposition 3, the seller's cost of providing high product safety exceeds the threshold $\bar{c}$. In these cases, setting $\beta = 0$ induces the platform to choose a price that induces low product safety by the seller. As we know from the previous section, there are two cases to consider: Firstly, as for part (iii), c is sufficiently close to $\bar{c}$, so that there exist liability rules for which the platform optimally prices at $\bar{p}(\beta)$. In this case, to induce the lowest possible price in the market, the policymaker optimally chooses $\beta^* = \hat{\beta}$ (see also Figure 2(c)).[19] Secondly, as for part (iv), c is sufficiently large such that the platform

[19] Inducing high product safety is always optimal for the social planner in this range of c. The reason is that under a price of $\bar{p}$ the seller's share of the reduction in expected accident costs is equal to c, which implies that the welfare gain of high product safety (which equals the total reduction in expected accident costs) is larger than c.

prefers the monopoly price $p^m_{\bar{x}}(\beta)$ for all $\beta \in [0,1]$, and the seller will only provide low product safety. In this case, the policymaker optimally chooses $\beta^* = 0$ to minimize the welfare loss due to the quantity distortion (see Figure 2(d)).

3.5 The impact of the cost of high product safety on equilibrium behavior

In a next step, we study in more detail the effect of the cost of high product safety (c) on behavior along the equilibrium path. The results are illustrated in Figure 3. Panel (a) illustrates again the non-monotone effect of c on the optimal liability rule (see the discussion of Proposition 3): When c is either very small or very large, the optimal liability rules shields the platform from liability and all the burden is placed on the seller ($\beta^* = 0$, and hence constant in c). For all other values of c, the optimal liability rule is $\beta^* = \hat{\beta}$, which decreases with c. Intuitively, a higher value of c ceteris paribus reduces the seller's willingness to provide high product safety. To countervail this incentive, the policymaker optimally shifts a larger liability share from the platform to the seller.

Figure 3 (b) illustrates the welfare effect of c along the equilibrium path. When c is very small, the seller provides high product safety even under the monopoly price $p^m_{\underline{x}}(\beta)$. In this range, both price and safety level are unaffected by c, and an increase in c leads to a direct negative (linear) welfare effect. As c increases further, to retain the seller's incentive to provide high product safety, the platform departs from monopoly pricing and chooses $\bar{p}(\beta)$ instead which, from (7), decreases in c thereby leading to higher demand and a lower quantity distortion. This gives rise to an additional positive (and also linear) indirect welfare effect. Furthermore, as a third effect, as just seen in panel (a), a higher c leads to a lower optimal liability share $\hat{\beta}$ for the platform, which in turn increases $\bar{p}(\beta)$. The interplay of these three effects gives rise to the observed hump-shaped welfare effect of c. Finally, for sufficiently large values $c > \bar{\bar{c}}$, the platform chooses the monopoly price $p^m_{\bar{x}}(\beta)$, the policymaker optimally assigns all liability to the seller (to minimize the quantity distortion), and the seller provides low product safety only. As a result, welfare is constant in c in this parameter range.

Figure 3 (c) and (d) illustrate the effect of c on the optimal platform pricing and profit. Intuitively, recall from (8) that $p^m_{\underline{x}}(\beta)$ is independent of c, so that there is no effect when monopoly pricing prevails (i.e. for c small and for $c > \bar{\bar{c}}$). By contrast, when the platform prices at $\bar{p}(\beta)$, the last two channels driving the welfare effect emerge also here, i.e. the (negative) effect of c on $\bar{p}(\beta)$ and in the indirect (positive) effect via $\hat{\beta}$, which render $\bar{p}(\beta)$ U-shaped in c.

The platform's profit is maximum under monopoly pricing and no liability (which occurs when c either very small or very large). As c moves away from zero, the platform eventually switches to price $\bar{p}(\beta)$, and the optimal liability rule jumps from $\beta^* = 0$ to $\hat{\beta} > 0$, leading to a (discontinuous) drop in the platform's profit. As c increases further, the platform's profit increases with c since the gain from a lower liability share (recall that $\hat{\beta}$ is decreasing in c) outweighs the reduction in the price $\bar{p}(\beta)$. In the range of c where $\bar{p}(c)$ increases with c both the effect of an increase in c on liability and on price contribute to higher seller profits.

Figure 3: Optimal liability rule, welfare, platform pricing and profit in equilibrium, depending on seller's cost of providing high product safety (c)

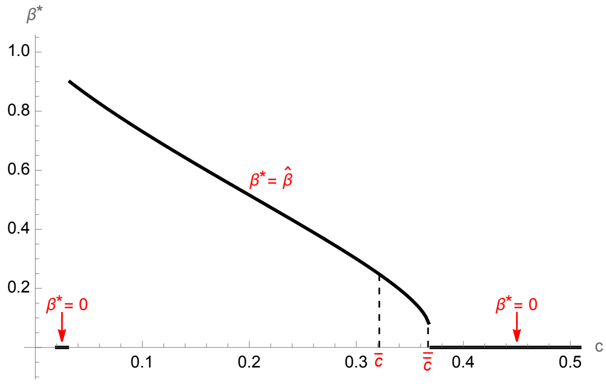

(a) Optimal liability rule

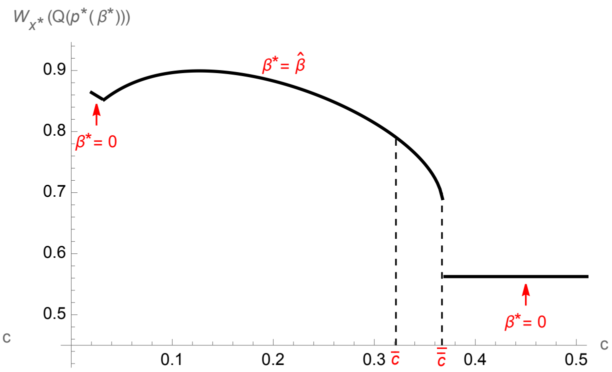

(b) Welfare

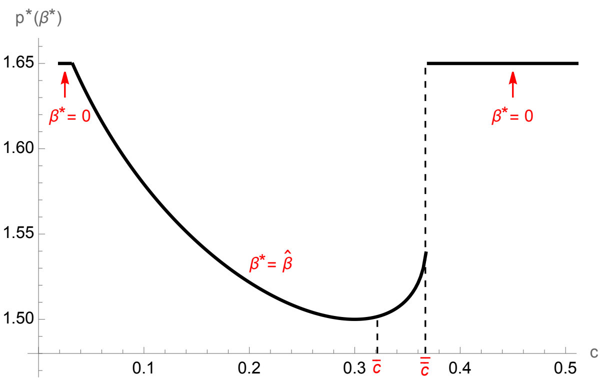

(c) Platform pricing

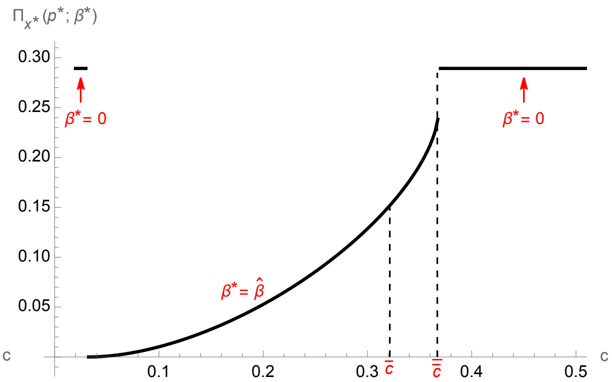

(d) Platform profit

Numerical specification for all panels: $t = 0.7$, $\gamma = 2.1$, $\alpha = 1.2$, $\bar{x} = 1$, $\underline{x} = 0.5$, leading to thresholds $\bar{c} = 0.32$ and $\bar{\bar{c}} = 0.37$.

4 Extension: Platform externalities

THIS SECTION IS UNDER CONSTRUCTION

- In the model considered so far, consumers' valuation for the good was unaffected by the supply side of the market. However, in many platform markets there exist across-group externalities in the sense that, say, a consumer's valuation for the platform good depends on characteristics on the supply side such as the number of active sellers or the variety of goods offered by them (see e.g. Belleflamme and Peitz, 2021). For example, in the context of ride-hailing platforms, a larger number of drivers/vehicles might reduce the average waiting time for consumers, which makes the service more valuable to them.

- To capture such externalities, in this section we consider an extension where consumers' valuation increases with the number of sellers active on the platform. As there was only one seller in the baseline model, we also need to allow for a larger (and endogenous) number of active sellers on the platform.

5 Discussion and conclusion

This paper examines platform liability in a setting where a platform controls the price paid by consumers, a third-party seller determines product safety, and a policymaker sets the liability rule. Ride-hailing platforms serve as an illustrative example. We characterize the platform's optimal pricing strategy as a trade-off between maximizing consumer revenue and minimizing expected liability costs, which depend on the seller's choice of product safety. This trade-off leads to deviations from monopoly pricing for intermediate levels of platform liability, where both the platform and the seller face significant expected liability costs. In such cases, the platform optimally sets a price below the monopoly level to increase demand and hence also the seller's incentive to improve product safety. We also find that the optimal liability regime is dichotomous: either the platform is fully exempt from liability or it bears the maximum liability share that still incentivizes the seller to provide high product safety. We also consider an extension with across-group externalities on the side of consumers and an endogenous number of active sellers on the platform.

Our framework might be useful to address some further issues in the context of platform liability: Firstly, in our setting, product safety is determined by the seller(s), a natural starting point since they produce the good or provide the service. In practice, however, also platforms can influence expected harm, for instance, through (costly) screening measures which help to identify and remove "bad" sellers (see e.g., Hua and Spier, 2023; Peitz, 2025). In the ride-hailing context, for example, a July 2022 lawsuit filed in the San Francisco County Superior Court on behalf of 550 women alleged assaults and related offenses by Uber drivers, accusing the company of inadequate safety measures and insufficient background checks, and of prioritizing "growth over customer safety (BBC, 2022). The framework of Hua and Spier (2025) considers bilateral investment (but not consumer pricing), and the two investments choices are strategic substitutes, i.e. more investment by the platform crowds out sellers' incentives. In contrast, in our setting with consumer pricing the two investments seem to be complements: Intuitively, a higher safety investment by the platform reduces its expected liability costs, leading to a lower monopoly price. In turn, this fosters consumer demand which ceteris paribus also improves the seller's investment incentives. Moreover, the lower monopoly also lowers the "distance" to the price $\bar{p}$ (i.e. the highest price for which the seller would provide high product safety). This makes a downward deviation to this price less costly for the platform, and hence increases the parameter range for which doing so is optimal.

Secondly, our model confines attention to the case of a monopolistic platform which raises the issue of platform competition. Karle, Peitz, and Reisinger (2020) provide a theoretical framework linking market structure (i.e. whether there is one dominant platform or rather an oligopolistic structure) to the intensity of product market competition between sellers, which depends on the degree of differentiation of their products. An oligopolistic platform structure arises when sellers' products are relatively homogeneous, as multiple platforms allow sellers to soften competition between them. To the best of our knowledge, the issue of potentially harmful products and platform liability has not yet been addressed in this literature.

References

Armstrong, Mark. 2006. "Competition in two-sided markets." *The RAND Journal of Economics* 37 (3):668–691.

Belleflamme, Paul and Martin Peitz. 2021. *The economics of platforms*. Cambridge University Press.

Buiten, Miriam, Alexandre De Streel, and Martin Peitz. 2023. "The law and economics of AI liability." *Computer Law & Security Review* 48:105794.

Buiten, Miriam C. 2024. "Product liability for defective AI." *European Journal of Law and Economics* 57:239–273.

Buiten, Miriam C., Alexandre de Streel, and Martin Peitz. 2020. "Rethinking liability rules for online hosting platforms." *International Journal of Law and Information Technology* 28:139–166.

Busch, Christoph. 2021. "Rethinking product liability rules for online marketplaces: A comparative perspective." In *TPRC49: The 49th Research Conference on Communication, Information and Internet Policy*.

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

Daughety, Andrew F and Jennifer F Reinganum. 1995. "Product safety: Liability, R&D, and Signaling." *American Economic Review* 85 (5):1187–1206.

Dawid, Herbert, Xuan Di, Peter M Kort, and Gerd Muehlheusser. 2024. "Autonomous vehicles policy and safety investment: an equilibrium analysis with endogenous demand." *Transportation Research Part B: Methodological* 182:102908.

Dawid, Herbert and Gerd Muehlheusser. 2022. "Smart products: Liability, investments in product safety, and the timing of market introduction." *Journal of Economic Dynamics and Control* 134:104288.

De Chiara, Alessandro, Idoia Elizalde, Ester Manna, and Adrian Segura-Moreiras. 2021. "Car accidents in the age of robots." *International Review of Law and Economics* 68:106022.

De Chiara, Alessandro, Juan José Ganuza, Fernando Gómez, Ester Manna, and Adrian Segura-Moreiras. 2024. "Platform Liability with Reputational Sanctions." *UB Economics Working Papers, NO. E24/467* .

Di, Xuan, Xu Chen, and Eric Talley. 2020. "Liability design for autonomous vehicles and human-driven vehicles: A hierarchical game-theoretic approach." *Transportation Research Part C: Emerging Technologies* 118:102710.

Di, Xuan, Herbert Dawid, and Gerd Muehlheusser. 2025. "When Engineering Meets Economics: AI-Powered Safe and Accountable Autonomous Driving." *Artificial Intelligence for Transportation* 1 (2):100020.

Friedman, Ezra and Eric Talley. 2019. "Automatorts: How should accident law adapt to autonomous vehicles? Lessons from Law & Economics." *Columbia University, mimeo* .

Galasso, Alberto and Hong Luo. 2022. "When does product liability risk chill innovation? Evidence from medical implants." *American Economic Journal: Economic Policy* 14 (2):366–401.

Geistfeld, Mark A. 2017. "A roadmap for autonomous vehicles: State tort liability, automobile insurance, and federal safety regulation." *California Law Review.* 105:1611–1694.

Gless, Sabine, Emily Silverman, and Thomas Weigend. 2016. "If Robots Cause Harm, Who is to Blame? Self-Driving Cars and Criminal Liability." *New Criminal Law Review* 19 (3):412–436.

Guerra, Alice, Francesco Parisi, and Daniel Pi. 2022a. "Liability for robots I: legal challenges." *Journal of Institutional Economics* 18 (3):331–343.

———. 2022b. "Liability for robots II: an economic analysis." *Journal of Institutional Economics* 18 (4):553–568.

Hay, Bruce and Kathryn Spier. 2005. "Manufacturer Liability for Harms Caused by Consumers to Others." *American Economic Review* 95 (5):1700–1711.

Hotelling, Harold. 1929. "Stability in competition." *Economic Journal* 39 (153):41–57.

Hua, Xinyu and Kathryn Spier. 2023. "Platform safety: Strict liability versus negligence." *HKUST Business School Research Paper, No. 2023-105* .

———. 2025. "Holding platforms liable." *American Economic Journal: Microeconomics* forthcoming.

Jullien, Bruno, Alessandro Pavan, and Marc Rysman. 2021. "Two-sided markets, pricing, and network effects." In *Handbook of Industrial Organization*, vol. 4. Elsevier, 485–592.

Karle, Heiko, Martin Peitz, and Markus Reisinger. 2020. "Segmentation versus agglomeration: competition between platforms with competitive sellers." *Journal of Political Economy* 128 (6):2329–2374.

Lefouili, Yassine and Leonardo Madio. 2022. "The economics of platform liability." *European Journal of Law and Economics* 53 (3):319–351.

Lyft. 2025. "Lyft reports record Q2 2025 financial results." `https://s27.q4cdn.com/263799617/files/doc_financials/2025/q2/FINAL-Q225-Press-Release.pdf`.

McGuire, Edward. 1988. "The Impact of Product Liability." *The Conference Board* Report 908.

Peitz, Martin. 2025. "Governance and Regulation of Platforms." In *Handbook of New Institutional Economics*. Springer, 565–593.

Polinsky, Mitchell and Steven Shavell. 2010. "The Uneasy Case for Product Liability." *Harvard Law Review* 123:1437–1492.

Regibeau, Pierre. 2025. "Current challenges in competition policy." *Journal of Institutional and Theoretical Econonomics (JITE)* forthcoming.

Rochet, Jean-Charles and Jean Tirole. 2003. "Platform competition in two-sided markets." *Journal of the European Economic Association* 1 (4):990–1029.

Schwartzstein, Joshua and Andrei Shleifer. 2013. "An activity-generating theory of regulation." *The Journal of Law and Economics* 56 (1):1–38.

Schweizer, Urs. 2024. "Liability for accidents between road users whose activity levels are verifiable." *Journal of Institutional and Theoretical Economics (JITE)* forthcoming.

Shannon, Brett, Lee Friedman, Andrew Hellinger, Kirsten Almberg, and Johnathon Ehsani. 2024. "Work-related crashes in rideshare drivers in the United States." *Journal of Safety Research* 89:13–18.

Shavell, Steven. 2020. "On the redesign of accident liability for the world of autonomous vehicles." *The Journal of Legal Studies* 49 (2):243–285.

Spier, Kathryn E. and Rory Van Loo. 2025. "Foundations for platform liability." *Notre Dame Law Review, forthcoming* .

Uber. 2024. "Uber announces results for fourth quarter and full Year 2024." `https://investor.uber.com/news-events/news/press-release-details/2025/Uber-Announces-Results-for-Fourth-Quarter-and-Full-Year-2024/default.aspx`.

Viscusi, Kip and Michael Moore. 1993. "Product Liability, Research and Development, and Innovation." *Journal of Political Economy* 101 (1):161–184.

Wagner, Gerhard. 2019. "Robot, inc.: personhood for autonomous systems?" *Fordham Law Review* 88:591.

Yasui, Yuta. 2022. "Platform liability for third-party defective products." *SSRN Working Paper No. 4224267* .

Zennyo, Yasuke. 2023. "Should platforms be held liable for defective third-party goods?" *CPRC Discussion Paper Series, CPDP-93-E* .

Zhang, Qingyu, Reza Langari, H Eric Tseng, Dimitar Filev, Steven Szwabowski, and Serdar Coskun. 2019. "A Game Theoretic Model Predictive Controller With Aggressiveness Estimation for Mandatory Lane Change." *IEEE Transactions on Intelligent Vehicles* 5 (1):75–89.

Online References

BBC. 2020. "ISIS 'still evading detection on Facebook', report says." URL https://www.bbc.com/news/technology-53389657.

———. 2022. "Uber sued in US over sexual assault claims." URL https://www.bbc.com/news/business-62158976.

CNN. 2022. "Uber releases safety data: 998 sexual assault incidents including 141 rape reports in 2020." URL https://edition.cnn.com/2022/06/30/tech/uber-safety-report.

CPSC. 2024. "CPSC finds Amazon responsible under Federal Safety Law for hazardous products sold by third-party sellers on Amazon.com." URL https://www.cpsc.gov/Newsroom/News-Releases/2024/CPSC-Finds-Amazon-Responsible-Under-Federal-Safety-Law-for-Hazardous-Products-Sold-by-Third-Party-Sellers-on-Amazon-com.

DeStatis. 2025. "European Union: 8 out of 10 internet users shop online." URL https://www.destatis.de/Europa/EN/Topic/Science-technology-digital-society/Onlineshopping_Products.html.

GAO. 2024. "Ridesharing and taxi safety: information on assaults against drivers and passengers." URL https://www.gao.gov/products/gao-24-106742.

Global Media Insight. 2025. "YouTube statistics 2025 (demographics, users by country and more)." URL https://www.globalmediainsight.com/blog/youtube-users-statistics/.

GovInfo. 2017. "Cohen v. Facebook." URL https://www.govinfo.gov/content/pkg/USCOURTS-nyed-1_16-cv-05158/pdf/USCOURTS-nyed-1_16-cv-05158-0.pdf.

Marketplace Pulse. 2025. "Amazon steers third-party seller share to all-time high." URL https://www.marketplacepulse.com/articles/amazon-steers-third-party-seller-share-to-all-time-high.

New York Post. 2024. "NY teen loses sight in one eye after buying 'defective' laser from Amazon: lawsuit." URL https://nypost.com/2024/11/30/us-news/ny-teen-loses-sight-in-one-eye-after-buying-laser-from-amazon-lawsuit/.

New York Times. 2025. "Uber's festering sexual assault problem." URL https://www.nytimes.com/2025/08/06/business/uber-sexual-assault.html.

NME. 2023. "Music publishers sue Amazon-backed AI platform for copyright infringement." URL https://www.nme.com/news/music/music-publishers-sue-amazon-backed-ai-platform-for-copyright-infringment-3518082.

Statista. 2025a. "E-commerce worldwide - statistics facts." URL https://www.statista.com/topics/871/online-shopping/.

———. 2025b. "In-store and e-commerce retail sales share in the United States from 2022 to 2028." URL https://www.statista.com/statistics/829220/share-of-retail-sales-by-channel-us/.

The Super Lawyer. 2025. “Who is liable in an Uber accident: the driver, Uber, or another party?” URL `https://thesuperlawyer.com/is-uber-liable-for-accidents/`.

UCLA Labor Center. 2023. “Analysis of high volume for-hire vehicle data for New York City.” URL `https://labor.ucla.edu/publications/for-hire-vehicle-data-new-york-city/`.

USA Today. 2022. “Uber driver accused of breaking into passenger’s home, raping her, after dropping her off.” URL `https://eu.usatoday.com/story/news/nation/2024/05/13/uber-driver-lemyrious-monaye-dawson-arrest-atlanta/73672932007/`.