Product–Consumer Substitution and Safety Regulation*

Konrad Grabiszewski[†] Alex Horenstein [‡]

April 2017

Abstract

We develop a theory of safety regulation where product safety and consumer skills are negatively correlated. Demand and supply drive this correlation. Given the option to choose any product, low skill consumers choose safer products. Given the option of selling to any consumer, producers chose to sell riskier products to high skill consumers. We validate our theory using a data set obtained from iRacing, an online racing simulator. A unique and important feature of our data is that it contains objective measures of product safety and consumer skills. This allows us to test our theory and analyze our policy implications.

Keywords: safety regulation; product safety; consumer skills; product-consumer substitution; moral hazard; adverse selection

*We thank Steve Myers, Executive Vice President and Executive Producer at iRacing, for sharing the data with us. We also thank him and Dean Marsh for helping us understand online racing and the iRacing data. We thank Gil de Ferran, professional racing driver and team owner, for helping us understand the world of racing. We thank Jeffrey Segal and his team at GPX Lab for helping us understand the differences between real-life and simulation racing. We thank Raphael Boleslavsky, Salvador Ortigueira, and Tomasz Piskorski for valuable comments. Nicolo Bates provided outstanding research assistance.

[†]Department of Economics, University of Miami, Coral Gables, FL 33124, USA; konrad.grabiszewski@gmail.com

[‡]Department of Economics, University of Miami, Coral Gables, FL 33124, USA; alexhorenstein@gmail.com

1 Introduction

We develop a theory of safety regulation and support it with an empirical analysis using over 2 million observations generated by iRacing, an online racing simulator. Safety regulation aims at making risky activities less dangerous by lowering the probability of a bad outcome. Examples of such activities include practicing sports (bad outcome is an injury or death), driving (bad outcome is an accident), and trading in financial markets (bad outcome is a loss of wealth).[1]

We divide safety regulation policies into two types. Product regulation determines the level of product safety. Safer products, ceteris paribus, reduce the likelihood of a bad outcome. Consumer regulation is about establishing specific consumer skills. Higher skills, ceteris paribus, make a risky activity less dangerous.[2] In other words, product safety regulates what is consumed while consumer regulation specifies who can consume.

The main concerns of the safety regulation literature are (a) whether higher product safety or consumer skills decreases the probability of a bad outcome[3] and (b) how people react to regulation. When regulation yields the intended result, that is, the probability of a bad outcome decreases, then we say that we observe the regular effect. However, when that probability increases, then we observe the Peltzman effect.[4]

[1]Safety regulation matters especially in the presence of behavioral biases and externalities. Biases might drive an agent to make sub-optimal decisions and externalities negatively affect those not directly involved in the agent's decision-making process. For example, overconfidence bias prompts drivers to behave more recklessly comparing to their true value of skills. Consequently, these drivers make the roads more dangerous and, if they cause an accident, this negatively affects other drivers.

[2]For instance, in the case of road safety, regulation is about reducing the probability of a fatal accident (bad outcome). Product regulation increases product safety by making cars safer (e.g., seat belts, driving aids, technological limits on the maximum speed a car can reach) or improving the road infrastructure (e.g., fewer potholes, impact-reducing barriers along the road). Consumer regulation screens out people with low driving capabilities and experience. This can be achieved by designing more difficult driving tests and requiring that the test be repeated every few years, or making sure that potential drivers have the capability to provide first aid to crash victims.

[3]Traditionally, in the safety regulation literature, a policy is evaluated (theoretically and empirically) by its impact on the probability of a bad outcome; the welfare analysis of regulation is omitted since what is observable is the safety measure rather than social welfare. In Appendix B, we discuss the problem of social optimality in the context of safety regulation.

[4]"The Peltzman effect" is a term we borrow from the literature on road safety. This effect was first studied by Peltzman (1975) who found that the product regulation introduced in the United States in the mid-1960s (e.g., installation of seat belts for driver and all passengers)"had no effect on the highway death toll. There is some evidence that regulation may have increased the share of this toll borne by pedestrians and increased the total number of accidents." The Peltzman effect is a potential problem for not only product regulation but also consumer regulation. For instance, in their summary of a large literature on driver education and its impact on road safety, Ker et al. (2005) conclude that there is "no evidence that driver education programmes are effective in preventing road traffic injuries or crashes."

Due to the changes in product safety or consumer skills, consumers adjust their effort (i.e., private safety precautions). This adjustment in behavior can be either positive (when consumers increase their effort) or negative (when effort decreases).

In this paper, we argue and empirically confirm that product safety and consumer skills are negatively correlated. When product becomes safer, then consumer skills (of representative consumer) decrease; and when consumer skills increase, then product (used by representative consumer) becomes less safe. We call this phenomenon the product-consumer substitution. Ours is the first paper to establish the presence of this substitution. In the standard theory of safety regulation introduced by Peltzman (1975), it is implicitly assumed that the correlation between product safety and consumer skills is zero.

Product-consumer substitution is driven by demand and supply, which strengthen each other. Demand-driven product-consumer substitution is due to the choices made by consumers: when the product becomes safer, then people with low skills who previously decided against consumption become consumers since, now, they finally can afford (in terms of cost of effort) to join the group of consumers (i.e., adverse selection). Supply-driven product-consumer substitution is due to the choices made by producers: less safe products are offered only to consumers with skills high enough.

As an example of the product-consumer substitution, driven by both demand and supply, consider the case of investors and financial instruments. Several papers establish the relevance of the demand-driven product-consumer substitution (adverse selection). van Rooij et al. (2011) find that people with lower levels of financial literacy are less likely to invest in complex financial instruments like stocks. According to Lusardi and Mitchell (2014), financial literacy captures "peoples' ability to process economic information and make informed decisions about financial planning, wealth accumulation, debt, and pensions." Amromin et al. (2013) show that "complex mortgages attract sophisticated borrowers." In our framework, financial literacy and sophistication are captured by the consumer skills. Complexity of financial instruments measures how difficult it is for an investor to understand and evaluate them (e.g., Treasury bonds are simpler than stock options).[5] That is, the inverse of complexity is product safety. The negative correlation between financial literacy and participation in financial markets is an example of the product-consumer

[5]Complex financial instruments properly used for hedging could decrease the overall risk of an investor's portfolio. However, using complex instruments without understanding the risks involved is likely to increase the risk faced by a low skill investor.

substitution driven by consumer choice.

The supply-driven product-consumer substitution is about the constraints associated with financial trading accounts. Stockbrokers impose limits on the complexity of assets their clients can trade based on their financial sophistication.[6] E-trade, an online broker, has four levels of trading suitability for trading stock options, which go from not allowing clients to trade options (level 0) to allowing all options strategies (level 3) (`https://www.optionshouse.com/blog/trading-strategies/trading-options-suitability-levels/`). To assess trading suitability, E-trade requires information such as years of experience trading options, yearly number of trades, and size of trades. Another broker, Interactive Brokers, has more specific rules for allowing a client to trade options. To be able to trade options, their client needs a specific liquid net worth and income wage depending on his/her age, two years of experience, and at least having performed 100 trades during his/her lifetime (`https://www.interactivebrokers.com/en/index.php?f=4945&p=tradingrequirements/`).

Neglecting product-consumer substitution has non-negligible consequences. First, disregarding this phenomenon might lead to erroneous conclusions. In particular, it is possible that while the standard theory predicts an increase (decrease) in effort or the probability of a bad outcome, we actually observe a decrease (increase).

Second, without product-consumer substitution, we obtain an incomplete understanding of safety regulation. In particular, we find that in the presence of product-consumer substitution, effort increases due to product regulation if and only if effort decreases due to consumer regulation. This result does not hold in the standard theory, but is important if a regulator considers consumer effort as part of his objective function.

Third, disregarding product-consumer substitution might lead to wrong regulation policies. To explain this observation, first, note that the standard theory is concerned only with the direct impact of regulation on effort which declines because product regulation makes people feel too safe and consumer regulation makes people feel overly confident about their skills. If that decline

[6] All brokers, as required by the Financial Industry Regulatory Authority (FINRA) rule 2111, recognize that "options involve risk and are not suitable for all investors" (see for example `https://www.optionshouse.com/blog/trading-strategies/trading-options-suitability-levels/`) and need to check for client's suitability. Therefore, all brokers require information from their clients to determine their level of sophistication and decide which assets their clients can or cannot trade (`http://finra.complinet.com/en/display/display_main.html?rbid=2403&element_id=9859`). However, FINRA is not a government regulator but a not-for profit self-regulatory organization. In that regard, the suitability rule can be interpreted as self-imposed by the FINRA members.

in effort is larger than, in absolute terms, the direct gain from regulation, then we observe the standard offsetting behavior. In contrast with the standard theory, we show that with the product-consumer substitution it is possible to observe two phenomena: (1) regular effect simultaneously with the standard offsetting behavior, and (2) the Peltzman effect without the standard offsetting behavior.

These two phenomena not only directly reject the standard theory but also show why relying on the standard theory might lead to wrong policies. Consider the first phenomenon and suppose that the regulator observes the regular effect. The regulator does not test for the presence of the standard offsetting behavior since, after all, the standard theory states that the regular effect and the standard offsetting behavior are mutually exclusive. Motivated by the initial success, the regulator might be tempted to introduce additional regulation to increase either product safety or consumer skills. However, once this additional regulation is implemented, the observed effect is surprisingly the Peltzman effect rather than the regular effect. Clearly, a wrong policy has been implemented.

Next, consider the second phenomenon. The regulator observes the Peltzman effect and, as indicated by the standard theory, his remedy is to implement an additional policy whose objective is to curb a decrease in effort. However, because there is no standard offsetting behavior, this additional policy has no impact and is just an inefficient allocation of resources. Again, the standard theory leads to wrong policy.

To address the problem of product-consumer substitution and its impact on policy design, we propose to design regulation as a two-dimensional policy: when implementing product regulation (consumer regulation), it is necessary to limit the decrease in consumer skills (product safety) by imposing a complementary consumer regulation (product regulation).

We empirically support our theory using data from iRacing (`http://www.iracing.com/`), an online car racing simulator. iRacing has been designed to replicate in the virtual world the three main components of physical racing—tracks, cars, and driver behavior—as realistically as possible. The only difference between the iRacing simulations and real world racing is that, in iRacing, nobody gets hurt or dies in an accident.

The purpose of using the iRacing data is *not* to analyze the problem of safety regulation in iRacing or, more generally, professional racing. We use the iRacing data because, given what and how data is collected, iRacing is an outstanding laboratory in which we can test our hypotheses. The iRacing data contains separate and objective measures of two key theoretical components, product safety and consumer skills. To the best of our knowledge, ours is the first empirical study using such measures.

Our data consists of several consumer skill and product safety regimes which permits us to go beyond the standard pre-and-post regulation analysis. As far as we know, only Cohen and Einav (2003) and Sobel and Nesbit (2007) use data with multiple regulation regimes.

With over 2 million observations, our data is robust to small-sample issues. In addition, our data is simulation-based and automatically collected, which avoids the usual problem of misreporting or missing data.

We are unaffected by one of the main empirical issues in the safety regulation literature: it takes time before a law is fully enforced, and this lack of immediate and full-enforcement might distort the empirical results. For example, Cohen and Einav (2003) recognize this problem in the case of the mandatory seat-belts. In iRacing, this problem is nonexistent.

Our hypotheses are of qualitative nature. Hence, our focus is not the magnitude of the estimated parameters but their sign. For example, it is not important that in iRacing the correlation between product safety and consumer skills (product-consumer substitution) is precisely -0.43; what matters is that this correlation is actually negative (and statistically significant).

We start with the traditional empirical exercise; i.e., testing for the Peltzman effect. In the case of both product and consumer regulation, we detect the regular effect. Then, we focus on hypotheses which differentiate ours from the standard theory. Except for the iRacing data, we are not aware of other data set used in the safety regulation literature that permits to test these hypotheses.

First, we analyze the problem of omitted variables. In the safety regulation literature, it is rather common to not control for either product safety or consumer skills. In particular, when testing for the Peltzman effect in the case of product regulation (consumer regulation), we exclude the measure of consumer skills (product safety) from the control variables.

We find that the omitted-variable bias is a serious problem. In each case, product regulation and consumer regulation, the relevant estimator is higher comparing to the estimation with the full set of controls. In fact, in the case of product regulation, we detect the presence of Peltzman effect when we do not control for consumer skills; however, once we control for consumer skills, we reject the hypothesis of Peltzman effect.

Our analysis of the omitted-variable bias leads us to detect the product-consumer substitution. We also show that this phenomenon is driven by both demand (adverse selection: low skill drivers choose safer cars) and supply (design of iRacing: low skill drivers do not have access to cars as unsafe as those which are available to high skill drivers).

Next, we focus on two hypotheses (mentioned above) which directly reject the standard theory. First, we find a sub-sample of our data with both the regular effect and the standard offsetting behavior. Second, we find a sub-sample with the presence of the Peltzman effect that is not due to the standard offsetting behavior.

Finally, we analyze a two-dimensional regulation that addresses the problems generated by the product-consumer substitution. To be more precise, we consider a sub-sample in which product regulation results with the Peltzman effect. However, when we analyze only drivers with high consumer skills, then the outcome of regulation turns out to be the regular effect. We conclude that, in this specific case, consumer regulation should complement product regulation in order to achieve the desired effect.

The problem of safety regulation is relevant in many fields. Some examples include, but are far from being limited to, recreational and professional sports (e.g., McCarthy and Talley (1999), Sobel and Nesbit (2007), Pope and Tollison (2010), Chong and Restrepo (2014)), consumer products (e.g., Viscusi (1996), Viscusi et al. (2005)), crime prevention,[7] and financial institutions and markets (e.g., Grossman (1992), Gorton and Huang (2004), Dam and Koetter (2012), Farhi et al. (2012), and Allen et al. (2015))[8].

In section 2, we develop our theory using a simple mathematical model. In section 3, we discuss in detail what iRacing is, describe how we translate the theory of safety regulation in the context of iRacing, and present the data we use in our empirical studies. Section 4 includes empirical results supporting our theory; in particular, we verify the hypotheses established in section 2. We end

[7]While Becker (1968) assumes that private crime-preventing expenditures decrease when public expenditures increase, Guha and Guha (2012) provide a model that derives the relationship between private and public expenditures.

[8]In fact, we can find the problem of bailouts and moral hazard discussed in work as early as Bagehot (1873).

with conclusions in section 5.

2 Theory

We develop a theory of safety regulation that adds the negative correlation between product safety and consumer skills (product-consumer substitution) to the standard theory. In the standard theory, this correlation is assumed to be zero which, as we show, is not an innocuous assumption. In the presence of product-consumer substitution, the standard theory might incorrectly predict the outcome of safety regulation as well as erroneously explain what drives that outcome. Consequently, relying on the standard theory could lead to invalid policy recommendations.

We use a simple mathematical model to motivate the theory. From the modelling perspective, the closest to our approach is Viscusi (2007); however, our mathematical model is more general and provides several new insights.

We highlight five hypotheses generated by our theory. Four of these hypotheses separate ours from the standard theory. We empirically confirm these hypotheses in section 4, where we also provide an empirical analysis of the two-dimensional policy we recommend in this paper.

2.1 Setup

Consider an activity that yields either a bad outcome (utility zero) or a good outcome (utility one). Let α denote the **product safety**; higher α means that the product becomes safer or easier to use. Let β denote the **consumer skills** of the representative consumer (he); higher β means that consumers become more capable and experienced. We think of α as capturing the safety of the representative product, that is, the product chosen by the representative consumer whose skills are represented by β.

We focus on *what* rather than *how* regulation achieves. That is, in the case of product (consumer) regulation, we are not interested in specific regulatory tools which increase product safety (consumer skills); rather, we are interested in the consequences of establishing a specific value of product safety (consumer skills) that the regulator targets. Consequently, product regulation means that the regulator chooses a specific value of α, and consumer regulation is synonymous with choosing a specific value of β.

We assume that regulations always achieve their intended level of product safety or consumer skills. That is, there is no issue of incomplete or delayed enforcement. Disregarding this issue is in accordance with the main research questions we ask in this paper as well as with our data in which we have complete and immediate enforcement.

Upon observing the level of regulation, a consumer chooses **effort** a where higher a means more effort. As it is standard in economics, we separate effort and skills. A low-skill consumer can compensate for his low β by exerting more effort, and a high-skill consumer can shirk his effort because of his high β. The cost of effort is $c(a)$ and the consumer is an expected-utility maximizer.

Consumer receives utility zero with probability λ and utility one with probability $1-\lambda$. **Probability of a bad outcome** λ is our safety measure. There are three factors that affect the probability of a bad outcome: product safety α, consumer skills β, and consumer effort a.

Our assumptions imposed on $c(a)$ and $\lambda(\alpha, \beta, a)$ are standard and straightforward. In particular, $c(a)$ is twice-continuously differentiable, strictly increasing ($c' > 0$), and convex ($c'' \geq 0$). We also assume that product safety, consumer skills, and effort are beneficial. That is, each of these three variables decreases the probability of a bad outcome; $\lambda_a = \frac{\partial \lambda}{\partial a} < 0$, $\lambda_\alpha = \frac{\partial \lambda}{\partial \alpha} < 0$, and $\lambda_\beta = \frac{\partial \lambda}{\partial \beta} < 0$. Finally, we assume that λ is a strictly convex function of a; i.e., $\lambda_{aa} = \frac{\partial^2 \lambda}{\partial a^2} > 0$.[9]

The main innovation of our theory is the negative correlation between product safety and consumer skills, which we call the **product-consumer substitution**. From the mathematical point of view, the product-consumer substitution means that β (consumer skills) is a function of α (product safety) such that the derivative $\frac{d\beta}{d\alpha}$ is negative. In addition, we assume that $\beta(\alpha)$ is a C^1 function (continuous, differentiable, and with continuous derivative). Since $\beta(\alpha)$ is injective, α can be expressed as a C^1 function of β with $\frac{d\alpha}{d\beta} < 0$.

According to the standard theory, the correlation between product safety and consumer skills is zero. We argue that this correlation is negative. This is exactly where we depart from the standard theory. We confirm the presence of product-consumer substitution in our data in section 4.2.

[9]This assumption is natural, especially if we analyze the probability of a good outcome. In the framework of safety regulation, $1 - \lambda$ is the revenue function from consumer effort a. As it is typical in economics, we assume that the revenue function is increasing and concave (in effort) which implies that λ must be decreasing and convex in a.

Hypothesis 1. *Product safety and consumer skills are negatively correlated.*

As we argued in Introduction, demand and supply are two forces behind the product-consumer substitution. In section 4.3, we detect both demand-driven and supply-driven product consumer substitution.

Hypothesis 2. *Product-consumer substitution is an outcome of demand and supply forces.*

1. *Demand-driven: less skilled consumers desire safer product.*

2. *Supply-driven: less skilled consumers are offered safer product.*

Knowing both α and β, the representative consumer chooses a in order to maximize his expected utility, $U(\alpha, \beta, a) = 1-\lambda(\alpha, \beta, a)-c(a)$. We focus on the unique interior solution of the optimization problem denoted by a^*.[10]

Let λ^* denote the probability of a bad outcome computed at optimal effort, $\lambda^* := \lambda(\alpha, \beta, a^*)$. From both empirical and theoretical perspectives, we are interested in how a^* and λ^* react to changes in regulation; these are the two fundamental problems addressed in the literature.

1. How does safety regulation change effort a^*?

2. How does safety regulation change the probability of a bad outcome λ^*?

Since the analysis of product regulation is the same as the analysis of consumer regulation, we discuss in detail only the former. Hence, hereafter we assume that the regulatory authorities increase product safety.

2.2 Safety regulation and effort

The optimal consumer effort a^* is a function of product safety and consumer skills; i.e., $a^* := a^*(\alpha, \beta)$. In addition, β is a decreasing function of α. We determine how a change in product safety α affects the optimal effort a^*.

$$\frac{da^*}{d\alpha} = \frac{\partial a^*}{\partial \alpha} + \frac{\partial a^*}{\partial \beta}\frac{d\beta}{d\alpha} \tag{1}$$

[10]We can either assume that the interior solution exists or impose the following standard assumptions, which guarantee that the interior solution exists: $\lim_{a\to 0} \frac{\partial U}{\partial a} > 0$ and $\lim_{a\to\infty} \frac{\partial U}{\partial a} < 0$.

Because of increased product safety, consumer changes his effort $\frac{da^*}{d\alpha}$; we call that change the **adjustment behavior**. If effort increases, then we have positive adjustment behavior and if effort decreases, then we talk about negative adjustment behavior (i.e., moral hazard).

While controversial for some authors, we believe that negative adjustment behavior is not surprising. In fact, it is a rather expected behavior. After all, product safety or consumer skills substitute consumer effort. Consider the following thought experiment. We take two different cars and analyze road safety. The first car is a typical modern vehicle that satisfies all regulatory standards. The second car is a fully autonomous car from the future. In the language of our theory, α is higher for the second car. Driving the future car is not really driving as we know today. The driver's role ends at providing directions to the vehicle. Then, he can relax, eat, read, or even sleep. His "driving" effort does not differ much from the effort the passengers of public transportation exert today. This zero effort is optimal because there is no reason for him to pay attention since his effort has no impact on whether or not there is a bad outcome. Hence, it is rather expected that higher α reduces effort.

Effort adjusts in two stages. First, there is the **primary adjustment behavior** $\frac{\partial a^*}{\partial \alpha}$. Upon observing an increase in product safety, the consumer modifies his effort. This is a direct impact of increasing product safety on effort. For instance, when cars become safer, the drivers might decide to lower their effort as the perceived safety increases. After all, even if an accident happens, the chances of getting hurt are smaller.

According the standard theory, adjustment behavior consists only of primary adjustment behavior. However, with the product-consumer substitution there also is the **secondary adjustment behavior**, $\frac{\partial a^*}{\partial \beta}$, which captures the change in effort due to the modification in consumer skills. This is an indirect impact of increasing product safety. For example, when the regulator makes cars safer, the drivers know that post-regulation less skilled people became drivers. Consequently, it is necessary to adapt to the new environment. As a precautionary measure, drivers may pay more attention as it is not wise to rely on other drivers to maintain a desired safety level; i.e., $\frac{\partial a^*}{\partial \beta}$ is negative.

A priori, neither the sign nor the size of the secondary adjustment behavior is clear.[11] Consequently, the standard theory incorrectly predicts the change in effort due to regulation. In particular, it is

[11] In Appendix A, we analyze the relationship between primary and secondary adjustment behaviors.

possible that effort increases (decreases) while the standard theory predicts a decrease (increase).

To elaborate, consider a laboratory test measuring how car safety affects effort. A driver is asked to test-drive cars with different safety/difficulty levels while the experimenter measures a variety of proxies for effort like heart rate, pulse, sweat, etc. Suppose that the observed effort decreases when cars become safer/easier to drive. According to the standard theory, increasing product safety results in a decrease in effort. If effort is a variable that is being used to determine whether or not to introduce regulation, then the regulator might decide against the regulation to avoid the Peltzman effect.

However, the experiment determines only that the primary adjustment behavior is negative; i.e, $\frac{\partial a^*}{\partial \alpha} < 0$. There still is the secondary adjustment behavior $\frac{\partial a^*}{\partial \beta}$, which has not been measured in the experiment. If the secondary adjustment behavior is also negative, then either the decrease in effort is smaller than measured in the experiment or, more importantly, the total change in effort is positive. Consequently, with this additional information, the regulator might opt for introducing the regulation.

When it comes to changes in effort, there is an important relationship between product regulation and consumer regulation. While the impact of an increase in product safety on effort is captured in equation (1), an increase in consumer skills affects effort in the following way: $\frac{da^*}{d\beta} = \frac{\partial a^*}{\partial \beta} + \frac{\partial a^*}{\partial \alpha}\frac{d\alpha}{d\beta}$. Suppose that $\frac{da^*}{d\alpha} > 0$; i.e., $\frac{\partial a^*}{\partial \alpha} + \frac{\partial a^*}{\partial \beta}\frac{d\beta}{d\alpha} > 0$. If we multiply both sides of this inequality by $\frac{d\alpha}{d\beta}$, then we observe that $\frac{\partial a^*}{\partial \beta} + \frac{\partial a^*}{\partial \alpha}\frac{d\alpha}{d\beta} < 0$; i.e., $\frac{da^*}{d\beta} < 0$. Consequently, in the presence of product-consumer substitution, effort increases due to product regulation if and only if effort decreases due to consumer regulation. This result has an important policy implication: if the regulator aims at not only lowering the probability of a bad outcome but also increasing effort, then the choice between product regulation and consumer regulation is an important task. This conclusion is not valid within the standard theory where there is no product-consumer substitution.

2.3 Safety regulation and probability of a bad outcome

We turn to the analysis of how product regulation changes the probability of a bad outcome when that probability is computed at optimal effort.

$$\frac{d\lambda^*}{d\alpha} = \underbrace{\underbrace{\lambda^*_\alpha}_{\substack{\textit{gain from}\\ \textit{regulation}}} + \underbrace{\lambda^*_a \frac{\partial a^*}{\partial \alpha}}_{\substack{\textit{primary}\\ \textit{impact of effort}}}}_{\textit{primary effect}} + \underbrace{\underbrace{\frac{d\beta}{d\alpha}\lambda^*_\beta}_{\substack{\textit{loss from}\\ \textit{regulation}}} + \underbrace{\frac{d\beta}{d\alpha}\lambda^*_a \frac{\partial a^*}{\partial \beta}}_{\substack{\textit{secondary}\\ \textit{impact of effort}}}}_{\textit{secondary effect}} \tag{2}$$

We say that safety regulation results in the **regular effect** if we observe a decrease in the probability of a bad outcome. Otherwise, regulation yields the **Peltzman effect**. In section 4.1, we test for the presence of Peltzman effect in our data in the context of both product and consumer regulation; this is a fundamental exercise in the safety regulation literature.

Hypothesis 3. *Product regulation yields the regular effect.*

Hypothesis 4. *Consumer regulation yields the regular effect.*

The change in the probability of a bad outcome due to product regulation depends on four elements. The sum of the first two elements in equation (2) can be called the **primary effect** or the standard effect of regulation; this is precisely the change in the probability of a bad outcome as predicted by the standard theory. The remaining two elements of equation (2) constitute the **secondary effect** of safety regulation and are driven by the product-consumer substitution; this is a new part.

First, λ^*_α is the **gain from regulation**. This is an unambiguous decrease in the probability of a bad outcome that assumes no change in behavior and characteristics of consumers. To measure the gain from regulation it is necessary to resort to laboratory tests where we can control for or eliminate the change in effort. An example of such a direct gain is a 51% reduction in road fatalities due to installed air bags and lap-shoulder belt (product regulation in the context of road safety).[12]

This 51% decrease is distorted by consumers (drivers) who modify their effort. In the context of road safety, effort captures how much attention people pay while driving. Upon learning that cars becomes safer (due to air bags and seat belts), drivers might feel too secure and decide to behave

[12]Exhibit 6 in “Fifth/Sixth Report to Congress, Effectiveness of Occupant Protection Systems and their Use,” National Highway Traffic Safety Administration U.S. Department of Transportation, Publication No. HS 809 442, 2001.

more recklessly (i.e., moral hazard). This primary adjustment behavior leads to the **primary impact of effort** $\lambda_a^* \frac{\partial a^*}{\partial \alpha}$. Since effort decreases the probability of a bad outcome (i.e., $\lambda_a^* < 0$), the sign of that impact depends on whether the primary adjustment behavior $\frac{\partial a^*}{\partial \alpha}$ is positive or negative. A priori, the sign of the primary impact of effort is unknown.

When cars become safer (due to already mentioned installed air bags and seat belts), then some of the people with low driving skills who used to avoid driving decide to become actual drivers. In consequence, an increase in product safety is associated with a decrease in consumer skills of the representative consumer. This decrease in skills negatively affects the safety (roads become more dangerous) and unambiguously increases the probability of a bad outcome. We call this effect the **loss from regulation** $\frac{d\beta}{d\alpha}\lambda_\beta^*$.[13] How big that loss is depends on the magnitude of the product-consumer substitution $\frac{d\beta}{d\alpha}$ and the importance of consumers skills in reduction of the probability of a bad outcome λ_β^*.

However, drivers know that the pool of consumers has changed because less qualified people have became drivers. Consequently, it is necessary to adapt to the new environment. For instance, as a precautionary measure, drivers may pay more attention as it is not wise to rely on other drivers to maintain a desired safety level. This secondary adjustment behavior leads to the **secondary impact of effort** $\frac{d\beta}{d\alpha}\lambda_a^* \frac{\partial a^*}{\partial \beta}$. Since the sign of the secondary adjustment behavior is unknown, it is impossible to determine whether the secondary impact of effort increases or decreases the probability of a bad outcome. However, since $\frac{d\beta}{d\alpha}\lambda_a^*$ is positive, we conclude that if effort increases (decreases) due to skills of representative consumer being lower, then the probability of a bad outcome decreases (increases).

Without additional assumptions or empirical tests, it is not possible to determine whether the secondary effect increases or decreases the probability of a bad outcome. Consequently, the standard theory not only incorrectly estimates the magnitude of the change in that probability[14] but also might mistakenly predict the Peltzman effect (regular effect) while the true outcome of regulation

[13] In the case of consumer regulation, an increase in consumer skills will result with a decrease in product safety. For instance, imposing restrictions on who can actively participate in financial markets increases financial literacy (higher consumer skills) but also increases the complexity of financial instruments (lower product safety). This is due to the negative correlation between financial literacy and market participation we already discussed in Introduction.

[14] For example, Cohen and Einav (2003) estimate that if 90% of drivers were to wear a seat belt, then about 1,500–3,000 lives would be saved on an annual basis. However, they also note that "although this estimate of the effect of increased seat belt usage on saved lives is substantial, it is considerably smaller than the estimate used by the federal government, which is 5,536 saved lives annually." Our theory suggests that this gap between pre-regulation expectation and post-regulation realization might be driven by the product-consumer substitution.

is the regular effect (Peltzman effect).

More importantly, the fact that the sign of secondary effect is uncertain implies that we might observe phenomena which directly contradict the standard theory. Recall that the primary effect is the only element of the standard theory. Whenever the primary effect is positive, we say that we observe the **standard offsetting behavior**; that is, $\lambda_\alpha^* + \lambda_a^* \frac{\partial a^*}{\partial \alpha} > 0$ in equation (2). Without the product-consumer substitution, observing the regular effect is equivalent with the lack of the standard offsetting behavior while the Peltzman effect is caused only by the standard offsetting behavior. However, each of these claims need not be true.

Hypothesis 5. *The regular effect does not rule out the standard offsetting behavior.*

Hypothesis 6. *The Peltzman effect does not entail the standard offsetting behavior.*

In order to explain the theoretical foundations of Hypothesis 5, suppose that product regulation results with the regular effect. According to equation (2), the following is possible: (a) the primary effect is actually positive (i.e., the standard offsetting behavior), and (b) the secondary impact of effort is negative and larger (in absolute terms) than the sum of the standard offsetting behavior and the loss from regulation. Consequently, we experience a decrease in the probability of a bad outcome. In section 4.4, we empirically confirm Hypothesis 5.

Hypothesis 6 assumes that we observe the Peltzman effect. As equation (2) indicates, it is possible that the primary effect is actually negative (i.e., no standard offsetting behavior), while it is the secondary effect that is positive. In this case, the Peltzman effect is driven only by the product-consumer substitution. In fact, it is also possible that the total effort increases—i.e., $\frac{da^*}{d\alpha} > 0$—but we still experience the Peltzman effect; this would be the case when the loss from regulation is large enough (in absolute terms). In section 4.5, we provide an empirical verification of Hypothesis 6.

Hypotheses 5 and 6 are important from the perspective of policy design since they indicate the dangers of disregarding the product-consumer substitution. If a regulator assumes zero correlation between product safety and consumer skills and observes the regular effect, then he dismisses the problem of the standard offsetting behavior. This could lead to generating wrong policies. Relying on the (incorrect) conclusion that the standard offsetting behavior is not an issue, the regulator might introduce additional regulations (increase α or β) and it is only then that the unexpected

Peltzman effect becomes a reality.

Not being concerned with the product-consumer substitution can also lead to inadequate policies trying to eliminate the Peltzman effect. According to the standard theory, the only remedy for the Peltzman effect is to correct the negative primary adjustment behavior. However, the regulatory effort spent on restraining the standard offsetting behavior might be futile. Rather than designing and implementing tools which curb the (non-existent) negative primary adjustment behavior, it is necessary to limit the decrease in consumer skills.

2.4 Two-dimensional regulation

When the product-consumer substitution is a serious concern, then, in order to mitigate a decrease in consumer skills, product regulation must be supported by consumer regulation. This requires implementing consumer regulation that introduces barriers to entry for low-skill consumers (e.g., driving test) or educates consumers in order to increase their skills. We call it **two-dimensional regulation** since, as opposed to traditional policies, the objective is to affect product safety and consumer skills simultaneously. In section 4.6, we empirically analyze an example of a two-dimensional policy in the iRacing data.

Our policy suggestion might seem counter-intuitive. Some would argue that an important reason for increasing product safety is to make the consumption safer, especially for consumers with low skills. However, as the product-consumer substitution indicates, higher product safety turns some non-consumers into consumers. Their skills are lower than the lowest consumer skills before regulation and this decrease in skills might not be compensated for by higher product safety.

Consider the case of financial regulation. Moore (2003), Campbell (2006), and Lusardi and Tufano (2015) show that people with lower financial literacy are more prone to commit mistakes in their investment decisions. As we already discussed, these people enter the market when financial instruments become less complex. Hence, if the regulator plans to ban certain instruments to make α higher, then, in order to avoid the Peltzman effect, it is necessary to prevent a too high decrease in β by prohibiting low-skill investors from entering the market.

3 Data

Our data was provided by iRacing (http://www.iracing.com/), an online racing simulator developed by iRacing.com Motorsport Simulations. We start with a description of the simulator in section 3.1.[15]

Since we are interested in analyzing human behavior, it is important that the members of iRacing behave as if they were participating in the real-life race. In section 3.2, we discuss how iRacing replicates the real-life racing in the virtual world.

In section 3.3, we interpret the data in light of the theory described in section 2 and discuss the variables used in empirical studies. Finally, in section 3.4, we explain the data-cleaning process.

3.1 iRacing: online racing simulator

iRacing is "the world's premier motorsports racing simulation. iRacing puts you in the driver's seat by allowing members to experience today's newest form of competitive motorsport: virtual racing" (http://www.iracing.com/overview/).

In order to use the simulator, it is necessary to become an iRacing member. The monthly membership fee is around $12, and the annual fee is around $110. There are no computer-simulated racers; all drivers are humans. Because of its competitive nature, iRacing is an example of an e-sport.

The simulator provides over 70 real racetracks from around the world and 50 cars ranging from easy-to-drive (Pontiac Solstice) to very challenging (Lotus 49) that have been re-created in iRacing. Once they log in, members of iRacing decide what race they want to join. The list of available races depends on the driver's skills which we will explain shortly. At any given moment, it is possible to be part of only one race. Races take place in real time.

Like in non-virtual racing, each race begins with a qualification that determines each racer's initial position. Prior to the race, drivers have the possibility to train on the track used in the race.

[15] If the reader is interested in learning more about iRacing, we suggest to watch live streaming of races at http://www.iracing.com/live/, or visit iRacing-dedicated YouTube channels like True Racer (https://www.youtube.com/user/TrueRacerAcademy/) or Empty Box (https://www.youtube.com/user/TacticalCardboard/).

As opposed to typical video games, the first-person point of view is the only available view for the racer in the simulator. As in a real race, each driver in iRacing has a view only from the seated position inside the car.

Before they choose a race, the drivers know the following about the race: the racetrack, type of track (oval or round), race conditions (night or day), the number of laps, and minimum requirements (in terms of driver's skills) to join the race. For about 90% of our data, for a given race, only one specific car is available; and the drivers know what car is being used in that race. In the remaining races, there are multiple cars available. However, these cars do not significantly differ in terms of their performance measure. Consequently, it is safe to say that in a given race everyone drives the same car.

3.2 iRacing: replicating real-life racing

There are three elements of the racing simulator that are important from our (research) perspective: tracks, cars, and drivers' behavior. We argue that the iRacing simulator replicates in the virtual world the real-world racing environment and behavior.

Behavior. Since we analyze human behavior, the most important question is whether iRacing members behave as if in a real-life race. After all, no matter how realistic the simulator is, there is always a risk that players drive recklessly on purpose or crash for fun. Fortunately, there are several strong incentives against such behavior which imply that the behavior of iRacing drivers mimics that of real-life drivers.

Drivers are divided into groups based on their license which captures the individual racing capabilities. New members start with the license Rookie and can be promoted depending on their performance. There are seven levels of license: Rookie, class D, class C, class, B, class A, class Pro, and class Pro World Class.

As already mentioned, each race is characterized by minimum requirements a racer must satisfy in order to join the race. This requirement is the the minimum license. If a race specifies that the minimum license is B, then nobody with license C, D, or Rookie is able to participate in the race.

The system of ranked licenses serves as an important incentive mechanism that prevents undesired racing behavior (e.g., crashing on purpose). Being promoted to a higher license is a driver's objective because such a promotion not only fulfills personal ambition and increases the driver's reputation but also, because of the minimum license requirement, allows access to more challenging cars and races against better-skilled drivers.

In addition, in the Pro and Pro World Class license, drivers have the opportunity to participate in NASCAR-sanctioned races with monetary prizes of up to $10,500 for the winner (see `http://www.iracing.com/nascar-iracing-com-series/` and `http://www.nascar.com/en_us/iracing.html`). Thus, reaching these levels of license has strong economic incentives.

The next mechanism that strongly incentives the iRacing members to behave as racers do in physical races is the system of incident points. During each race, drivers accumulate incident points for their involvement in on-track incidents. In Table 1, we present how many incident points a driver receives for each possible racing incident.

[Table 1 about here.]

In each race, there is a maximum number of incident points that a driver can accumulate. Reaching that limit results in an immediate race disqualification.

The system of incident points is related to the system of licenses. In order to be promoted to a higher level of license, a racer has to not only win races but also maintain a certain level of safe racing. It is possible that someone wins many races but does not receive a promotion because of reckless behavior. In fact, collecting too many cumulative incident points results in a demotion to a lower license.

In order to be able to use iRacing, "the member will require a controller to enjoy the full range of experiences afforded by iRacing's racing simulator. A host of steering wheel/pedal combos, gamepads, joysticks, mouse-based control systems, and any version of the Microsoft Windows operating system, supporting touch screen driving are compatible with iRacing.com" (`http://www.iracing.com/membership/system-requirements/`). It is not possible to drive in iRacing using a keyboard. This significantly differentiates iRacing from typical video games. According to the company that owns and manages iRacing, more than 95% of racers use a wheel and pedal set. The authors of this paper tried iRacing using a mouse but failed miserably (which is obviously due to lack of not only proper equipment but also skills). Since members of iRacing are required to purchase additional equipment, we believe that this works as a pre-selection mechanism that screens out those whose main objective is to crash on purpose.

Track and cars. When it comes to the tracks and cars, the objective is to have them designed as virtual replicas of the real-life tracks and cars. That is, the virtual cars are to behave on virtual tracks in the same way their real-life counterparts behave on physical tracks. iRacing has been successful in realistically replicating cars and tracks so that professional racers use iRacing for training purposes.[16]

In order to design the tracks, the company relies on "its pioneering, proprietary application of three-dimensional laser-scanning technology to create two key features: highly detailed sight-pictures and precise physical features of each track's racing surface."[17] According to Dale Earnhardt, Jr., a famous NASCAR driver, "every inch of every track [in iRacing] is modeled perfectly."

When it comes to the design of virtual cars, the most important technological aspect is a physics engine, which is "a complex system of high-speed mathematical functions that replicate and deliver dynamic forces using data-driven calculations, thereby leading to a series of instantaneous dynamic actions and reactions. Consequently, iRacing's virtual world leverages the same physical dynamics and that drivers experience in the real world."[18] Virtual cars have the same specifications as the real cars which is important from our perspective since we use the weight and horse power ratio as a proxy measure for a car's difficulty/safety.

To sum, iRacing is a simulated environment that not only replicates the physical attributes of the real-life racing such as tracks and cars but also provides incentives to the members so their behavior simulates the behavior of real-life racers.

3.3 iRacing data in the light of the theory: our variables

Our data consists of individual observations. For example, if N drivers participated in a given race, then we obtain N observations. After the cleaning procedure discussed in section 3.4, the data used in our study consists of 2,274,192 observations and includes 41,010 different individuals

[16]See testimonials available at `http://www.iracing.com/testimonials/`.

[17]For more information, see `http://www.iracing.com/cars-and-tracks/track-technology/` and a video available at `http://www.iracing.com/track-technology/` that explains the design of virtual tracks.

[18]For more information, see `http://www.iracing.com/car-technology/`. On that website, there also is a video explaining in detail the design of virtual cars.

from 106 countries who, between January and December 2015, participated in 150,598 races. Each observation consists of the seven variables we discuss below.

The main three variables that we use are Incidents Point per Mile, Weight to Horse Power Ratio, and License. We also add four track- and race-related variables that serve as our controls (Traffic Density, Laps in the Race, Oval, and Night).

Incident points capture how unsafe a driver's performance was during a race; more incident points implies less safety. Since races differ in the length of the track and racers do not always complete the whole race, we define **Incidents Point per Mile** as the number of incident points per mile driven. This is our safety measure λ. Higher Incidents Point per Mile means higher λ.

There are two car-related variables that are of our interest: weight and horse power. In the racing world, a car's **Weight to Horse Power Ratio** is one of the fundamental metrics in assessing how controllable a car is. The lower this ratio is, the faster the car accelerates and the higher its final speed is. In our paper, the Weight to Horse Power Ratio, denoted also by WHP, serves as a proxy measure of product safety α. Higher WHP means higher α.

The variable **License**, which we already discussed above, captures the driving capabilities of iRacing members. This is our proxy measure of consumer skills β. Higher License means higher β. To each class of License, we assign a numerical value in the increasing order: Rookie (1), class D (2), class C (3), class B (4), class A (5), class Pro (6), and class Pro World Class (7).

Traffic Density is the average number of drivers per lap in a race. During a race, some drivers might drop out of the race or be disqualified. Imagine that a race consists of 50 laps, $LR = 50$, with 20 drivers starting the race, $N = 20$. Then, assume that driver i completed 40 of the 50 laps in the race, $LC_i = 40$. We define Traffic Density as the sum of all completed laps by all drivers in a race divided by the number of laps in that race $\sum_{i=1}^{N} \frac{LC_i}{LR}$. This measure takes into account the fact that two races that started with the same number of drivers and had the same number of total laps do not represent the same traffic density if the quantity of existing drivers per lap during the race differs.

Laps in the Race is the number of laps in a given race, **Oval** is a binary variable that is one if the track is oval and zero otherwise, and **Night** is also a binary variable that is one if the race conditions are night conditions and zero otherwise.

Table 2 presents summary statistics.

[Table 2 about here.]

Figures 1 and 2 depict the distribution of average Incidents per Mile with respect to Weight to Horse Power Ratio (product safety) and License (consumer skills).

[Figure 1 about here.]

[Figure 2 about here.]

3.4 Cleaned data

First, we eliminated observations generated by races in which there was at least one Rookie driver. Hence, if a race includes twenty drivers and one of them is a Rookie, then we remove all twenty observations from our data. This is motivated by the fact that the Rookie racers could just be testing the simulator (as we did) and the incentives we describe in section 3.2 might not be affecting them. Consequently, the presence of Rookie drivers might turn an iRacing race into a video game rather than a racing simulation.

Second, we eliminated all unofficial races. Races are divided into official and unofficial. Unofficial races have no impact on an iRacing member's career and need not motivate the drivers to behave as if they would in a physical race. Hence, we focus only on official races, which replicate real-life racing behavior.

Finally, we keep only those races in which the maximum number of incident points allowed in a race for a driver before being disqualified equals seventeen. About 94% of the official races allow for exactly this maximum.

4 Empirical Analysis

4.1 Hypotheses 3 and 4: regular effect vs. Peltzman effect

Hypotheses 3 and 4 are about the traditional exercise in the safety regulation literature: evaluating how regulation (product and consumer) affects the safety measure.[19] The impact of regulation (product or consumer) on the safety measure can be evaluated within a simple regression framework. In Model 1 (Table 3), we analyze whether product regulation and consumer regulation result in the Peltzman or regular effect. The estimation equation is

$$Y = a_0 + a_1 \times WHP + a_2 \times License + Xb + \varepsilon, \tag{3}$$

where Y is the number of Incident Points per Mile (safety measure), WHP is Weight to Horse Power Ratio, License is the level of License, and X is a vector of four control variables (Traffic Density, Laps in the Race, Oval, and Night). Our focus is the sign of a_1 and a_2.

[Table 3 about here.]

In the case of product regulation, a_1 measures the impact of product regulation on the safety measure. If the sign of a_1 is positive (negative), then we detect the Peltzman (regular) effect. The control variables are License and X.

In the case of consumer regulation, a_2 measures the impact of product regulation on the safety measure. If the sign of a_2 is positive (negative), then we detect the Peltzman (regular) effect. The control variables are Weight to Horse Power Ratio and X.

We find that the coefficient of Weight to Horse Power Ratio is negative (-0.0002) and significant, and the coefficient of License is also negative (-0.0671) and significant. This means that we detect the regular effect for both product regulation (Hypothesis 3) and consumer regulation (Hypothesis 4).

[19] We start with Hypotheses 3 and 4 rather than Hypothesis 1 because, as it will become evident later, Hypothesis 1 is a discovery due to the detailed analysis of the regular/Peltzman effect.

4.2 Hypothesis 1: product-consumer substitution

Hypothesis 1 states that product safety and consumer skills are negatively correlated.

In Model 1 (Table 3), we are able to control for all relevant variables; when we consider product regulation we control for consumer skills (License) while when we analyze consumer regulation we control for product safety (Weight to Horse Power Ratio). However, in the literature, it is rather common that we are unable to control for either product safety or consumer skills. Thus, it is only natural to ask whether omitted-variable bias is a significant problem. We answer this question in the next exercise and uncover a new and important phenomenon that is a fundamental element of our theory; namely, product-consumer substitution. We use the same data and empirical strategy as in the previous section 4.1.

In Model 2, we analyze product regulation but do not control for consumer skills; i.e., comparing to Model 1, the only change we make is to omit the variable License. In Model 3, we evaluate consumer regulation without controlling for product safety; i.e., comparing to Model 1, variable Weight to Horse Power Ratio is omitted. Both models are presented in Table 3.

In Model 2, we detect the Peltzman effect as the coefficient of Weight to Horse Power Ratio is positive (0.0068) and significant. This means that when cars become safer/less difficult to drive, then we expect drivers to accumulate more incident points per mile. Recall that in Model 1, where we add License, we observe the regular effect. Hence, the Peltzman effect observed in Model 2 disappears as long as we control for consumer skills. The fact that the magnitude of the bias changes not only quantitative but also qualitative results serves as a warning that the omitted-variable bias is a significant problem for the empirical literature on safety regulation.

The omitted-variable bias is less important when it comes to consumer regulation. In Model 3, we omit product safety (Weight to Horse Power Ratio) and find that consumer regulation results in the regular effect; that is, the coefficient of License is negative (-0.0669) and significant. In Model 1, where we add Weight to Horse Power Ratio, we also detect the regular effect.

When comparing Model 2 to Model 1, we note that the coefficient of Weight to Horse Power Ratio increases. Comparing Model 3 to Model 1 shows that the coefficient of License increases as well. In both cases, we observe an upward bias. This bias is due to the negative correlation between Weight to Horse Power Ratio and License which is -0.43 (p-value < 0.001) and the fact that the coefficients of License and Weight to Horse Power Ratio are both negative. In other words, we detect the product-consumer substitution.

4.3 Hypothesis 2: forces behind the product-consumer substitution

Hypothesis 2 states that there are two forces, demand and supply, that drive the product-consumer substitution.

4.3.1 Supply-driven product-consumer substitution (design of iRacing)

In the case of iRacing, the "supply" side is the design of the simulator which determines the car difficulty, as measured by Weight to Horse Power Ratio, available for each level of License. Recall that in order to participate in a given race, the members of iRacing have to meet specific minimum License requirements. The lowest level of License that we consider in our empirical analysis is License D. Drivers with License D choose cars with Weight to Horse Power Ratio ranging from 5.08 to 15.90.

When the racers are promoted to License C, then they have access to the same cars as before; however, there are new cars added as well. These new cars range, in terms of Weight to Horse Power Ratio, from 2.24 to 15.29. That is, the cars with Weight to Horse Power Ratio in the range $[2.24, 5.08)$ are not available to drivers License D but accessible for drivers with License C.

When promoted from License C to License B, then again the new cars become available. The Weight to Horse Power Ratio of these new cars is from 1.80 to 5.46. Hence, by design, comparing to the drivers with License B, the drivers with License C are not permitted to drive cars which are not safe enough; in particular, these are the cars with Weight to Horse Power Ratio in the range $[1.80, 2.24)$.

Finally, when promoted from License B to License A, the Weight to Horse Power Ratio of new cars range is $[1.77, 3.77]$; again, less safe cars are added. In short, due to the design of iRacing, the relationship between Weight to Horse Power Ratio and License is negative: in order to be granted an access to less safe products, consumer must have skills high enough.

4.3.2 Demand-driven product-consumer substitution (adverse selection)

In iRacing, we also observe the demand-driven negative correlation between product safety and consumer skills; that is, racers with lower skills choose safer cars. In order to detect the presence of adverse selection in the context of iRacing, we need to show that given the same set of cars to choose from, high-skill drivers choose cars with lower Weight to Horse Power Ratio compared to the cars chosen by the low-skill drivers.

Obviously, we can not compare choices made by the drivers with different levels of License because, as explained above, by design the available choices of cars depend on the level of License. Hence, even though the average Weight to Horse Power Ratio is 11.13 for License D and 8.75 for License C, it would be inaccurate to state that the difference between these numbers is driven only by drivers' choice. Consequently, to overcome the problem of how iRacing has been designed, we rely on the following empirical strategy.

We know that, in iRacing, drivers with the same level of License have access to the same cars. Hence, our objective is to analyze the drivers with the same level of License but different level of skills. Our analysis consists of two parts. In part I, for a cohort of drivers with the same License J, we identify high-skill drivers and low-skill drivers. In part II, we compare the values of Weight to Horse Power Ratio chosen by these drivers. If high-skill drivers choose lower Weight to Horse Power Ratio, then this proves that, in the context of iRacing, consumers with higher skills choose less safe products, which confirms the presence of adverse selection.

Part I: Identifying high-skill and low-skill drivers within the same License. To differentiate the skills of drivers within the same License level, we exploit the fact that drivers can be demoted from License $J+1$ to License J, or promoted from License $J-1$ to License J. As we explain shortly, for a given License J, we say that the drivers who have been demoted from License $J+1$ are high-skill drivers while the drivers who have been promoted from License $J-1$ are low-skill drivers.

For a driver promoted or demoted to License J, we pick the first 5 races[20] in which the driver competed in License J. It might happen that after promotion/demotion, a driver gets promoted/demoted again in fewer than k races, where $k < 5$. In this case, we keep these k races

[20] All the results we present are robust to keeping any fixed number of races after promotion/demotion as well as keeping all the races until the following change of a driver's License level, if any.

of the driver in our sample.

A driver might be demoted more than once. For example, Ann might be demoted from License $J+2$ to License $J+1$ to License J and, finally, to License $J-1$. If Ann then gets promoted from License $J-1$ to J, then, under our identifying assumption, she would be called a low-skill driver while Bob, who was demoted from $J+1$ to J but has never been promoted to $J+2$, would be a high-skill driver. Clearly, this is not appropriate. Consequently, we restrict our attention to the drivers who have been demoted at most twice.

We argue that the drivers who were demoted from License $J+1$ to License J are high-skill drivers and the drivers promoted from License $J-1$ to License J are low-skill drivers. To support our argument, first, observe that the demoted drivers "return" to the environment they already know. For the promoted drivers, there are cars in License J which they have not used before. Comparing to the drivers promoted to License J, the drivers demoted to License J not only have more experience with cars available in License J but also proved to have more driving capabilities as they have already been promoted to a higher License $J+1$.

We now provide empirical proof that, for a given License J, demoted drivers from License $J+1$ are more skilled than promoted drivers from License $J-1$. In iRacing, to demonstrate high skills a driver needs to perform well in two dimensions: (i) high final position in a race, and (ii) low number of Incident Points. Therefore, we create three race-specific performance-related variables to measure a driver's performance within a race.

1. Relative Incident Points (Y_1): it is the Incidents per Mile of a driver in a race divided by the Average Incidents per Mile of all drivers in the same race.

2. Relative Final Position (Y_2): it is the final position of a driver in a race divided by the number of drivers in the race.

3. Quartile Position (Y_3): it is the quartile position of a driver in a race. (We remove races with less than 4 drivers.)

Recall that in our data we keep only observations generated by promoted and demoted drivers. However, a driver's final outcome in a race will depend also on his skills relative to the skills of other drivers who need not be in our sample, as they might have never been promoted or demoted within the time span of our data, or they might have been demoted more than twice. To control the outcome of a driver in a race relatively to the average skills of the drivers in that same race, we compute the difference between his/her License level and the Average License level of all drivers in the same race.

Once we control for the average skills, we check whether drivers demoted from $J+1$ to J are of higher skills than those promoted from $J-1$ to J. To that end, we create a dummy variable High Skills which is equal to 1 if the driver has been demoted and 0 if promoted. Therefore, our two independent variables are the following.

1. License Difference: it is the difference between the demoted/promoted driver's License and the average driver's License in a race. Since License is a measure of driving skills, this variable controls for the fact that a driver with higher than average License is expected to perform better than average.
2. High Skills: it is a dummy variable equal to 1 for players demoted from License $J+1$ to J and equal to 0 for players promoted from License $J-1$ to J.

We consider three estimation equations

$$Y_i = a_0 + a_1 \times \textit{High Skills} + a_2 \times \textit{License Difference} + \varepsilon, \tag{4}$$

where Y_i (race-specific performance-related dependent variable), License Difference, and High Skills are described above. Our focus is the sign of a_1.

[Table 4 about here.]

Table 4 depicts the results and clearly shows that those we defined as high-skill drivers perform better (in terms of final position and safety) than the group of low-skill drivers. This is because, for each value of License (D, C, and B), and for each race-specific performance-related variable Y_i, the coefficient of High Skills is negative and significant.

Part 2: Comparing selection of cars by high-skill and low-skill drivers within the same License. In the first part of our analysis, we established that, for a given level of License J, the drivers demoted from $J+1$ to J are of higher skills compared to the drivers promoted from $J-1$

to J. Finally, we show that the drivers with higher skills choose cars with lower Weight to Horse Power Ratio. To that end, we perform an observational study by comparing means and related statistics of Weight to Horse Power Ratio of the cars selected by high-skill and low-skill drivers. The results are presented in Table 5.

[Table 5 about here.]

For each License J, where J stands for D, C, and B, we observe that the 95% intervals of Weight to Horse Power Ratio (of cars chosen by the drivers) are disjoint: the lower bound for low-skill drivers is higher than the upper bound for high-skill drivers. In fact, the mean Weight to Horse Power Ratio for low-skill drivers is higher than the mean Weight to Horse Power Ratio for high-skill drivers; this result is significant at the 1% level. In short, our results indicate the presence of adverse selection.

4.4 Hypothesis 5: regular effect with standard offsetting behavior

Hypothesis 5 stipulates that it is possible to observe both the regular effect and the standard offsetting behavior. We find an example of such phenomena in our data. Our empirical strategy consists of several stages.

First, we truncate the data and consider only the observations with Weight to Horse Power Ratio greater than 9.5. This sub-sample of the whole data set consists of 493,454 observations. We denote this sub-sample by $\mathcal{R}$ (as in $\mathcal{R}$egular effect). Next, we establish the presence of the regular effect within $\mathcal{R}$; that is, we estimate (3) using only the sub-sample $\mathcal{R}$. The results are presented in column "All Licenses" (Table 6). Since the coefficient of Weight to Horse Power Ratio is negative (-0.0013) and significant, we observe the regular effect.

[Table 6 about here.]

Once we detect the regular effect, we test for the presence of the standard offsetting behavior. Our strategy is to analyze the impact of product regulation on the probability of a bad outcome under the assumption that consumer skills remain unchanged. This approach is driven by equation (2): when β (consumer skills) is fixed, then the change in the probability of bad outcome becomes $\lambda^*_\alpha + \lambda^*_a \frac{\partial a^*}{\partial \alpha}$. If that change is negative then we detect the standard offsetting behavior.

We partition $\mathcal{R}$ into subsets with identical License value. Let $\mathcal{R}_J$ denote a subset of $\mathcal{R}$ where the value of the variable License is J; for such a subset, the value of β (consumer skills) is fixed. Out of 493,454 observations, only 171 (i.e., 0.3% of the whole data) are with License Pro or Pro World Class. Consequently, these two level of License have negligible impact on our results. Hence, $J =$ D, C, B, A. For each $\mathcal{R}_J$, our estimation equation is

$$Y = a_0 + a_1 \times WHP + Xb + \varepsilon, \tag{5}$$

where Y is the number of Incident Points per Mile (safety measure), WHP is Weight to Horse Power Ratio, and X is a vector of four control variables (Traffic Density, Laps in the Race, Oval, and Night). Comparing to regression (3), we exclude the variable License as this variable is constant for a given $\mathcal{R}_J$.

We look at the sign of a_1, the coefficient of Weight to Horse Power Ratio. If that coefficient is negative and significant, then we reject the hypothesis of the standard offsetting behavior. If that coefficient is positive and significant, then we detect the standard offsetting behavior.

We are unable to reject the hypothesis of the standard offsetting behavior for levels of License B and A: in both cases, the coefficient of Weight to Horse Power Ratio is positive, 0.0019 and 0.0051, respectively and significant (see columns "License B" and "License A").[21] This confirms Hypothesis 5.

4.5 Hypothesis 6: Peltzman effect without standard offsetting behavior

Hypothesis 6 postulates the possibility of observing the Peltzman effect without the presence of the standard offsetting behavior. In order to empirically confirm this possibility, we use the same strategy employed to confirm Hypothesis 5; of course, the objective is to obtain different qualitative results.

Now, we consider only the observations with Weight to Horse Power Ratio smaller than 7. This

[21] It would seem that when the car difficulty is relatively low (i.e., product safety is high), then drivers with high racing skills behave more recklessly when the car difficulty decreases (i.e., when product regulation is implemented). These consumers might feel overly confident—after all, the product is safe and their skills are high—which makes them exert less effort while racing.

sub-sample is denoted by $\mathcal{P}$ (as in the $\mathcal{P}$eltzman effect) and consists of 1,362,597 observations. We test for the presence of the Peltzman effect by estimating (3) using $\mathcal{P}$. The results presented in column "All Licenses" (Table 7) show that the coefficient of Weight to Horse Power Ratio is positive (0.0028) and significant. Hence, we detect the Peltzman effect.

[Table 7 about here.]

Next, we test for the presence of the standard offsetting behavior. We consider four values of License, $J =$ D, C, B, A. This is because out of 1,362,597 observations, only 4,143 (i.e., 0.3% of the whole data) are with License Pro or Pro World Class. A subset of $\mathcal{P}$ consisting of all observations with the value of the variable License equal to J is denoted by $\mathcal{P}_J$. We estimate (5) for each $\mathcal{P}_J$ ($J =$ D, C, B, A).

The coefficient of Weight to Horse Power Ratio is positive (0.0869) only for the sub-sample $\mathcal{P}_D$ (column "License D"). That is, we detect the standard offsetting behavior only for the lowest level of License. Note that, in our sub-sample $\mathcal{P}$, there are only 62,779 observations with the value of License being D. In other words, we detect the standard offsetting behavior in only 4.6% of our data. For 95% of observations, we do not detect the standard offsetting behavior. This confirms Hypothesis 6.

4.6 Two-dimensional regulation: empirical example

In our last empirical exercise, we implement the two-dimensional policy recommendation from section 2.4. Our starting point is the Peltzman effect detected in section 4.5 (see Model "All Licenses" in Table 7). Our goal is to convert that Peltzman effect into the regular effect. As already established in section 4.5, it is not the standard offsetting behavior but rather the product-consumer substitution that we need to deal with.

We suggest to complement product regulation with consumer regulation that introduces new restrictions on who can participate in a race. In particular, we remove all drivers with License less than C. To see the impact of this policy, we estimate (3) using the sub-sample $\mathcal{P}$ (see section 4.5) and add one restriction: we consider only observations with the value of License being at least C. Table 8 presents the results of our exercise.

[Table 8 about here.]

In column "License at least C," we observe that the coefficient of Weight to Horse Power Ratio becomes negative (-0.0006) which means that product regulation supported by consumer regulation results in the regular effect. If we limit access to racing to those with the License of at least B, then our policy becomes even more efficient: the coefficient of Licence becomes -0.001 and is significant at 1% (column "License at least B").

5 Conclusions

The standard theory of safety regulation focuses on the role of moral hazard and its negative effect. We propose an alternative theory whose main innovation is the product-consumer substitution; i.e., the negative correlation between product safety and consumer skills.

We support our theory with an empirical analysis using more than 2 million observations from iRacing, an online racing simulator. In the virtual world, iRacing replicates the physical racing. Most importantly, the incentives mechanism induces iRacing drivers to behave as if they were participating in real-world races. We consider iRacing to be an ideal laboratory since it provides objective measures of consumer skill and product safety.

During each race, the drivers accumulate incident points for being involved in racing incidents. Lower safety means more incident points. We use the variable Incidents Points per Mile as a safety measure. This is our dependent variable.

The variable Weight to Horse Power Ratio is a measure of product safety. This ratio is a proxy measure for a car's controllability: the lower the ratio is the faster the car accelerates and the higher its probable maximum speed. The variable License is a measure of consumer skills. Drivers are divided into seven groups based on their racing capabilities. To achieve a higher level of License, it is necessary to both show driving proficiency by finishing races in high positions and maintain a low level of incidents. Weight to Horse Power Ratio and License are our main independent variables. We also use four additional controls (Traffic Density, Laps in the Race, Oval, and Night).

In our benchmark regression we observe the regular effect for both product and consumer regulation (Hypotheses 3 and 4). Next, we analyze the importance of the omitted-variable bias and conclude that it is a serious problem for the empirical literature on safety regulation. This leads us to detect the presence of the product-consumer substitution (Hypothesis 1) which, consequently, necessitates the development of a new theory.

We empirically confirm that the product-consumer substitution is driven by demand and supply forces (Hypothesis 2). In the case of iRacing, high skill drivers choose riskier cars (adverse selection) and, by design, more dangerous cars are accessible only to drivers with high skills.

Neglecting product-consumer substitution creates incomplete or even incorrect depiction of safety regulation. In particular, assuming that the product-consumer substitution is zero results in incorrectly estimated changes in consumer effort and probability of a bad outcome. In fact, it is possible that while the standard theory predicts an increase (decrease) in effort or the probability of a bad outcome, we actually observe a decrease (increase).

Under the presence of product-consumer substitution, product regulation and consumer regulation have opposite impacts on effort. This result is important if consumer effort is an objective of policy; for example, the regulator cares about not only the safety measure but also the response of the consumers.

Product-consumer substitution implies that we might observe the regular effect together with the standard offsetting behavior (Hypothesis 5) and the Peltzman effect without the standard offsetting behavior (Hypothesis 6). Each of these hypotheses could not happen according to the standard theory. In order to test for the standard offsetting behavior it is necessary to keep the consumer skills unchanged. In our case, we achieve that by considering observations with the same License value.

We confirm Hypothesis 5 by constructing a sub-sample in which we detect the regular effect as well as the standard offsetting behavior for two groups of consumers (Licenses B and A). This example shows that we should not disregard the problem of the standard offsetting behavior even if the Peltzman effect is not detected.

To confirm Hypothesis 6, we construct another sub-sample with the presence of the Peltzman effect. However, we observe the standard offsetting behavior in only 4.6% of our data. This result stresses the importance of product-consumer substitution.

Finally, we propose a two-dimensional policy as a solution to the problem of product-consumer substitution. In the case of product regulation (consumer regulation), the objective is to not only increase product safety (consumer skills) but also control the decline in consumer skills (product safety) by simultaneously imposing consumer regulation (product regulation). We provide an empirical analysis of this combined product-consumer regulation.

References

Allen, F., E. Carletti, I. Goldstein, and A. Leonello (2015): "Moral Hazard and Government Guarantees in the Banking Industry," *Journal of Financial Regulation*, 1, 30–50.

Amromin, G., J. C. Huang, C. Sialm, and E. Zhong (2013): "Complex Mortgages," *AFA 2012 Chicago Meetings Paper.*

Bagehot, W. (1873): *Lombard street: A description of the Money Market*, London: Henry S. King & Co.

Becker, G. S. (1968): "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 76, 169–217.

Campbell, J. Y. (2006): "Household Finance," *Journal of Finance*, 61, 1553–1604.

Chong, A. and P. Restrepo (2014): "Regulatory Protective Measures and Risky Behavior: Should We Be Saved From Ourselves?" *working paper, http://economics.mit.edu/files/10791.*

Cohen, A. and L. Einav (2003): "The Effects of Mandatory Seat Belt Laws on Driving Behavior and Traffic Fatalities," *Review of Economics and Statistics*, 85, 828–843.

Dam, L. and M. Koetter (2012): "Bank Bailouts and Moral Hazard: Evidence from Germany," *Review of Financial Studies*, 25, 2343–2380.

Farhi, E., , and J. Tirole (2012): "Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts," *American Economic Review*, 102, 60–93.

Gorton, G. and L. Huang (2004): "Liquidity, Efficiency, and Bank Bailouts," *American Economic Review*, 94, 455–483.

Grossman, R. S. (1992): "Deposit Insurance, Regulation, and Moral Hazard in the Thrift Industry: Evidence from the 1930s," *American Economic Review*, 82, 800–821.

Guha, B. and A. S. Guha (2012): "Crime and Moral Hazard: Does More Policing Necessarily Induce Private Negligence?" *Economics Letters*, 115, 455–459.

Ker, K., I. Roberts, T. Collier, F. Beyer, F. Bunn, and C. Frost (2005): "Post-licence Driver Education for the Prevention of Road Traffic Crashes: A Systematic Review of Randomised Controlled Trials," *Accident Analysis and Prevention*, 37, 305–313.

Lusardi, A. and O. S. Mitchell (2014): "The Economic Importance of Financial Literacy: Theory and Evidence," *Journal of Economic Literature*, 52, 5–44.

Lusardi, A. and P. Tufano (2015): "Debt Literacy, Financial Experiences, and Overindebtedness," *Journal of Pension Economics and Finance*, 14, 332–368.

McCarthy, P. S. and W. K. Talley (1999): "Evidence on Risk Compensation and Safety Behaviour," *Economics Letters*, 62, 91–96.

Moore, D. (2003): "Survey of Financial Literacy in Washington State: Knowledge, Behavior, Attitudes, and Experiences," *Washington State University Social and Economic Sciences Research Center Technical Report*, 03-39.

Peltzman, S. (1975): "The Effects of Automobile Safety Regulation," *Journal of Political Economy*, 83, 677–726.

Pope, A. T. and R. D. Tollison (2010): ""Rubbin' is Racin'": Evidence of the Peltzman Effect from NASCAR," *Public Choice*, 142, 507–513.

Sobel, R. S. and T. M. Nesbit (2007): "Automobile Safety Regulation and the Incentive to Drive Recklessly: Evidence from NASCAR," *Southern Economic Journal*, 74, 71–84.

van Rooij, M., A. Lusardi, and R. Alessie (2011): "Financial Literacy and Stock Market Participation," *Journal of Financial Economics*, 101, 449–472.

Viscusi, W. K. (1996): *Fatal Tradeoffs: Public and Private Responsibilities for Risk*, Oxford University Press.

——— (2007): "Regulation of Health, Safety, and Environmental Risks," in *Handbook of Law and Economics, Volume 1*, ed. by A. M. Polinsky and S. Shavell, Palgrave Macmillan, 592–645.

Viscusi, W. K., J. M. Vernon, and J. E. Harrington Jr. (2005): *Economics of Regulation and Antitrust*, The MIT Press.

Appendix A Primary adjustment behavior and secondary adjustment behavior

In the context of product regulation, we discuss the relationship between the primary adjustment behavior $\frac{\partial a^*}{\partial \alpha}$ and the secondary adjustment behavior $\frac{\partial a^*}{\partial \beta}$. First, we derive both adjustment behaviors.

$$\frac{\partial a^*}{\partial \alpha} = -\frac{\lambda^*_{a\alpha}}{\lambda^*_{aa} + c''(a^*)} \tag{6}$$

$$\frac{\partial a^*}{\partial \beta} = -\frac{\lambda^*_{a\beta}}{\lambda^*_{aa} + c''(a^*)} \tag{7}$$

Since $\lambda_{aa} + c'' > 0$, the sign of each derivative depends on their numerators. Note that $\lambda_{a\alpha}$ and $\lambda_{a\beta}$ measure the marginal impact of product safety and consumer skills, respectively, on the marginal gain from effort λ_a. One could either estimate the signs of $\lambda_{a\alpha}$ and $\lambda_{a\beta}$, or impose the additional assumptions.

However, regardless of what we assume about $\lambda_{a\alpha}$ (positive or negative), it is only appropriate to assume the same about $\lambda_{a\beta}$. That is, product safety and consumer skill have the same qualitative impact on the marginal gain from effort. This implies that $\frac{\partial a^*}{\partial \alpha}$ and $\frac{\partial a^*}{\partial \beta}$ have the same sign.

Recall from equation (1) that the change in effort is $\frac{da^*}{d\alpha} = \frac{\partial a^*}{\partial \alpha} + \frac{\partial a^*}{\partial \beta}\frac{d\beta}{d\alpha}$. Since $\frac{d\beta}{d\alpha}$ is negative, we conclude that primary and secondary adjustment behaviors work in opposite ways. That is, primary adjustment behavior decreases effort if and only if secondary adjustment behavior increases effort.

Appendix B Optimal regulation

We discuss the problem of optimality in the context of safety regulation. Again, we focus on product regulation; that is, the regulator increases α.

Let $U^* = 1 - \lambda(\alpha, \beta, a^*) - c(a^*)$ denote the utility of representative consumer computed for the optimal level of effort. For simplicity, we assume that policy design, implementation, and enforcement are costless.

First, as in the standard theory, we assume no product-consumer substitution; that is, β is not affected by changes in α.

$$\frac{dU^*}{d\alpha} = -\lambda_\alpha(\alpha, \beta, a^*) \tag{8}$$

Since $\lambda_\alpha < 0$, we observe that U^* is an increasing function in α no matter what the values of other parameters are; more regulation is always desired.

With the product-consumer substitution, we have $\frac{d\beta}{d\alpha} < 0$. Hence, we obtain the following.

$$\frac{dU^*}{d\alpha} = -\lambda_\alpha(\alpha, \beta, a^*) - \lambda_\beta(\alpha, \beta, a^*)\frac{d\beta}{d\alpha} \tag{9}$$

Since $\lambda_\alpha < 0$ and $\lambda_\beta \frac{d\beta}{d\alpha} > 0$, the sign of $\frac{dU^*}{d\alpha}$ is, a priori, unknown. Consequently, whether or not regulation is desirable depends on both gain and loss from regulation.

Figure 1: Average Incident Points per Mile for each value of WHP (Weight to Horse Power Ratio)

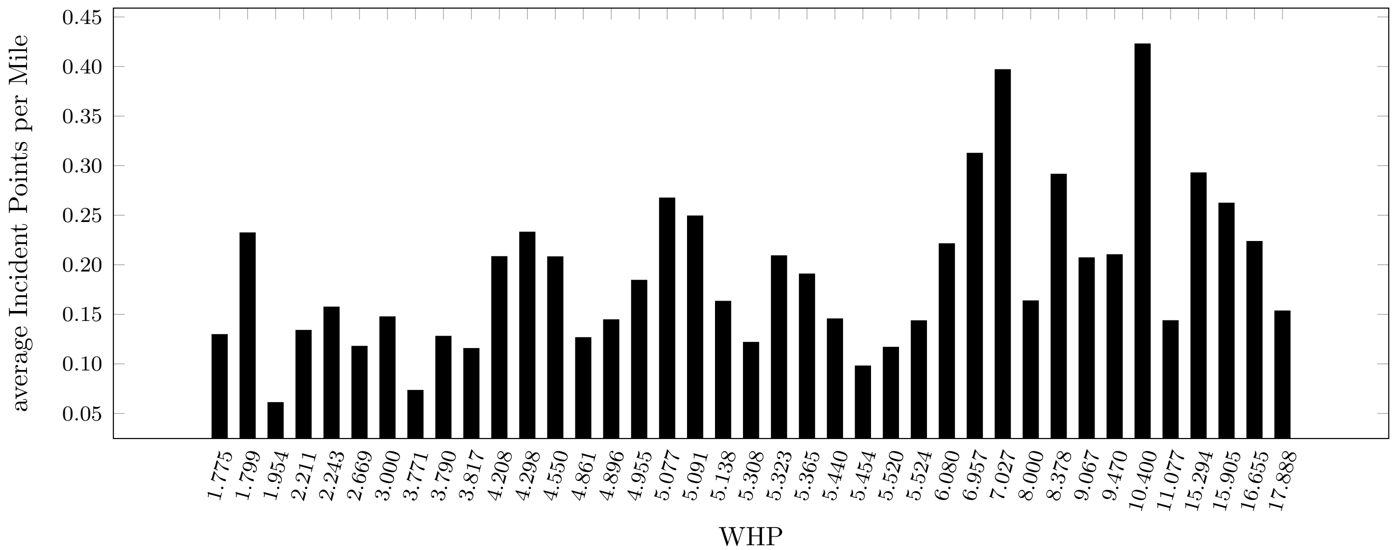

Figure 2: Average Incident Points per Mile for each value of License

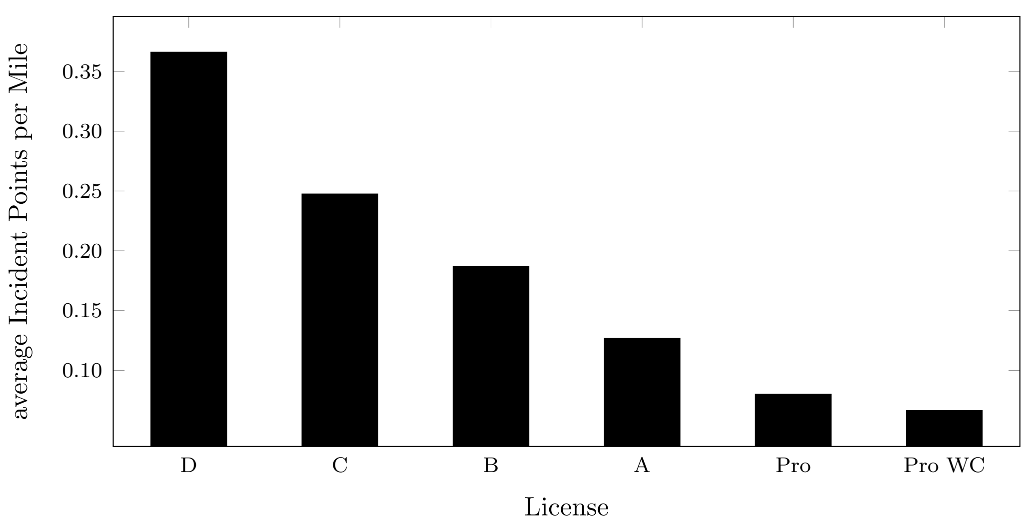

Table 1: Incidents and Incident Points

This table presents the number of incidents points per each possible racing incident. Drivers accumulate incident points during a race.

incident	incident points
Wheels off the racing surface	1
Loss of control	2
Contact with other objects	2
Heavy contact with another driver	4

Table 2: Summary Statistics

This table presents summary statistics of variables in the cleaned data set (2,274,192 observations from January 1st to December 31st, 2015). We translate alphabetical License levels to numerical values: class D (2), class C (3), class B (4), class A (5), class Pro (6), and class Pro World Class (7).

	Mean	Std. Dev.	Min	Max
Incidents per Mile (safety measure)	0.21	0.48	0.00	48.00
WHP (product safety)	7.78	4.11	1.77	17.89
License (consumer skills)	3.83	1.12	2	7
Traffic Density (control)	14.33	4.21	1.00	45.96
Laps in the Race (control)	29.93	16.69	1	160
Oval (control)	0.49	0.50	0	1
Night (control)	0.20	0.40	0	1

Table 3: Hypotheses 3 and 4

In each regression, the dependent variable is Incident Point per Mile. Estimates are obtained through OLS regressions. Significance levels have been calculated using heteroscedastic robust standard errors (*** 1%, ** 5%, * 10%). In Model 1, we test whether product regulation and consumer regulation result with the regular or Peltzman effect. In Models 2 and 3, we determine whether omitting relevant variables affects our results. In Model 2, we analyze product regulation but do not control for consumer skills, measured by the variable License. In Model 3, we analyze consumer regulation but do not control for product safety, measured by the variable WHP (Weight to Horse Power Ratio).

	Model 1	Model 2	Model 3
WHP	-0.0002^{***}	0.0068^{***}	
License	-0.0671^{***}		-0.0669^{***}
Traffic Density	-0.0083^{***}	-0.0115^{***}	-0.0083^{***}
Laps in the Race	0.0015^{***}	0.0011^{***}	0.0015^{***}
Oval	-0.0534^{***}	-0.0405^{***}	-0.0529^{***}
Night	0.0083^{***}	0.0216^{***}	0.0082^{***}
N	2,274,192	2,274,192	2,274,192
Adj R^2	0.0378	0.0189	0.0378

Table 4: Hypothesis 2 (part I)

Estimates are obtained through OLS regressions. Significance levels have been calculated using heteroscedastic robust standard errors (*** 1%, ** 5%, * 10%). In column "License J" (where J stands for D, C, and B), we consider only the observation with the value of License equal to J.

	License D	License C	License B
Dependent variable: Relative Incident Points			
License Difference	-0.3356^{***}	-0.2850^{***}	-0.2664^{**}
High Skills	-0.2204^{***}	-0.0411^{***}	-0.0721^{**}
Adj R^2	0.0142	0.0132	0.0126
Dependent variable: Relative Final Position			
License Difference	-0.0870^{***}	-0.0672^{***}	-0.0596^{**}
High Skills	-0.0630^{***}	-0.0369^{***}	-0.0329^{**}
Adj R^2	0.0267	0.0199	0.0164
Dependent variable: Quartile Position			
License Difference	-0.3602^{***}	-0.2816^{***}	-0.2648^{**}
High Skills	-0.2233^{***}	-0.1419^{***}	-0.1188^{**}
Adj R^2	0.0286	0.0232	0.0209
N	46,995	60,504	66,917
Promoted	40,025	50,869	53,168
Demoted	6,970	9,635	13,749

Table 5: Hypothesis 2 (part II)

For License D, C, and B, we provide mean of Weight to Horse Power Ratio and related statistics for high-skill drivers and low-skill drivers. At License J, the high-skill drivers are those who have been demoted from License $J+1$ to License J, and the low-skill drivers are those who have been promoted from License $J-1$ to License J. Lower Weight to Horse Power Ratio means that the car is less safe/more difficult to control. For every License, the difference of the means is significant at the 1% level.

	Mean	Std. Deviation	Std. Error	95% Interval	
License D					
Low skills	11.601	3.930	0.0196	11.562	11.639
High skills	10.545	3.907	0.0468	10.453	10.637
License C					
Low skills	8.618	4.235	0.0188	8.581	8.655
High skills	8.053	4.228	0.0431	7.968	8.137
License B					
Low skills	7.074	3.667	0.0159	7.042	7.105
High skills	6.631	3.490	0.0298	6.573	6.690

Table 6: Hypothesis 5

In each regression, the dependent variable is Incident Point per Mile. Estimates are obtained through OLS regressions. Significance levels have been calculated using heteroscedastic robust standard errors (*** 1%, ** 5%, * 10%). We consider a sub-sample consisting only of the observations with Weight to Horse Power Ratio (WHP) greater than 9.5. In column "All Licenses," we test whether product regulation results with the regular or Peltzman effect. In column "License J" (where J stands for D, C, B, and A), we truncate our sub-sample and consider only the observation with the value of License equal to J.

	All Licenses	License D	License C	License B	License A
WHP	-0.0013^{***}	-0.0037^{***}	-0.0032^{**}	0.0019^{***}	0.0051^{***}
License	-0.0790^{***}				
Traffic Density	-0.0094^{***}	-0.0147^{***}	-0.0094^{***}	-0.0075^{***}	-0.0039^{***}
Laps in the Race	0.0007^{***}	0.0015^{***}	0.0003	0.0002	0.0003
Oval	0.0842^{***}	0.0835^{***}	0.0702^{***}	0.0315^{*}	0.0293^{**}
Night	0.0256^{***}	0.0228	0.0329^{*}	0.0789^{***}	-0.0345^{***}
N	493,454	198,476	122,467	96,266	76,074
Adj R^2	0.0327	0.013	0.010	0.008	0.006

Table 7: Hypothesis 6

In each regression, the dependent variable is Incident Point per Mile. Estimates are obtained through OLS regressions. Significance levels have been calculated using heteroscedastic robust standard errors (*** 1%, ** 5%, * 10%). We consider a sub-sample consisting only of the observations with Weight to Horse Power Ratio (WHP) smaller than 7. In column "All Licenses," we test whether product regulation results with the regular or Peltzman effect. In column "License J" (where J stands for D, C, B, or A), we truncate our sub-sample and consider only the observation with the value of License equal to J.

	All Licenses	License D	License C	License B	License A
WHP	0.0028***	0.0869***	−0.0028**	−0.0028***	−0.0010***
License	−0.0629***				
Traffic Density	−0.0055***	−0.0106***	−0.0096***	−0.0064***	−0.0037***
Laps in the Race	0.0010***	0.0104***	0.0027***	0.0011***	0.0001***
Oval	−0.0777***	−0.2475***	−0.1351***	−0.0907***	−0.0405***
Night	0.0146***	−0.0207***	0.0109***	0.0204***	0.0206***
N	1,362,597	62,779	220,966	401,031	673,678
Adj R^2	0.0351	0.024	0.021	0.014	0.006

Table 8: Two-dimensional regulation

In each regression, the dependent variable is Incident Point per Mile. Estimates are obtained through OLS regressions. Significance levels have been calculated using heteroscedastic robust standard errors (*** 1%, ** 5%, * 10%). In column "License at least J" (where J stands for C or D), we truncate our data and consider only the observation with the value of License equal or higher than J.

	License at least C	License at least B
WHP	-0.0006^{*}	-0.0010^{***}
License	-0.0505^{***}	-0.0504^{***}
Traffic Density	-0.0054^{***}	-0.0046^{***}
Laps in the Race	0.0007^{***}	0.0004^{***}
Oval	-0.0705^{***}	-0.0581^{***}
Night	0.0194^{***}	0.0209^{***}
N	1,299,818	1,078,852
Adj R^2	0.0236	0.0146