How Do People React to Income-Based Fines? Evidence from Speeding Tickets Discontinuities*

Martti Kaila

University of Glasgow

June 22, 2026

Abstract

This paper studies the impact of income-based criminal punishments on crime. In Finland, speeding tickets become income-dependent if the driver's speed exceeds the speeding limit by more than 20 km/h, leading to a substantial jump in the size of the speeding ticket. Contrary to predictions of a traditional Becker model, individuals do not bunch below the fine hike. Instead, the speeding distributions are smooth at the cutoff. However, I demonstrate that the size of the realized speeding ticket has sizable, but short-lived, impacts on reoffending. I use a regression discontinuity design to show that fines that are, on average, 200 euros larger decrease reoffending by 15 percent in the following six months. The drop in reoffending is driven by high-income individuals facing the highest fine at the cutoff. I estimate that a fixed fine hike that matches the current deterrence effect would raise the fine increase faced by the bottom income quartile by about 300% at the cutoff, relative to the increase they face under the current income-based schedule. My empirical results are consistent with an explanation that people operate under information frictions. To illustrate this, I construct a Becker model with misperception and learning that can explain all the empirical findings.

* *martti.kaila@glasgow.ac.uk* marttikaila.com I am particularly thankful for Kristiina Huttunen, Emily Nix, Roope Uusitalo, and Markus Jäntti, who have provided guidance, help, and inspiration for this project. I thank Jennifer Doleac, Daniel Hauser, Randi Hjalmarsson, Ramin Izadi, Tuomas Kosonen, Dave Macdonald, Tuomas Matikka, Teemu Pekkarinen, Tuomas Pekkarinen, Mikko Silliman, Christian Traxler, and Tuuli Vanhapelto for helpful comments and discussions. In addition, I thank seminar participants at Aalto University, Bocconi University, CES IFO Labor Economics Conference, EALE Conference, Helsinki GSE Labor and Public Economics Workshop, Criminological Seminar at the University of Helsinki, Vrije Universiteit Amsterdam, and University of Glasgow.

1 Introduction

One of the primary roles of government is to deter people from committing actions that cause harm to others. For example, since traffic accidents are the most common unnatural cause of death globally (WHO, 2017), governments worldwide use speeding tickets to discourage drivers from driving at speeds that pose a danger to others. Such sanctions typically increase with offense severity, consistent with standard economic theory (Becker, 1968). In some countries, fines also scale with offenders' incomes. Yet little is known about how these income-based fines affect behavior across the income distribution.

This paper studies how income-based fines affect behavior, using high-quality register data and a unique context where the magnitude of fines increases with a person's income. Specifically, in Finland, speeding tickets become income-dependent if the driver's speed exceeds the limit by more than 20 km/h. Fines are allowed to increase without an upper limit at the cutoff. This policy means that high-income drivers have a substantial incentive to slow and bunch below the income-based threshold, since failing to avoid the income-based cutoff can be very expensive. For example, in 2019, the police assigned NHL ice hockey player Rasmus Ristolainen an income-based speeding ticket equal to approximately 120,000 euros.[1]

The first key result of the paper is that individuals ignore the vast discontinuous jumps in the severity of the punishment. This directly contradicts the seminal model of Becker (1968), which predicts that people should react to discontinuities in the price of speeding at the 20 km/h cutoff by slowing down and bunching just below the income-based fine cutoff. Specifically, working with the police in Finland, I obtained detailed police speeding ticket data which I link to administrative tax data. Using this data, I find zero excess mass below the income-based fine cutoff in the speeding distributions. Surprisingly, even high-income drivers fail to bunch below the cutoff, even though they face a considerable incentive to avoid the income-based fine.

The rest of the paper focuses on understanding why people ignore this stark discontinuity in punishments and the broader implications of an income-based versus fixed fine system. I propose two potential explanations that could attenuate individuals' reactions to incentives, both of which have been previously applied in other non-crime contexts (Chetty, 2012; Kleven and Waseem, 2013). First, I examine the possibility of information frictions: drivers are unaware of the system's details, leading to suboptimal speed decisions around the cutoff. Second, I consider the possibility of adjustment costs that cause individuals not to react to

[1] Finnish and international media frequently cover the Finnish income-based fine system. See, for example, articles by The New York Times and The Atlantic.

the cutoff. The adjustment costs explanation posits that individuals operate under perfect information but neglect the cutoff since optimization is mentally too costly relative to the expected punishment.

Both of these possible frictions can generate smooth speed distributions around the discontinuity *ex-ante*, but the two friction types imply very different predictions on how realized fines should affect recidivism *ex-post*. If individuals react to a realized fine, then this reaction suggests imperfect information, with individuals learning from the punishment. In contrast, under the adjustment cost story, the experienced fine should not change an individual's behavior since the ticket is just a realization of a rational bet.

The second main result of the paper is that people *do* react to the severity of realized punishments ex-post. Since the speeding distributions are continuous at the cutoff, I use a regression discontinuity design (RDD) to study how larger realized fines, due to the income-based system, affect reoffending. According to my RDD estimates, individuals assigned larger fines due to the income-based fine system are less likely to commit another traffic offense in the short term. The impact of larger fines on reoffending peaks around 4-8 months after the income-based fine. Those assigned, on average, a 200 euro larger fine are approximately 2-3 percentage points less likely to commit another traffic crime in the following 4-8 months. Compared to the average speeding behavior of the speeders who receive a smaller fixed fine, this estimate implies a 15-20 percent reduction in recidivism. In addition, reoffenders' speeding distributions evolve smoothly at the income-based fine cutoff. Hence, larger fines trigger an extensive margin response.

The RDD results lead to two conclusions. First, they suggest that individuals learn from punishments, implying that information frictions at least partly cause the smooth distribution around the discontinuity. Second, this learning response results in less speeding in the short term after fines are realized, but the fine hikes do not provide deterrence ex-ante.

Importantly, the specific deterrence response is sharply heterogeneous by income. I show that high-income individuals receiving larger fines at the income-based cutoff are less likely to reoffend in the short term than low-income individuals receiving smaller fines. To do this, I study how the RDD estimates differ by income quartiles. The fine jumps by around € 50 (500) in the bottom (top) income category. For cumulative reoffending over six months, the point estimate is close to zero (0.1 pp.) for the bottom quartile, while in the top quartile the point estimate of -4.5 pp. is double the main estimate. In addition, the different income quartile groups respond similarly to small fixed-fine jumps. This suggests that high-income individuals' stronger reaction to larger fines is mainly because of the larger fines and not because low- and high-income drivers differ on some other dimension.

Lastly, my RDD results suggest that the impact of the larger fine on recidivism fades out over time. The effect of larger fines becomes statistically insignificant 12 months after the initial speeding incident. This result suggests that despite the heftier fines leading to a substantial short-term learning reaction, drivers may forget in the long run.

What model could simultaneously explain smooth speeding distributions, reactions to realized fines, and possible long-run fade-out in the effect of fines? To rationalize these findings, I construct a Becker model with misspecification and learning to capture drivers' behavior. The model builds on the idea that if individuals find pricing schemes too complex to understand, they may use the so-called "ironing" heuristic, where people approximate marginal prices with average prices (Rees-Jones and Taubinsky, 2019). Drawing from this insight, I assume that drivers find the actual penalty function too complex. Therefore, they make speeding decisions using a rule of thumb that the speeding ticket increases linearly with speed. Based on their speed, drivers may receive a fine. If the fine is larger than predicted, they conclude that the relationship between speed and fine is steeper than expected and update their beliefs accordingly.

I show that this Becker model with misperception and learning rationalizes all my empirical findings. First, the "ironing" heuristic causes drivers to ignore the fine discontinuity, leading to a smooth speeding distribution. Second, due to the jump in the true penalty function, drivers may receive larger fines than they expected. This steepens their linearized penalty function and reduces speeding in the next period. Furthermore, the larger the true discontinuity in the fine, the more significant the driver's reaction *ex-post*. Finally, I show that when individuals learn with the misspecified model, there are drivers whose beliefs and actions do not converge but follow cycles.[2] This is consistent with the fade-out over time in the impact of receiving a larger-than-expected fine at the discontinuity.

In the third and final main empirical result, I use my heterogeneous treatment-effect estimates to quantify the distributional consequences of replacing income-based fines with a fixed fine increase that holds reoffending constant at the 20 km/h cutoff. The implied safety-neutral increase is €215. Although the size is close to the average observed increase (€204), the fixed fine reallocates penalties sharply across the income distribution. Bottom-quartile drivers would face an increase of €215, up from €52 under the current system, while top-quartile drivers would see their increase fall from €499 to €215. In levels, the implied fixed fine above the cutoff would amount to 40% of monthly net income for bottom-quartile drivers, compared with about 25% under the current income-based system.

This paper provides new evidence on how income-based fines shape deterrence and the

[2] An extensive theoretical literature studies learning when an individual's model is misspecified (See, for example, Berk (1966), Bohren and Hauser (2021), Esponda and Pouzo (2016), or Nyarko (1991)).

incidence of monetary sanctions across the income distribution. First, I show that income-based fines do not generate a deterrence response at the 20 km/h cutoff, despite large financial incentives to slow down. This speaks to the literature on *general deterrence*, which examines how the *threat of punishment* discourages crime in the general public, a central question in the economics of crime. Economic theory predicts that harsher punishments should deter crime, but the evidence is mixed (Chalfin and McCrary, 2017). Lee and McCrary (2017) find little behavioral response to the sharp increase in punishment severity at age 18. In contrast, Traxler *et al.* (2018) show clear evidence of bunching below speeding fine notches in Germany, where the fine schedule is stable and identical across drivers.

The contrast between my findings and those of Traxler *et al.* (2018) underscores the importance of institutional context. Individuals may react strongly to simple and predictable incentives but ignore large ones when the structure is complex or not salient. This distinction matters for policy design. If individuals misperceive or fail to understand complex sanctioning rules, as my results suggest, then policy effectiveness may depend less on the size of penalties and more on their salience and comprehensibility.[3]

Second, I show that income-based fines reduce reoffending, with the strongest responses among high-income drivers who face the largest fine increases. This finding contributes to the literature on *specific deterrence* effects of punishment, which examines how *experiencing punishment* deters individuals (Gehrsitz, 2017; Hansen, 2015; Finlay *et al.*, 2024; Drago *et al.*, 2009).[4] My paper is most closely related to Dušek and Traxler (2022) and Goncalves and Mello (2023).[5] Dušek and Traxler (2022) find that speeding tickets reduce reoffending at an extensive margin (no ticket vs. ticket) but not at an intensive margin (low vs. high ticket). In contrast, Goncalves and Mello (2023) observe that a speeding ticket hike decreases future speeding at the intensive margin. I extend this evidence by showing that the behavioral response scales sharply with fine size across income groups. High-income drivers are also more likely to speed and reoffend, suggesting that income-based fines concentrate behavioral change among individuals with higher baseline risk.

Combining these results, I quantify the distributional consequences of replacing income-based fines with a safety-neutral fixed fine, thereby providing a direct, policy-relevant comparison of the two systems while holding reoffending constant. A safety-neutral fixed fine

[3] Prior work highlights the importance of speed limits and enforcement for traffic safety (Bauernschuster and Rekers, 2022; DeAngelo and Hansen, 2014; van Benthem, 2015).

[4] The term "specific deterrence" is more common in criminology (Doleac, 2023), while economists typically frame such behavior as information updating (Chalfin and McCrary, 2017). Identifying specific deterrence effects is difficult, as punishments may influence recidivism through alternative channels, such as peer effects. Studies of incarceration often capture multiple mechanisms, making results context-dependent. For example, Aizer and Doyle (2015) and Mueller-Smith (2014) find that incarceration increases recidivism, whereas Bhuller *et al.* (2020) and Rose and Shem-Tov (2021) find the opposite.

[5] Other notable papers that study learning in the context of crime and law enforcement are by Banerjee *et al.* (2019), Philippe (2024), Rincke and Traxler (2011), Drago *et al.* (2020), and Lochner (2007).

would be substantially more regressive, shifting a larger share of the financial burden toward low-income drivers, who are more likely to experience financial hardship from fines (Mello, 2024). Because a fixed fine would sharply reduce penalties for high-income drivers while pushing penalties for low-income drivers further into the range in which fine collection rates may decline (Norris and Rose, 2024), moving to a fixed-fine cutoff could also reduce government revenue.

Finally, my paper also relates to the literature studying individuals' reactions to discontinuities in the choice sets and the role of optimization frictions (for a review, see Kleven (2016)). A common finding is that discontinuities, for example, in tax rates, produce evident bunching responses (Bastani and Selin, 2014; Chetty *et al.*, 2011, 2013; Kleven and Waseem, 2013; Harju *et al.*, 2019). However, these bunching reactions often translate into small elasticities due to optimization frictions. I contribute to this literature by documenting that people may completely ignore substantial price jumps due to information friction but learn from price signals. This evidence is consistent with the findings in other contexts that people do not use marginal prices in decision-making (Rees-Jones and Taubinsky, 2019; Kostøl and Myhre, 2021; Ito, 2014).

The remainder of the paper is organized as follows. Section 2 provides institutional details and describes the data. Section 3 shows the speeding distribution. Section 4 investigates whether the size of the speeding ticket impacts recidivism. Section 5 introduces a framework that seeks to rationalize my findings. Section 6 uses the heterogeneous treatment-effect estimates to quantify the distributional implications of the safety-neutral fixed fine across the income distribution. Section 7 concludes.

2 Institutional Setting and Data

2.1 The Relationship Between Speed, Income, and Speeding Tickets

Figure 1 gives an illustrative example of the speeding punishment schedule when the speed limit is 100 km/h. Other possible speed limits are 20, 30, 40, 50, 60, 70, 80, and 120 km/h, and the implications are similar. Given the 100 km/h limit and a 6 km/h grace speed, so long as an individual drives at or below 107 km/h, there is no ticket given.

Two factors determine the size of the speeding ticket in Finland. First, if a speeding violation is minor, i.e. only 7-20 km/h above the posted limit, then the individual receives a fixed fine equal to €140-200.

Second, if an individual's speeding violation is more than 20 km/h over the limit, the day-fine system kicks in, and speeding tickets become income-based. Under the income-based fine system, the size of the fine is determined by the product of the offense severity and an individual's income. Depending on the severity of the excess speeding, the individual receives 1-120 fine units called day-fines. In cases of speeding, the higher the observed speed, the more severe the offense, and hence the higher the number of day-fines the individual receives. The following rule governs the monetary value of a single day-fine

$$I = \frac{Y - 255}{60} - 3 \times D, \tag{1}$$

where I is the monetary value of a single fine, Y is the net monthly income, and D is the number of dependents.[6]

Finland introduced the income-based fine system in 1921 for two reasons. The first motivation was equity. The government wanted to design a fine system that treated everyone similarly, regardless of income. Second, the government was concerned that the high inflation at the beginning of the 20th century would erode the deterrence effect of the fine system. One way to overcome this problem was to link the value of fines to income. The name "day-fine system" originates from the feature that the value of a single day-fine was initially set to be equal to the amount of salary a worker would lose if he had to spend one day in prison instead of working. (HE, 1920)[7]

Figure 1(a) provides a theoretical example of how the size of the fine changes at the income-based fine cutoff for different income quartiles.[8] The fixed fine amount is identical across all income groups and applies for 7-20 km/h above the speed limit. Once the speed exceeds the limit by 20 km/h, the income-based fine system takes effect, and the ticket size can jump substantially, depending on the driver's income. For the top income quartile, the fine jumps by around 400 euros at the cutoff. In contrast, for the bottom quartile, the fine does not change at the cutoff. Although Figure 1(a) illustrates the system using discrete income groups, in practice, the income-based fine schedule produces a continuum of discrete jumps.

To illustrate the relative magnitude of fines across income groups, Figure 1(b) relates the fine size to an individual's net monthly income. Before the income-based fine cutoff, lower-income individuals pay a larger share of their income as fines. However, after the income-based fine threshold, the fine equals around 15-20 percent of net monthly income for each income group, except for the lowest income group. In this sense, the income-based fine system may be viewed as more equitable, as it imposes fines representing a similar share

[6]Net income consists of the sum of all taxable earned income, capital income, employee benefits, pensions, and most social benefits. Police always calculate the fine based solely on the driver's income. The income of the driver's family members does not affect the fine.

[7]The day-fine system is used in courts as well.

[8]The quartile-specific income stated in the figure is the average net income over the years 2015-2017 calculated using microdata.

Figure 1: Relationship Between Income and the Size of Speeding Ticket

(a) Theoretical example when the speeding limit is 100 km/h

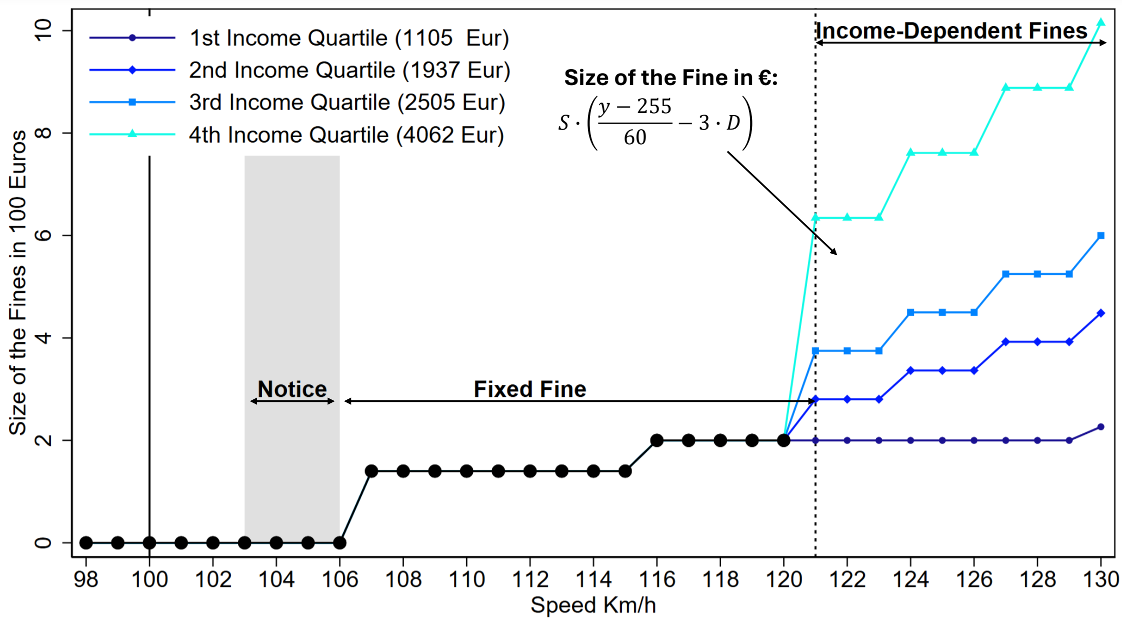

(b) Size of the Fine w.r.t. Monthly Net Income

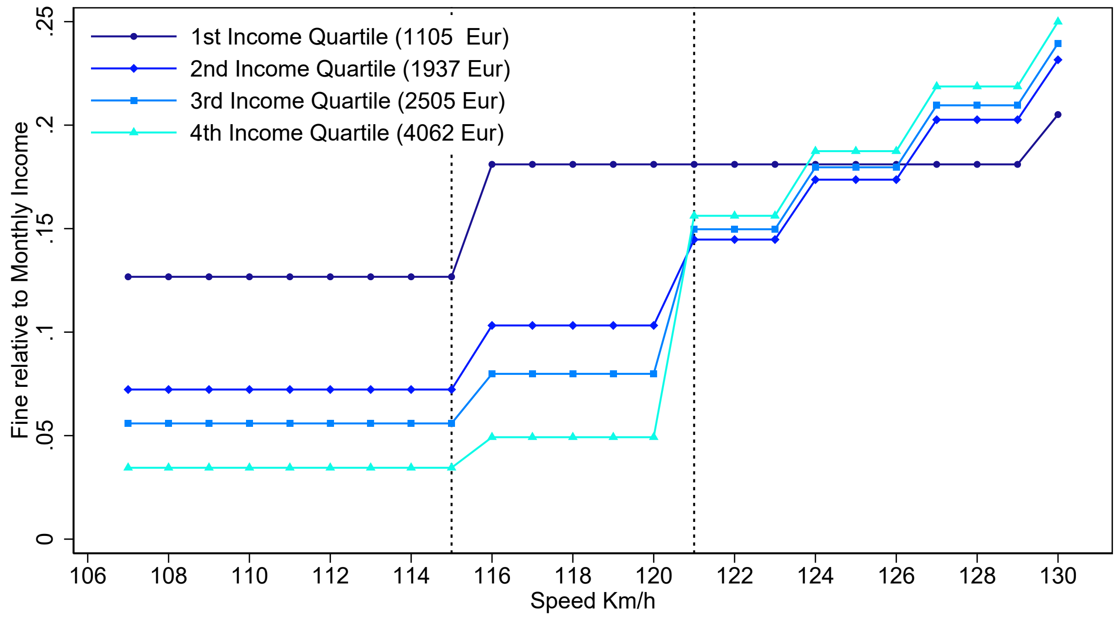

Notes: Figure Panel a shows the theoretical relationship between the size of the speeding ticket, speed, and net income for different income quartile groups when the speeding limit is 100. Figure Panel b relates the speeding ticket's size to the quartile's average net income. Average net incomes within quartiles are calculated using microdata. In the equation on top of the figure, S refers to a number of income-based fine units, y to net monthly income, and D to a number of dependents. The example presented in the figure assumes that the number of dependents is zero.

of income across income groups.

Individuals just on the right-hand side of the income-based fine cutoff usually receive around 10-12 income-based fine units. After the cutoff, the number of income-based fine units increases in a stepwise manner. Appendix Table B.1 provides the exact punishment schedules that the police use in Finland to punish speeding.

Two additional aspects of the system are essential to this study. First, the income-based ticket cannot be lower than the fixed fine equal to 200 euros.[9] Second, the income-based speeding ticket schedule does not have an upper limit in euros. This has resulted in some extraordinarily large speeding tickets, with some of the largest tickets in the past equalling around 100,000-120,000 euros. The current unofficial record holder is professional National Hockey League player Rasmus Ristolainen, who was fined €120,680 in 2019 for speeding.

2.2 Speeding surveillance

Speeding is monitored in two ways in Finland. First, the police use fixed surveillance stations or automatic traffic surveillance vehicles that measure the speeds of passing cars. If the surveillance station observes a speeding car, it takes a picture that the police use to assign a fine. Second, traditional police patrols monitor speeding with radar and laser speed guns. When a patrol observes a driver that exceeds the speeding limit, they may stop the car and issue a notice or fine to the driver. I use only data created by the cameras, since their measurements are not affected by decisions such as whom they should stop.[10] Appendix Figure A.3 presents a picture of a typical police speeding surveillance camera.

After the surveillance machine has taken a picture, it is sent to the Police Traffic Safety Center, where an officer identifies the car owner based on the registration number. Next, the police adjust the speed downward due to measurement error and then decide, based on the adjusted speed, whether to issue a notice or fine to an individual.[11]. All the numbers I report in this paper are adjusted speeds. Finally, if the police assign a fine or notice, the owner receives an electronic message and letter at home by post within 30 days from the date of the offense. The message notifies the individual of the fine and its size.[12]

The police calculate income-based fines using an individual's income from the most recent tax decision. Family members' incomes do not affect the fine size. If an individual's income has changed, it can be adjusted to reflect more recent information. I do not observe

[9] The fine stays constant at the cutoff for a person whose net income is around 1455 euros per month. If an individual earns less than this, the fine does not change at the cutoff.

[10] For example, Goncalves and Mello (2021) provides evidence that in Florida, U.S. officers discriminate against minorities by not discounting their speed, which leads to larger fines for minority groups.

[11] The correction is 3 km/h if the speed is less than 100 km/h and 3% if the speed is larger than 100 km/h

[12] The police always send the ticket to the owner of the car. The owner must request an administrative review if the owner was not the driver when the speeding incident occurred.

individuals' monthly earnings at the time of a speeding incident. Instead, I proxy monthly earnings using annual net income, which may introduce modest measurement error.

Frequent speeding or extreme speeding can also lead to a temporary suspension of an individual's driver's license. Three speeding tickets within a year or at least four tickets within two years lead to a driving suspension lasting between 6-18 weeks. However, the suspension policy does not change at the income-based fine cutoff, which makes it easier to interpret the regression discontinuity design estimates I will estimate and describe. Further, police may suspend an individual's driver's license for 1-6 months if the individual's speed exceeds the speeding limit of more than 35 km/h.

2.3 Data

Data sources My primary data set comes from the National Police Board of Finland and includes all the speeding tickets that the police gave between April 2018 and May 2020 (National Police Board of Finland, 2020).[13] The police data contains rich information such as the recorded speed, the speeding event's location, the prevailing speed limit, and whether an officer or an automatic camera system measured the speed. The data also include the monetary value of fixed fines, the value of a single day-fine, and the number of day-fines assigned.

A key component of the police data is that each speeding ticket in the data includes a unique personal identifier, which I use to link the police speeding data to other administrative registers. First, I merge the speeding data into Statistics Finland's crime report register, which spans between 2006-2020 (Statistics Finland, 2022a). The crime report data contains the same speeding events as the police speeding data, but it does not include information about the exact speed. Thus, I cannot use it to plot speeding distributions, but I can construct the recidivism outcomes using this data. This approach increases my sample size.[14]

Further, I merge the speeding data into Statistics Finland's FOLK data module, which contains a full population of Finnish residents between 1987-2019. From this data set (Statistics Finland, 2022b), I observe individual's labor market outcomes such as labor earnings, income, and employment, and basic demographics like age, municipality of residence, marital status, and whether a household owns a car.

To investigate the impacts on accidents (Statistics Finland, 2023), I use traffic accident data from Statistics Finland, which covers police-reported accidents between 1989 -

[13]Due to a reform, I do not observe fixed fines after May 2020.

[14]With this approach, I can use observations from the beginning of my sample period for which I would not observe any pre-outcomes without the crime report data

2022. The dataset includes information on the parties involved, the number of injuries, the date of the accident, and the municipality in which it occurred. Importantly, it contains unique personal identifiers that allow me to link individuals involved in accidents to other administrative records.

I also conduct analysis using traffic data collected by the Finnish Transport Agency using automatic traffic monitoring system (TMS) stations (Fintraffic, 2020). There are around 500 TMS stations scattered around Finland. These TMS stations observe if a vehicle passes the station and record information such as time, speed of the vehicle, and vehicle class. Using the data collected by TMS, I can plot the speed distribution of all drivers, not just those caught speeding by cameras. Unfortunately, TMS data is anonymous. Hence, I cannot observe information on the identity of the driver.

Sample Restrictions The first restriction I impose is that I only focus on speeding incidents measured by automatic cameras. I ignore police patrol speeding data because prior evidence shows that police officers may manipulate the observed speed, complicating the interpretation of possible bunching (Goncalves and Mello, 2023). Second, I omit speeding tickets given in areas where the limit is 120 km/h since there are too few observations. Finally, I only take individuals whose personal identifier numbers are available in the police data and who are at least 18 years old at the time of the speeding ticket.[15]

Outcome Variables In the second part of the paper, I study whether the size of the fine impacts recidivism. For this analysis, I use police crime report data to follow individuals for 12 months after the initial speeding ticket and calculate a cumulative reoffending probability for each month. I measure reoffending by constructing an indicator variable that takes the value one if an individual has committed a new traffic crime in the current month or any past months subsequent to the initial speeding ticket.[16] I am able to follow the same individuals for five months, but after 6 months my sample starts to decrease. However, in the robustness section, I show that my results are insensitive to this limitation. I also use crime report data to build variables measuring non-traffic criminal activity. Like the primary outcome, the non-traffic crime indicator takes the value of one if an individual has committed a non-traffic crime 1-12 months after the initial speeding ticket. Finally, I examine whether the magnitude of the fine affects the probability of being involved in a traffic accident as a driver within six months of receiving a speeding ticket.

Background Variables. Finnish register data contain a rich set of background variables such as an individual's gender, education level, age, and employment status. I use these variables to check the validity of the setting, as control variables, and to perform

[15]Personal identifiers are missing for foreigners who do not live in Finland permanently.

[16]Appendix Figure A.17 shows that I obtain similar results if the outcome is just speeding.

Figure 2: The Relationship between Income, Probability of Speeding, and Crime Report

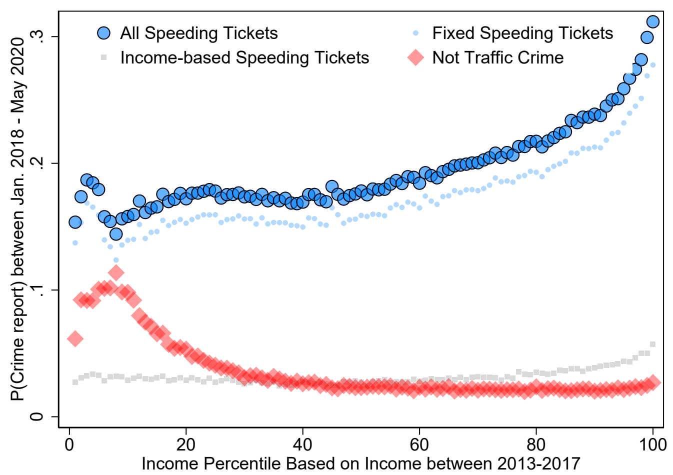

Notes: The figure presents the relationship between income percentile and the probability that an individual appears in police crime report records between January 2018 and May 2020. The sample includes individuals from birth cohorts 1953 to 1998. Income percentiles are calculated by comparing each individual's mean income over 2013–2017 with that of others in the same birth cohort. The analysis is restricted to individuals residing in households with a car. Blue dots with circles indicate the share of individuals in each percentile receiving a speeding ticket between January 2018 and May 2020. Red diamonds indicate the share of individuals appearing in police crime report records for suspected crimes other than traffic offenses during the same period. Small dots and grey squares further decompose the speeding ticket averages into the shares attributable to fixed fines (blue) and income-based fines (grey). Appendix Figure A.6 reports the average earnings within each percentile. Appendix Figures A.5 and A.7 demonstrate the robustness of the results across multiple specifications.

heterogeneity analysis. The most important variable is the net yearly income, which I use to approximate an individual's net monthly income at the time of the speeding incident.[17]

Descriptive Statistics Table 1 provides descriptive statistics of the background variables. Column 1 reports the means of all background variables in the sample that I obtain after imposing the restrictions described above. Column 2 reports the same statistics but in the small window around the income-based fine cutoff. On average, the individuals in my sample are more likely to be male, older, and have high net income.[18]

To give an idea of how common speeding tickets are and who receives fines, Figure 2 plots a relationship between crime and income among the Finnish adult population with a car. The figure reveals two interesting patterns. First, speeding is a common crime. Around 15 percent of individuals in the median income percentile have received a speeding ticket over a two-and-a-half-year follow-up period. By contrast, when the suspected offense is not a traffic crime, only around 2 percent of the median-income individuals show up in the police crime report data. Second, the association between an individual's income and speeding tickets differs substantially from the relationship between income and non-traffic

[17]For those who receive income-based fines, I can calculate their monthly net income from the income-based fine. Appendix Figure A.4 illustrates that these two measures are highly correlated.

[18]I find that the average net income of individuals who are at least 18 years old is approximately €2,000 in 2018.

Table 1: Descriptive Statistics and Balance Check

	All	**Window +-3**	**RDD**	**CCT RDD**
	Mean	Mean	RDD	RDD
	(1)	(2)	(3)	(4)
Traffic offence, months t-1 to t-6	0.122	0.134	0.014	0.018
			(0.011)	(0.013)
Non-Traffic offence, months t-1 to t-6	0.014	0.018	-0.002	-0.003
			(0.004)	(0.005)
Employed	0.706	0.708	0.002	0.004
			(0.013)	(0.016)
Unemployed	0.032	0.037	0.016	0.019
			(0.006)	(0.007)
Outside the Labor Force	0.262	0.255	-0.019	-0.024
			(0.012)	(0.015)
Monthly Net Income	2,590	2,572	-77.374	-75.269
			(42.134)	(52.966)
Primary Education	0.164	0.168	0.003	0.007
			(0.011)	(0.013)
Secondary Education	0.434	0.442	0.002	0.006
			(0.013)	(0.016)
Tertiary Education	0.402	0.391	-0.005	-0.008
			(0.014)	(0.018)
Female	0.331	0.298	-0.002	-0.001
			(0.011)	(0.014)
Age	47.296	45.926	-0.947	-0.895
			(0.432)	(0.536)
N of Children	0.834	0.847	0.062	0.081
			(0.037)	(0.046)
Finnish Speaking	0.886	0.884	0.021	0.024
			(0.010)	(0.012)
Married	0.501	0.481	0.000	0.006
			(0.014)	(0.017)
Urban Municipality	0.729	0.716	0.014	0.021
			(0.013)	(0.016)
Semi-urban Municipality	0.160	0.166	-0.015	-0.017
			(0.009)	(0.012)
Rural municipality	0.111	0.118	-0.002	-0.007
			(0.010)	(0.013)
Capital region	0.338	0.307	-0.003	0.002
			(0.012)	(0.016)
Observations	338,191	23,026		

Notes: The table shows the means of predetermined characteristics among individuals in the speeding ticket sample and results from balance checks. Column (1) shows the means of background variables in the estimation sample. Column (2) presents the means of background variables in the estimation sample after restricting individuals within +-3 km/h from the income-based cutoff. Column (3) shows results from the balance check where the dependent variable is a predetermined characteristic. The column reports the estimates of β obtained using the equation 2. Column (4) reports results from the balance check conducted using the approach of Calonico *et al.* (2014). I cluster standard errors at the individual level. I omit speed limit zones of 80 km/h and 120 km/h. See Section 2.3 for more details on sample construction.

crimes. Higher-income individuals are much more likely to receive a speeding ticket than lower-income individuals. This fact suggests that speeding is likely a luxury crime. In contrast, there is a strong negative gradient between income and non-traffic crimes.[19]

In addition, Figure 2 decomposes the speeding tickets into fixed fines (small blue dots) and income-based fines (small grey squares). The decomposition shows that the income gradient is much steeper among fixed fines than income-based fines. The disparity in the speeding probability by income is consistent with the idea that if individuals are risk-averse, fixed fines deter low-income people more than high-income people. Of course, other things may contribute to the gap. For example, wealthy individuals may drive more.

3 Do People Bunch at the Income-based Fine Cutoffs?

This section tests whether the income-based fine cutoff deters excessive speeding through general deterrence. A simple Becker model, in which individuals compare gains from speeding against expected punishment, predicts that drivers should react to a jump in the threat at the cutoff by bunching. In the model, the cutoff makes speeds above it discontinuously more expensive, implying that some drivers should respond by slowing down and bunching below the threshold.[20] If this prediction holds, we should see excess mass below the income-based fine cutoff in the speeding distribution.[21]

Speeding Distributions in Traffic Data I begin by plotting the speed distributions in the Finnish Transport Infrastructure Agency's (FTIA) traffic data. The automatic traffic monitoring stations scattered around Finland generate the data. These stations measure the speed of each passing car. The data allows me to see the total speed distribution, not only the speed distributions conditional on a driver going faster than the posted speed limit.

Figure 3 shows the first key result of the paper. The figure plots total speed distributions in the FTIA data. Figure Panel 3(a) plots the raw distributions in different speed limit zones. Figure Panel 3(b) pools together all speed limit zones and normalizes the x-axis variable as a distance from the income-based fine cutoff. The figures reveal two interesting patterns. First, each speed distribution peaks around the speed limit, suggesting that, on average, speed limits do strongly shape drivers' behavior.

Second, even though the speed distributions peak at the posted speed limits, Figure 3(b) shows that speed distributions exhibit no bunching around the points where the police's

[19]Appendix Figures A.5 and A.7 show that the results are robust to multiple alternative ways to construct income percentiles.

[20]Section 5 presents this result using a formal model.

[21]A common phenomenon in various contexts is that people react to discontinuous changes in incentives by bunching. Examples of this behavior include studies by Kleven and Waseem (2013), Einav *et al.* (2019), and Best *et al.* (2019)

tolerance level (-14) ends, fixed fines increase (-5), or income-based fines turn on (0). The fact that we do not see any mass just below these fine thresholds suggests that drivers do not react to fine discontinuities as the simple Becker model would predict.

However, it is possible that high-income individuals react to fines, but the behavior of drivers whose incentives change little at the cutoff masks the bunching by high-income individuals. Unfortunately, the FTIA data does not contain information on the driver. To overcome this shortcoming, next I move to use the police speeding data, which allows me to observe each driver's income.

Speeding Distributions in Police Data Figure 4 plots the speeding distributions using the police speeding camera data.[22] Figure 4(a) shows the distributions at different speed limit zones. The figure illustrates that speed distributions evolve strikingly smoothly. In other words, drivers do not bunch below the speeding ticket hikes highlighted by vertical lines, despite significant incentives to do so.

In Figure 4(b), I plot the pooled and normalized speeding distributions for the income quartile groups, which I define using predetermined income. Light (dark) blue dots show the speeding distribution for the top (bottom) income quartile group. We see that even the top income quartile's speeding distribution evolves smoothly at the cutoff, although they have the largest incentives to bunch below the threshold. Overall, my findings suggest that the discontinuous increase in fines induced by the income-based system does not generate general deterrence effects at the cutoff.

It is worth emphasizing that drivers may face different incentives at traffic monitoring stations compared to camera sites, as the former generally lack cameras. Nonetheless, the probability of enforcement remains positive at the stations due to traditional police crackdowns. Thus, while incentives to bunch may be weaker outside camera regions, we would expect to see bunching in a frictionless world.

Figure 4(a) shows that we observe some excess mass at speed 96 km/h when the speed limit is 80 km/h. However, this spike is unlikely to arise from the driver's behavioral response. The mass point is just after a fixed fine hike, implying that people do not bunch to avoid larger fines.

One explanation for the lack of bunching is that individuals know the locations of police surveillance stations and slow down just before cameras to avoid speeding tickets. If this is the case, the speeding ticket sample contains individuals who either do not care about the cameras or do not know their location. However, the fact that the total speed distributions shown by Figure 3 behave continuously everywhere suggests that the smoothness of the

[22] The police assign a fine only if the driver's speed crosses the police's tolerance level. This means that the speed ticket data reveals the speed distribution's right tail.

Figure 3: The Speed Distributions in Traffic Data

(a) Raw Speed Distributions

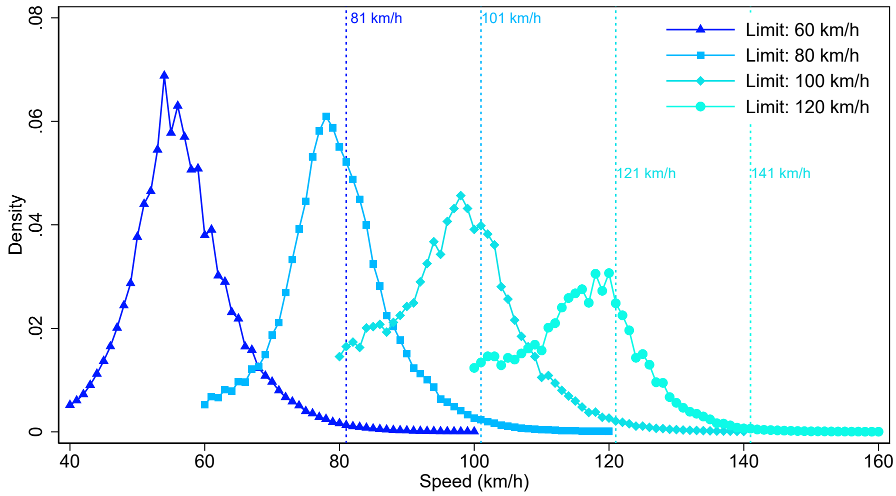

(b) Pooled Speed Distribution

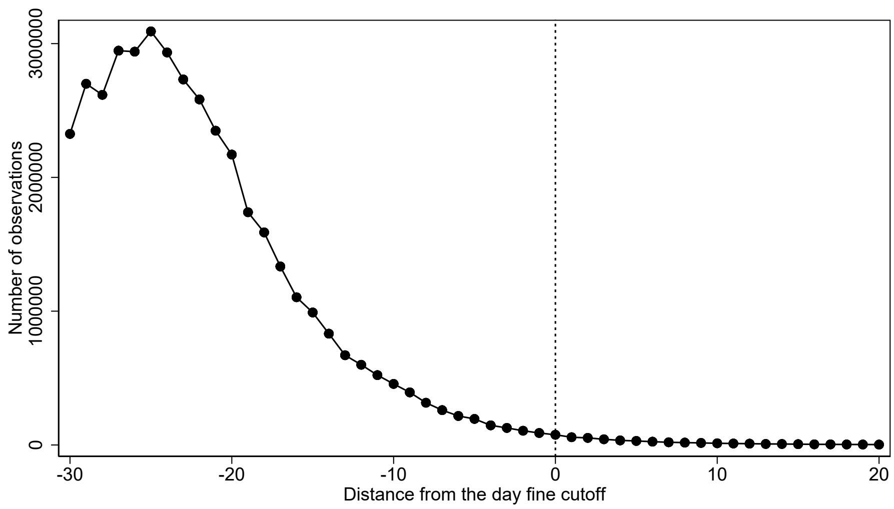

Notes: The figure presents the speed ticket distribution in the data created by the Finnish Transport Agency's automatic traffic monitoring stations. Figure panel (a) shows the speed distributions separately in different speed limit zones. Figure panel (b) pools all speed distributions together and normalizes the x-axis variable as a distance from the income-based fine cutoff.

Figure 4: The Speeding Distributions in Police Camera Data

(a) Speeding Distributions

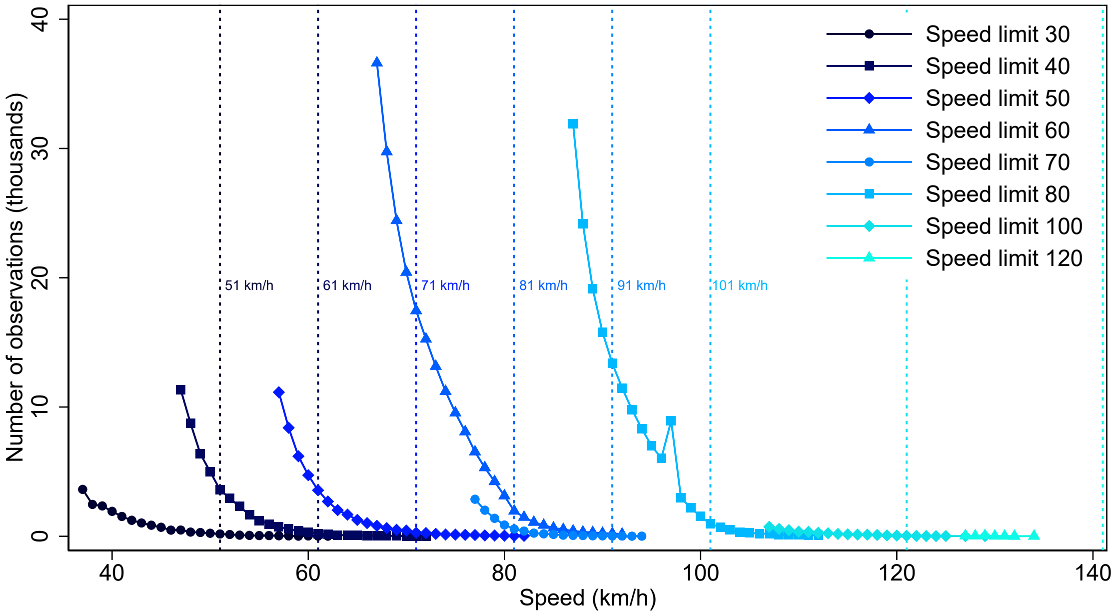

(b) Normalized Distributions for Income Quartile Groups (limit 80 km/h omitted)

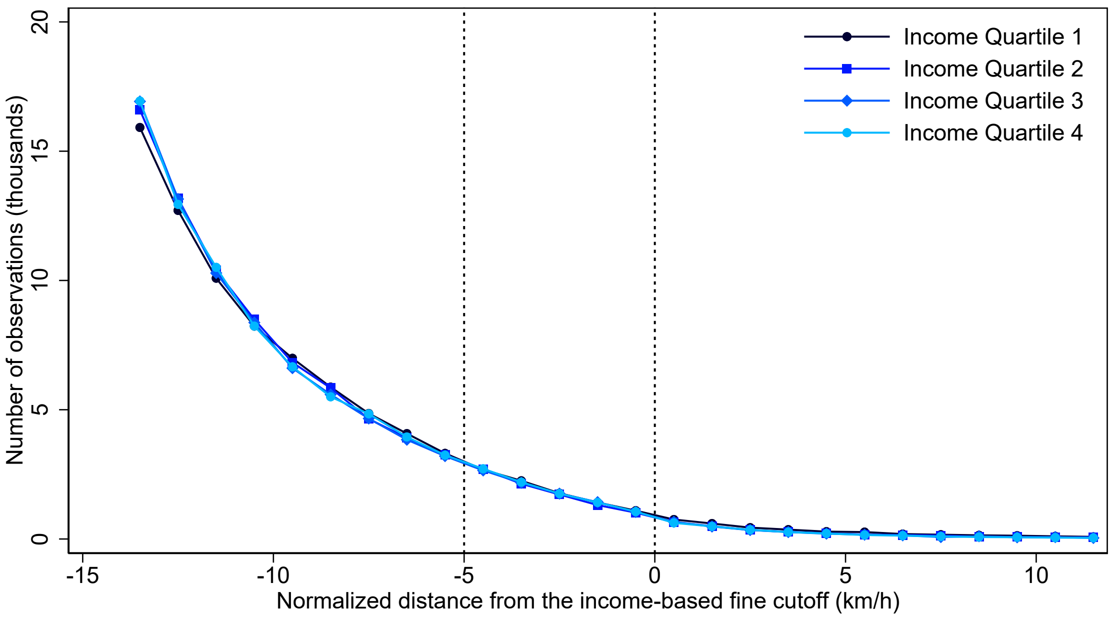

Notes: The figure presents the speeding ticket distributions in the sample containing speeding tickets assigned by cameras between Jan 2018 - May 2020. Figure Panel (a) shows the distributions in different speeding limit areas. The x-axis refers to the speed in km/h, and the y-axis to the number of speeding tickets per km/h. Vertical dashed lines indicate the points where the average speeding ticket jumps due to the income-based fine system. Figure Panel (b) pools all the speeding limit zones except 80 km/h together, normalizes the x-axis to measure the distance from the income-based fine cutoff, and plots the speeding distributions for different income quartiles. I define income quartiles using an individual's average net income 1-2 years before the speeding incident. The first vertical line highlights the point at which the fine jumps a fixed amount. The second vertical line highlights the point at which the income-based fine system kicks in.

speeding distributions is not just an artifact of the police camera data.

To shed additional light on how cameras shape behavior, Appendix Figure A.21 presents data from a location equipped with both a traffic monitoring station and a police camera. While the speed distribution peaks sharply at the speed limit, it remains smooth around the income-based fine cutoff, with no evidence of bunching.[23]

Starting from the next section, the rest of the paper tries to understand why people ignore the cutoff.

4 Does a Larger Speeding Ticket Influence Re-offending?

The previous section documents that drivers do not behave as a simple Becker model predicts. Drivers ignore the discontinuous changes in the price of speeding at the income-based fine cutoff *ex-ante* and do not bunch as this theory would predict. In other words, changes in the threat of punishment at the cutoff do not have a general deterrent effect.

There are two possible explanations for the lack of bunching. First, drivers may fail to bunch because of information frictions. In other words, people do not understand how the system works, leading to suboptimal decisions. Second, people may operate under perfect information but ignore the cutoff because of adjustment costs. To bunch, an individual has to consider many things, such as the measurement error of a car's speedometer, the downward correction made by the police, and the exact location of the income-based cutoff. Due to the complexity of the problem, individuals may feel that gains from bunching are smaller than the adjustment costs, implying that they neglect the cutoff.

Although both information frictions and adjustment costs can generate smooth speeding distributions, they make different predictions on how realized speeding tickets should affect reoffending. If information frictions characterize an individual's behavior, then assigned fines should reveal information about the punishment system and affect recidivism. However, under adjustment costs, individuals should not react to fines since they already operate with perfect information.

This section examines whether fines can still exert specific deterrence effects when general deterrence effects are absent. Specifically, this section asks whether being assigned a larger versus a smaller speeding ticket changes the driver's behavior *ex-post*. To answer this question, I use a regression discontinuity design to compare individuals just below and above the cutoff. The smoothness suggests that drivers do not manipulate whether or not they receive a large fine.

[23]Traffic monitoring stations generally lack cameras. However, I was able to identify one station that has a camera: Station 145, located on Ring Road 1, which runs through Helsinki and Espoo.

How individuals react to the size of the fine helps us understand how drivers make speeding decisions. Adjustment costs will likely explain the lack of bunching if the assigned speeding ticket size does not influence recidivism. However, if larger speeding tickets reduce reoffending, then information frictions and learning characterize the individual's behavior. In Section 5, I will present a theoretical framework that rationalizes the empirical findings presented in this and the previous section.

4.1 Empirical Specification

I investigate the impact of larger realized fines on recidivism using a sharp regression discontinuity design (RDD). My RDD equation takes the form

$$Y_{il,t} = \beta_t Z_{il,0} + f(S_{il,0}) + f(S_{il,0}) \times Z_{il,0} + \alpha_{l,0} + \epsilon_{il,t}, \tag{2}$$

where the dependent variable $Y_{il,t}$ measures the cumulative recidivism t months after the initial speeding incident. The indicator variable $Z_{il,0}$ equals 1 if an individual i crossed the income-based fine cutoff in the speeding limit region l in period zero (normalized to be the time of the speeding incident). The running variable $S_{il,0}$ controls the distance from the fine cutoff. The interaction term $f(S_{il,0}) \times Z_{il,0}$ allows the relationship between the outcome and the running variable to change at the cutoff. Further, I control for the speed limit fixed effects $\alpha_{l,0}$. The error terms ($\epsilon_{il,t}$) are clustered at the individual level. I use triangular weights centered at income-based fine cutoffs, and select optimal bandwidth around the cutoff using methods of Calonico *et al.* (2014).

The coefficient of interest, β_t, captures the specific deterrence effect of larger fines on recidivism. Interpreting this estimate as causal requires that the potential outcome evolves smoothly at the cutoff. If this identifying assumption holds, individuals at the cutoff are, on average, similar but receive speeding tickets of very different magnitudes.

A threat to identification could be that fixed and income-based fines would differentially affect the risk of license suspension or the size of future fines. However, as explained in Section 2.1, fixed fines and income-based fines accumulate points similarly. Moreover, the type of fine does not influence the size of future fines. Consequently, any discontinuity in reoffending at the cutoff can be attributed solely to variation in monetary penalties, rather than to differences in license suspension risk or the structure of future fines.[24]

Treatment Figure 5(a) clarifies the treatment for those who cross the income-based fine cutoff. Black connected dots in the figure reveal that those on the right-hand side of the cutoff (treatment group) receive a fine that is, on average, 200 euros larger than the

[24]Appendix Figure A.20 shows that about 20% of individuals received a speeding ticket in the previous year. Since suspension requires three violations in 12 months, few are at immediate risk of license suspension.

Figure 5: Graphical Evidence: Validity of the Setting

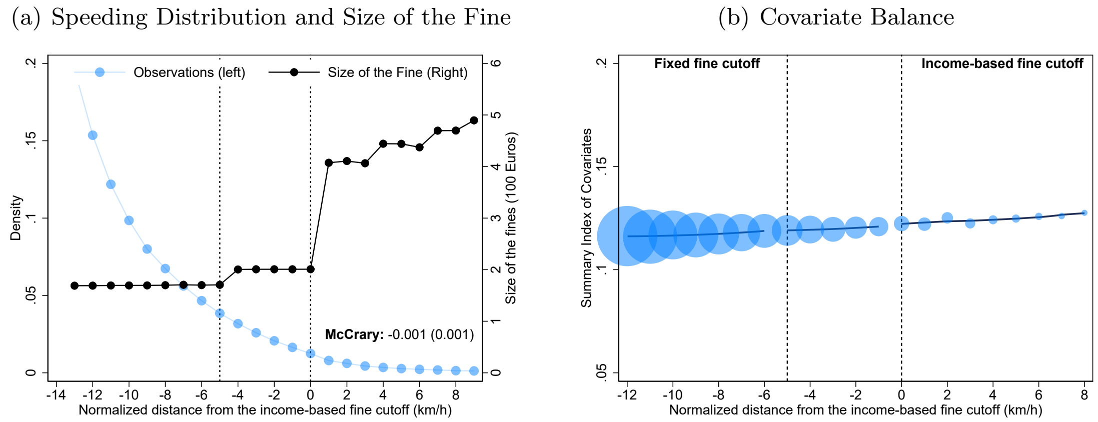

Notes: The figure presents how speeding distribution (a), size of the fine (a), and predetermined background characteristics (b) evolve at the income-based fine cutoff. The x-axis measures the distance from the income-based fine cutoff in km/h. In Figure Panel (a), the left-hand side y-axis refers to the density per km/h, and the right-hand side y-axis measures the average size of the speeding ticket. Figure Panel (b) plots a summary index of predetermined covariates against normalized speed. The index is equal to the predicted values I obtain by regressing the indicator variable, which takes value one if an individual commits another traffic offense within six months after the speeding tickets against covariates listed in table 1. I only used individuals who received a fixed fine to estimate the model. I use coefficients from this model to construct the index for everyone in my sample. The sample constructed as the section 2.3 describes.

fine for those on the left-hand side of the cutoff (control group). However, the size of the fine varies considerably with an individual's income (see, Appendix Figure A.19).

Validity of the Setting I perform two types of tests to assess whether the identifying assumption is satisfied. First, if people do not manipulate the running variable, the distribution of the running variable should be continuous. Figure 5(a) provides evidence that this first prediction holds in my setting. Based on the figure, the number of observations decreases rapidly with speed, but there is no evidence of excess mass below the cutoff. In other words, I do not find evidence that individuals manipulate the speed perfectly.

Second, if potential outcomes evolve smoothly at the cutoff, individuals just below and just above the cutoff should be similar in their observable characteristics on average. To determine whether predetermined characteristics evolve smoothly at the cutoff, I investigate whether a predicted probability of reoffending that I construct using predetermined covariates jumps at the cutoff.[25] The predicted probability is an index that summarizes all background variables succinctly into one variable. I construct the index as follows. First, I take individuals who received a fixed fine and regress a dummy variable for reoffending on a set of predetermined variables, excluding speed. Then, I use regression coefficients to calculate each individual's predicted probability of reoffending.

Figure 5(b) illustrates that the index score behaves smoothly at the cutoff, providing another piece of evidence supporting the key assumption of RDD. In addition, I study how

[25] Rose and Shem-Tov (2021) and Londoño-Vélez *et al.* (2020) conduct similar validity checks.

individual variables behave at the cutoff. Columns (3) and (4) of Table 1 show that most of the predetermined background variables behave continuously at the income-based fine cutoff. Out of 18 variables tested, only three are statistically significant at a 5 percent level, which is slightly more than one would expect to show up due to randomness. However, these significant variables do not seem to be jointly related to the predicted probability of committing traffic crimes, since the summary index of covariates evolves so smoothly in Figure 5(b). Furthermore, in the robustness section 4.5, I show that the main estimates barely move when I add these background characteristics as controls in the estimation.

4.2 The Impact of Speeding Ticket Size on Recidivism

Descriptive Evidence Figure 6 presents the first evidence of how people respond to larger fines that arise from the income-based fine system. The x-axis of the figure measures the distance from the income-based cutoff in km/h. The dots plot the probability of committing another traffic offense within six months after a speeding ticket. We observe sharp drops in the probability of reoffending at the cutoffs where the fines jump discontinuously. The drop is large at the income-based fine cutoff, but we also see some action at the fixed fine cutoff. Before and after the fine cutoffs, the probability of reoffending evolves smoothly. Under the key assumption I stated above, which seems plausible given the evidence, the discontinuous jump in the fines causes the observed drops in the probability of reoffending. Next, I use an RDD to quantify the size of the drop in reoffending at the income-based fine cutoff.[26]

RDD Estimates Figure 7 presents the second key result of the paper. The figure shows how the larger speeding tickets impact the probability of committing another traffic offense in the future. In Figure 7(a), I graph the RDD estimates, i.e. the β_ts from equation 2. These estimates capture the specific deterrence effect of a larger speeding ticket on the cumulative reoffending probability. My main outcome tracks individuals' cumulative reoffending probabilities 1-12 months after a fine.[27]

Figure 7(a) shows that the short-term cumulative reoffending probability drops when a driver receives a larger fine due to an income-based system. The impact peaks around 5-9 months after the speeding incident. At this time, drivers who received a larger fine are around 2.1-2.9 percentage points less likely to commit another traffic offense than individuals who received a fixed fine.[28]

[26] Appendix Figure A.26 shows the descriptive pattern after 3, 6, 9, and 12 months.

[27] I can follow the same sample for six months, after which the sample size starts to decrease. I show that my results are robust to this limitation in Section 4.5.

[28] I exclude the 80 km/h limit zone from my main analysis. Despite the lack of incentives to do so, I observe bunching at the speed of 96 km/h when the limit is 80 km/h. One explanation for this excess mass point is an error in the data, so I do not use the limit zone in the main analysis. However, Appendix Figure A.13 shows that I get similar results when I use all speed zones.

Figure 6: Graphical Evidence: Recidivism Within 6 Months

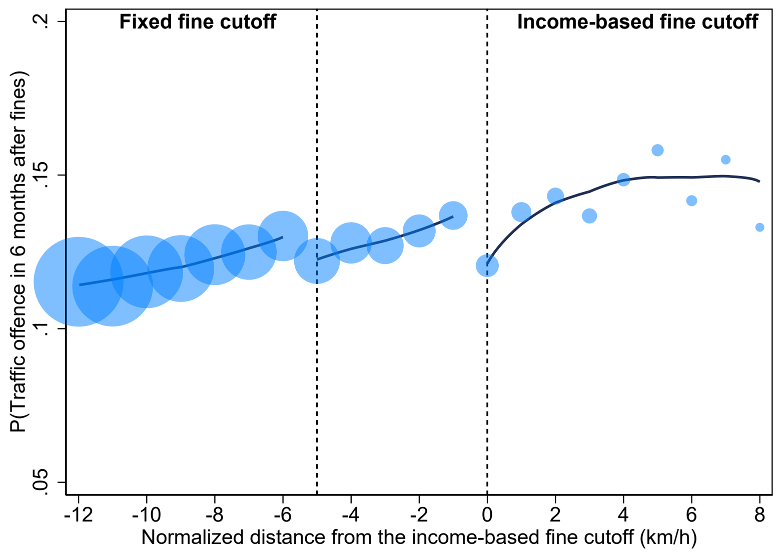

Notes: The figure presents a recidivism outcome as a function of normalized speed. The x-axis measures the distance from the income-based fine cutoff in km/h. The y-axis refers to the probability of committing another traffic offense within six months after the initial speeding ticket. Vertical lines highlight the fine discontinuity points. The income-based (fixed) fine cutoff locates at 0 (-5). The sample is constructed as the section 2.3 describes.

I also provide suggestive evidence that drivers do not react to larger realized fines by bunching below the income-based fine cutoff. Appendix Figure A.9 plots the speeding distributions for individuals reoffending within a year. The figure shows that individuals who received an income-based fine and reoffended do not bunch below the cutoff. These results suggest that a larger speeding ticket triggers only an extensive margin response.

To understand the relative magnitudes of these estimates, Figure 7(b) plots cumulative reoffending probabilities in the control group, and Figure 7(c) relates the point estimates shown in Figure 7(a) to the reoffending probabilities in the control group. Reoffending is common: of those who are just below the income-based fine cutoff, more than 20 percent commit another traffic offense within a year. However, larger income-based fines decrease recidivism in the short term. Figure 7(c) shows that those who receive an income-based fine are around 15-20 percent less likely to commit another traffic offense in the following 6-8 months than individuals who receive a smaller fixed fine.

Lastly, Figure 7 provides some suggestive evidence that the decrease in speeding as a result of a larger speeding ticket starts to fade out around nine months after the initial ticket was assigned. Although the point estimate is still negative twelve months after the initial speeding ticket, it is not statistically significant from zero at a 95% confidence level. This result suggests that if the target of larger fines is to achieve a permanent reduction in speeding, then fines must be assigned frequently to those who speed.[29]

[29] Appendix Figure A.23 replicates Figure 7, extending the follow-up period to 18 months. The point

Figure 7: RDD Estimates at the Day Fine cutoff - 80 km/h limit omitted

(a) RDD Point Estimates

(b) The Mean of Outcome in Control Group

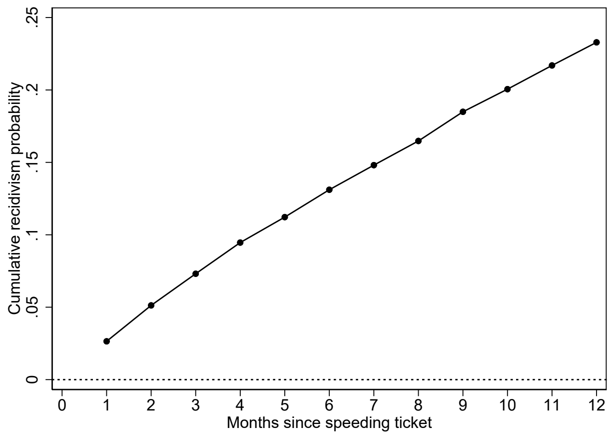

(c) The Relative Magnitude of Point Estimates

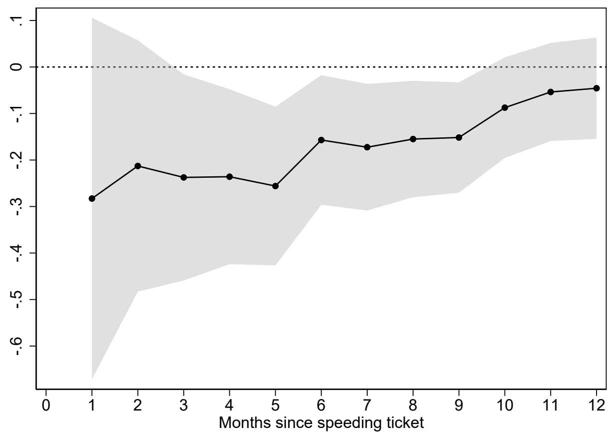

Notes: Figure panel (a) plots the RDD estimates of $\beta_{il,t}$ obtained using Equation (2). Estimates plot the impact of the higher fine on the cumulative probability of committing a traffic offense in t months after the initial fine. The estimates are obtained using conventional local linear regression with a triangular kernel. Vertical solid lines are the 95 percent confidence intervals. Standard errors are clustered at the individual level. Panel (b) shows the mean of the outcome variable in the control group, and panel (c) relates the point estimate shown by the connected dots in panel (a) to reoffending probabilities in the control group. Sample construction and data as defined in section 2.3.

Table 2: Regression Discontinuity Design Estimates

Dep. Variable: P(Traffic Offence between time 1 and t)

Time (t) (1)	**Mean** (2)	**RDD** (3)	**RDD** (4)	**RDD** (5)	**BW** (6)	**Obs.** (7)
1	0.028	-0.007	-0.008	-0.011	3	23,026
		(0.005)	(0.005)	(0.007)		
2	0.056	-0.011	-0.011	-0.015	3	23,026
		(0.007)	(0.007)	(0.009)		
3	0.078	-0.017	-0.018	-0.022	3	23,026
		(0.008)	(0.008)	(0.010)		
4	0.102	-0.022	-0.023	-0.027	3	23,026
		(0.009)	(0.009)	(0.011)		
5	0.120	-0.029	-0.030	-0.034	3	23,026
		(0.010)	(0.010)	(0.011)		
6	0.137	-0.021	-0.022	-0.025	4	31,912
		(0.009)	(0.009)	(0.011)		
7	0.155	-0.026	-0.026	-0.030	4	30,627
		(0.010)	(0.010)	(0.013)		
8	0.173	-0.026	-0.026	-0.030	4	29,323
		(0.011)	(0.010)	(0.013)		
9	0.192	-0.028	-0.030	-0.033	4	27,715
		(0.011)	(0.011)	(0.014)		
10	0.207	-0.018	-0.019	-0.020	5	34,908
		(0.011)	(0.011)	(0.014)		
11	0.224	-0.012	-0.012	-0.014	5	33,062
		(0.012)	(0.011)	(0.015)		
12	0.242	-0.011	-0.012	-0.013	5	30,592
		(0.013)	(0.013)	(0.016)		
Controls			✓			
CCT Estimator				✓		

Notes: The table shows the RDD estimates from the income-based cutoff displayed in Figure 7. The analysis omits the limit zone 80 km/h. The outcome is the probability that an individual commits another traffic offense 1 to t months after the initial speeding ticket. Column (2) shows the mean of the outcome for those just on the left-hand side of the cutoff. Column (3) presents the RDD estimates obtained using equation 2. Column (4) shows results from a similar analysis but with controls. Column (5) reports biased corrected RDD estimates obtained using the robust approach by Calonico *et al.* (2014). Column (6) shows the bandwidth selected using methods of Calonico *et al.* (2014). Column (7) shows the number of observations. In columns (3) and (4), standard errors in the parentheses are clustered at the individual level. For bias-corrected estimates, the confidence intervals are given by the approach of Calonico *et al.* (2014). Sample construction and data as defined in Section 2.3.

The potential fade-out appears even more pronounced in Figure 7(c), which illustrates the relative magnitude of the estimates. However, this pattern should be interpreted with caution. One possible explanation is that crossing the income-based fine threshold results in an immediate reduction in the treatment group's probability of reoffending compared to the control group. Thereafter, both groups accumulate offenses at a similar rate. This implies that the RDD estimate could remain stable, although its magnitude declines relative to the control group's mean over time.

4.3 Does the Effect Vary by Income?

The results above indicate that larger speeding tickets, triggered by crossing the income-based fine cutoff, lead to a reduction in reoffending. However, average effects may mask important variation in the impacts of fines as the fine discontinuity varies substantially with income. This raises two key questions: (1) Do high- and low-income drivers exhibit similar behavioral responses to their respective fine increases? (2) Does the marginal deterrence effect of fines differ by income? These questions are central to evaluating whether income-based fines are both equitable and effective.

I address these questions in two steps. First, I examine the threshold-crossing effects by income. Second, I estimate how the marginal effects of fines vary by income using an instrumental variable design. I start by splitting my sample into quartiles using individuals' annual net income two years before the ticketing year. Then, I estimate the effect separately in each quartile group, first using equation 2, and then using the IV approach.

Jump in the Fine by Income Figure 8(a) illustrates how the size of the fine develops as a function of speed for different income quartile groups. We see that different income quartiles experience very dissimilar speeding ticket hikes at the income-based fine cutoff. The fine increases by around 500 euros for the top quartile, whereas the hike equals around 50 euros for the bottom quartile. For the bottom group (top group), the fixed fine equals approximately 20 (4) percent of the net monthly income. In relative terms, the income-based fine is similar across all income groups.

Descriptive Evidence by Income Figure 9 plots reoffending probabilities around the fine thresholds, disaggregated by income group. Each blue dot represents the share of drivers who reoffend within six months of receiving a speeding ticket, plotted against driving speed, separately by income quartile. Panel (d) shows a sharp decline in reoffending at the income-

estimates appear to converge toward zero by the end of the period, although they become increasingly imprecise over time.

Figure 8: Impact of Fines and Income Quartile

(a) Average Size of Fine

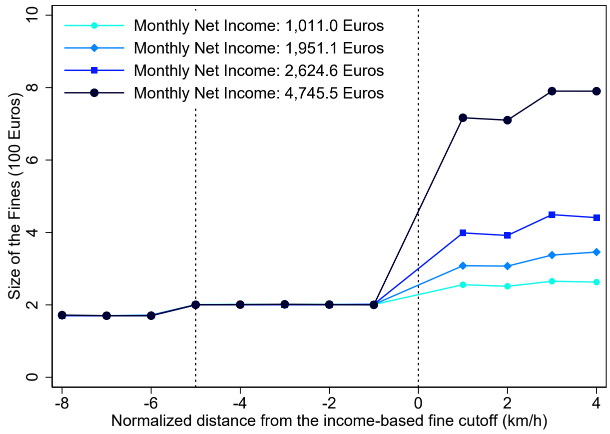

(b) RDD Estimates by Income (Income-Based Cutoff)

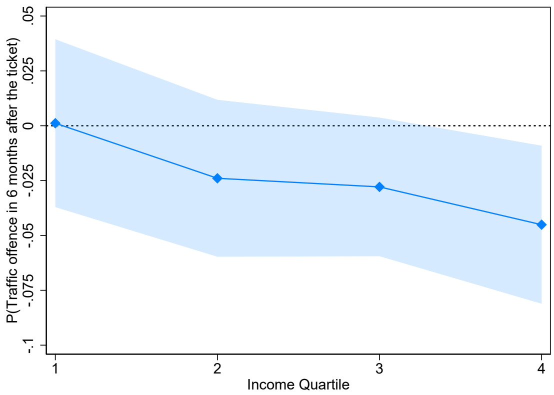

(c) RDD Estimates by Income (Fixed Fine Cutoff)

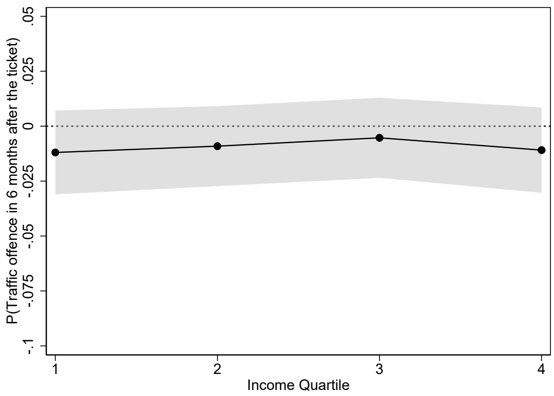

Notes: Figure Panel (a) shows a jump in the fine's size at the income-based fine cutoff for different income quartiles. The income quartiles are constructed using an individual's annual net income measured 1-2 years before the speeding incident. The legend in the left upper corner shows the approximated monthly income obtained by dividing the individual's annual net income by twelve. The x-axis measures the distance from the income-based fine cutoff. Figure Panel (b) reports RDD estimates for different income quartiles at the income-based fine cutoff. Figure Panel (c) restricts the sample to a fixed fine cutoff. The outcome is the probability that an individual commits another traffic crime within six months after the original speeding incident. The estimates are obtained using equation 2. The shaded bands show 95 percent confidence intervals. Standard errors are clustered at the individual level. Sample construction defined in 2.3.

based threshold for the top income group, which faces the steepest fine increase. In contrast, lower-income groups exhibit a smaller decline in reoffending.

Figure 9 also clarifies why Figure 6 suggests that only individuals exactly at the cutoff react to the fine's size. Figure 9 shows that the entire reoffending curve to the right of the cutoff shifts downward for the top income group. However, the strong heterogeneity by income causes the averaged pattern in Figure 6 to suggest that only individuals exactly at the cutoff respond to larger fines.[30]

Figure 9: Normalized Speed and Reoffending by Income Quartile

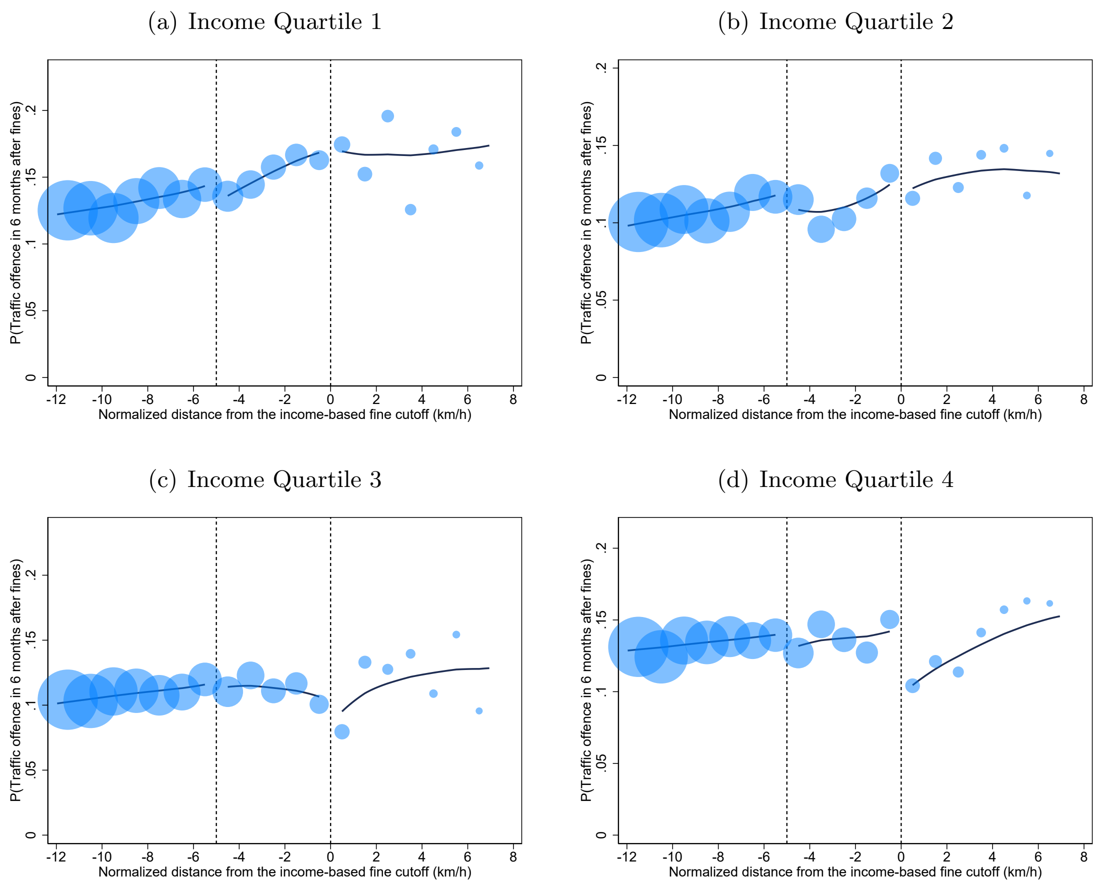

Notes: The figure presents a recidivism outcome as a function of normalized speed for different income quartile groups. The x-axis measures the distance from the income-based fine cutoff in km/h. Income quartiles are constructed using the mean disposable income 1-2 years before the speeding incident. The y-axis refers to the probability of committing another traffic offense within six months after the initial speeding ticket. Vertical lines highlight the fine discontinuity points. The income-based (fixed) fine cutoff locates at 0 (-5). The sample is constructed as the section 2.3 describes.

Threshold-crossing Effects by Income Figures 8(b) and 8(c) then plot RDD estimates for different income quartiles when the outcome is the cumulative reoffending probability six months after the initial speeding ticket. Blue diamonds in Figure 8(b) show the RDD

[30] A nonlinear relationship between measured speed and reoffending can also explain why the effect in Figure 6 appears localized. If this nonlinearity is strong around the cutoff, it may seem that only individuals at the threshold are affected, even though the treatment impacts all those above it. Appendix Figure A.24 illustrates this with two counterfactual scenarios based on the RDD estimates: in Panel (a), the treatment effect is constant; in Panel (b), it diminishes with observed speed.

estimates for different income quartiles at the income-based fine cutoff. Although the estimates are relatively imprecise, they suggest that high-income individuals react very strongly to their fines. The point estimate is zero for the bottom quartile, after which the estimates increase almost linearly. For the top income group, the RDD estimate (-0.045) is double the size of the main estimate I presented in the previous section.

The black dots in the Figure 8(c) present similar RDD estimates as above but use variation from the fixed fine cutoff where an individual's income does not affect the size of the fine.[31] Even though the estimates are imprecise, the magnitudes are smaller and more comparable between groups.

Together, these estimates suggest that individuals who receive larger fines are less likely to recidivate than individuals who receive small fines. Because all income groups react similarly to fixed fine hikes, the sizable reaction to income-based fines by high-income people is due to larger fines and not because they are different on some unobserved dimension that impacts speeding behavior and is correlated with income.[32]

Instrumental Variable Estimates The RDD estimates across income groups are challenging to compare as the fine size varies with income. Crossing the threshold increases fines by €50 for low-income drivers and by €500 for high-income drivers. As a result, treatment effects are not directly comparable across groups. To address this, I treat crossing the income-based cutoff as an instrumental variable for the realized fine amount. This approach allows me to estimate the marginal effect of a fine on recidivism.

In the instrumental variable (IV) analysis, my first stage equation is

$$F_{il,0} = \pi_0 Z_{il,0} + f(S_{il,0}) + f(S_{il,0}) \times Z_{il,0} + \alpha_{l,0} + \epsilon_{il,0}, \tag{3}$$

where $F_{il,0}$ stands for the fine size, $S_{il,0}$ controls for distance from the income-based cutoff, $Z_{il,0}$ takes value 1 if an individual crosses the income-based cutoff, $\alpha_{l,0}$ are limit fixed effects, and $\epsilon_{il,0}$ is the error term. The first stage coefficient π_0 identifies the average increase in fine size when crossing the income-based cutoff.

Reduced form regressions are equivalent to the main specification I present in equation 2. These estimates capture the effect of crossing the income-based cutoff on reoffending. IV estimates equal the ratio between the reduced form coefficients and the first stage coef-

[31] Appendix Figure A.14 conducts the same analysis as the main Figure 7 but using variation from the fixed fine cutoff (15 km/h), where the fine is independent of income. The jump in the fixed fine has a small and imprecise effect on reoffending.

[32] Appendix Table B.12 examines how the estimates vary with recent income changes. The results suggest that individuals with recent positive income growth exhibit a stronger response to the size of the realized fine than those with negative income growth. This pattern is consistent with individuals failing to adjust their expectations about fine amounts as their income increases prior to being fined.

ficients. IV estimates capture how reoffending changes in response to a one-euro increase in fines.[33]

The IV estimates have a causal interpretation under three assumptions. First, the instrument must be strongly correlated with the treatment. Figure 5(a) shows that this assumption holds. Second, the instrument must be independent of individuals' potential outcomes. This condition follows from the key RDD assumption that potential outcomes evolve smoothly at the cutoff. Figure 5(b) provides supporting evidence that this assumption is satisfied in my setting. Finally, the exclusion restriction requires that the instrument can affect the outcome only through the treatment. This assumption would be violated, for example, if crossing the cutoff would trigger extra license suspension points. However, as explained in Section 4.1, this is not the case.

Table 3 reports the results from the instrumental variable (IV) analysis. Column (1) reports results for the full sample, while columns (2)–(5) present results by income quartile, defined based on predetermined net income. Panels A and B of Table 3 show the first stage and reduced form estimates for the full sample and for each income quartile. Panel C presents the IV estimates. For ease of interpretation, all IV coefficients are scaled by a factor of 100. Column (1) indicates that a € 100 increase in the fine reduces reoffending by approximately one percentage point.

Columns (2)–(5) of Table 3 report IV estimates separately by income quartile. The IV estimate for the top income quartile (0.009), reported in column (5), is slightly smaller than the full-sample estimate (0.011). In contrast, columns (3) and (4) show the largest marginal effects for the two middle-income quartiles, although these estimates are imprecise. This pattern also appears in the elasticities reported at the bottom of the table, which I calculate as the ratio of the relative change in reoffending to the relative change in fine size.

In summary, the threshold-crossing effects and IV estimates offer important insights into how income-based fines influence reoffending behavior. The reduced form effects indicate that high-income individuals, who face the largest fines, reduce reoffending the most. Yet, while high-income individuals exhibit the strongest response, the IV estimates suggest that the marginal deterrence effects are largest for the middle-income groups, although the estimates are imprecise. Taken together, these findings suggest that income-based fines may help reduce income-related disparities in speeding behavior, which I document in Figure 2. However, because marginal deterrent effects decline with income, fines for high-income individuals must be sufficiently large to generate a comparable reduction in reoffending.

[33] I estimate the impact of fines on reoffending using the two-stage least squares (2SLS) approach that gives the correct standard errors. In my setting, 2SLS estimates will equal the ratio between reduced form and first stage coefficients.

Table 3: Instrumental Variable Analysis

		Income Quartile			
	All	Q1	Q2	Q3	Q4
	(1)	(2)	(3)	(4)	(5)
Panel A: First Stage (Fines)					
	204.113	51.582	103.299	187.268	499.060
	(7.427)	(1.790)	(2.850)	(4.807)	(28.539)
Panel B: Reduced Form (Reoffending within 6 months)					
	-0.022	0.001	-0.024	-0.028	-0.045
	(0.009)	(0.020)	(0.018)	(0.016)	(0.018)
Panel C: IV (Reoffending within 6 months)					
	-0.011	0.002	-0.023	-0.015	-0.009
	(0.005)	(0.038)	(0.018)	(0.009)	(0.004)
Controls	✓	✓	✓	✓	✓
Control mean	0.137	0.163	0.132	0.101	0.150
Elasticity	-0.157	0.024	-0.352	-0.296	-0.120
Bandwidth	4	4	4	4	4
Observations	31,912	8,418	7,708	7,887	7,899

Notes: The table reports first stage, reduced form, and instrumental variable (IV) estimates for the full sample and by income quartiles. Panel A presents first stage results from equation 3, where the outcome is the fine amount. Panel B reports reduced form coefficients from equation 2, where the outcome is an indicator for committing any traffic offense within six months after a speeding ticket. Panel C presents IV estimates using 2SLS with the same outcome as in Panel B. Column (1) shows estimates for the full sample. Columns (2)–(5) report results by income quartiles, defined using average net income over the two years before the speeding incident. Elasticities are calculated as the ratio of: (1) the reduced form coefficient divided by the mean reoffending rate in the control group, and (2) the first stage coefficient divided by the mean fine amount in the control group. IV estimates are multiplied by 100 for ease of interpretation. All regressions include controls to improve statistical precision. Standard errors are clustered at the individual level. See Section 2.3 for details on sample construction.

4.4 Additional Results

Appendix Figure A.22 illustrates how RDD results vary by prior speeding history. In the Bayesian learning framework, individuals who have received more speeding tickets should respond less to new information, as their priors are more precise. To test this, I estimate the RDD separately for reoffenders (those who received at least one speeding ticket in the preceding four years) and first-time offenders. Contrary to theoretical predictions, Appendix Figure A.22(b) shows that reoffenders' response to the fine hike is approximately three times larger than that of first-time offenders.

However, the results should be interpreted with caution. As Appendix Table B.11 shows, reoffenders and first-time offenders differ along several dimensions that may influence their responsiveness to fines. Reflecting this, Appendix Figure A.22(c) shows that reoffenders in the control group are substantially more likely to reoffend than first-time offenders. Nonetheless, these results remain informative, as they suggest that specific deterrence effects are strongest among individuals most likely to engage in behavior that poses risks to others.

Another important question is whether specific deterrence effects vary with monitoring intensity. When enforcement is absent, the fine size should have little impact, as individuals face a negligible risk of detection. Conversely, when the probability of being caught approaches one, most drivers are already deterred, making the fine amount less relevant. This reasoning suggests that fines should exert the strongest deterrence effect at intermediate levels of monitoring.

To test this prediction, I examine how recidivism varies with enforcement intensity in individuals' home municipalities. I proxy enforcement intensity by the number of speeding tickets issued per car-owning resident, averaged over 2018–2019. Appendix Figure A.28 illustrates substantial variation in monitoring intensity across municipalities.[34] I then divide the sample into tertiles based on this measure.

Appendix Table B.10 presents estimates of specific deterrence effects by monitoring intensity tertiles. Although imprecise, the point estimates are consistent with the hypothesis that specific deterrence effects peak when monitoring intensity is intermediate. Columns (3)–(4) suggest that fine size has a substantially larger impact on recidivism in intermediate-monitoring regions than in either low- or high-monitoring areas. Importantly, Panel A shows that this pattern is not driven by differences in fine size across regions.

Finally, I examine whether the magnitude of the realized fine affects the probability of a traffic accident using the RDD approach. My outcome variable is an indicator equal to one if an individual was involved in a traffic accident as a driver within six months of receiving a fine. Appendix Table B.13 presents the results. I find no evidence that the magnitude of the fine affects the probability of an accident. One caveat is that traffic accidents are rare, as reflected by the low control group mean in Table B.13. This may limit statistical power to detect anything but very large effects.

4.5 Robustness and Validity Tests

This section shows that my results do not arise from arbitrary RDD specification choices. First, one concern is that the selected specification approximates the unknown conditional expectation function poorly. To alleviate these concerns, Appendix Table B.8 shows that I obtain similar results when I use a specification with a second-order polynomial or conduct the analysis using the robust approach proposed by Calonico *et al.* (2014). Further, Appendix Figure A.18 shows "honest" Armstrong and Kolesár (2020) confidence intervals that incorporate the potential bias in the estimates.

[34] Unfortunately, I do not observe the number of speed cameras directly and must rely on this proxy.

Second, Appendix Figure A.10 demonstrates that the point estimates are insensitive to the bandwidth choice. Although I tie my hands by selecting the optimal estimation window using the methods of Calonico *et al.* (2014), Appendix Figure A.10 adds additional confidence that an arbitrary choice of window does not drive the results.

I also carry out a set of additional validity checks. First, Appendix Figure A.12 and Table 2 demonstrate that the RDD estimates are insensitive to the control variables. Second, Appendix Figure A.11 shows results from a placebo test where I randomly reallocate the cutoff to places where fines do not jump and carry out RDD analysis. I do not find statistically significant effects using placebo cutoffs.

Additionally, I test whether the effects potentially differ for the general population not included in my RDD sample. Individuals near the cutoff may be less informed about the enforcement regime than those driving at lower speeds or those with no prior infractions. If so, they may respond more strongly to the fine amount than the broader population.

To explore this possibility, I reweight the estimation sample to reflect the population of car owners who do not receive speeding tickets. I begin with all car owners in the register data. Then, using a logit model and an individual's background characteristics, I estimate the probability that an individual does not appear in the speeding sample. Finally, I re-estimate the RDD using the predicted probabilities as weights.[35] Column 5 of Appendix Table B.8 presents the reweighted RDD results. The standard and reweighted estimates are nearly identical. This suggests that, based on observed background characteristics, individuals at the cutoff and in the broader population likely respond similarly to the fine.

One limitation is that I cannot directly test for avoidance behavior, such as drivers altering routes to evade speed cameras. A natural test would examine whether drivers avoid specific cameras that have previously issued fines for them. Unfortunately, I cannot implement this test because my data lack camera-level identifiers. However, the fact that my identification strategy exploits intensive-margin variation, where both groups face the same surveillance environment and have equal exposure to camera location information, mitigates concerns that the results are driven by avoidance.

Finally, I conduct an analysis where my outcome is the cumulative probability of committing a non-traffic crime. Since my main sample contains individuals who rarely commit non-traffic crimes, it is unlikely that the size of the speeding ticket would impact non-traffic crimes. Indeed, I find that the RDD point estimates are very small and not statistically significant when my outcome is the cumulative probability of non-traffic criminality. Appendix Figure A.15 reports these results.

[35] Appendix Table B.9 lists the background characteristics used in the regression. Appendix Figure A.27 shows the predicted probabilities from the model.

5 What Explains the Smooth Distributions and Reactions to a Larger Fine?

So far, I have presented four empirical findings. First, I demonstrate that despite significant incentives, individuals do not bunch just below the income-based fine cutoff. Second, my findings indicate that realized fine size reduces recidivism in the short term. Third, the larger the realized fine, the larger the ex-post reaction by drivers. Specifically, higher-income drivers who face the largest jumps in assessed fines experience the largest drop in recidivism. Fourth, the effect on reoffending fades out over time.

In this section, I present a simple theoretical framework that seeks to rationalize these puzzling results, building on previous models in related literature (Traxler *et al.*, 2018). My conceptual framework incorporates two types of optimization frictions that could reconcile the main results.[36] First, I demonstrate that adjustment costs may attenuate drivers' response to fine hikes and result in little or no bunching below the cutoff, which indicates a discontinuous increase in the assessed speeding fine. If adjustment costs are high enough, the distribution may be completely smooth. However, under a model with adjustment costs, individuals should not react to realized fines. The intuition for this result is that drivers already know precisely how the system works, but because it is too costly to monitor speed perfectly, they choose not to do so. As a result, a realized fine does not bring any new information that could change his/her behavior next period.

Second, I present a model with information frictions that can generate the observed patterns in the data. The model builds on the idea that individuals use the so-called "ironing heuristic" and approximate marginal prices with average prices if the cognitive cost of perceiving complex pricing schemes is considerable (Rees-Jones and Taubinsky, 2019). In the model, drivers replace a complex true penalty scheme with a rule of thumb that the magnitude of the fine increases linearly with speed. Because of this heuristic, drivers ignore the jump in the marginal price of speeding at the cutoff, leading to a smooth speed distribution, as I observe in the data. However, when drivers are hit with a fine, they may change their behavior in the future. If an individual receives a larger fine than expected based on the linear rule, she concludes that the relationship between the size of the ticket and speed must be steeper than she expected. She will use this steeper penalty function in the next period and drive more slowly. Finally, based on the idea of Nyarko (1991), I show that there are drivers whose beliefs about the slope of the penalty function will never converge. This may explain the fade-out in effect.

[36] These frictions have been explored in related literatures (Chetty *et al.*, 2011; Chetty, 2012; Kostøl and Myhre, 2021) to explain attenuated responses to incentive changes.

5.1 A Stylized Model of Speeding with Adjustment Costs

Individuals choose how fast over the speed limit to go $x \in [0, X]$. They value speed via the function $u(x, \theta) = \theta x - x^2$ that is strictly concave in x and increasing in type θ. Individuals' taste for speed has a continuous and smooth distribution $G(\theta)$. With probability p, the driver is caught speeding and receives a fine, which is a step function where

$$f(x) = \begin{cases} 0 \text{ if } x = 0 \\ f^l \text{ if } 0 < x \leq x^h \\ f^h \text{ if } x > x^h. \end{cases}$$

Finally, optimization frictions ψ make adjusting speed costly.[37] Since my empirical part focuses on the intensive margin variation, I also ignore the extensive margin in the model.

Drivers' optimization problem takes the following form

$$\max_x EU = u(x, \theta) - p(f^l + f^h \cdot \mathbb{1}[x > x^h]) - \psi \mathbb{1}[x \neq x^*].$$

The last term of the problem implies that individuals pay an adjustment cost ψ if they do not locate at their interior optimum x^* given by the first-order condition.

Frictionless Model To illustrate the impact of fines and optimization frictions on optimal speed, let us first consider the problem without them. First, without fines, drivers locate at their interior optimum

$$x^* \quad \text{such that} \quad \partial u(x, \theta)/\partial x = 0.$$

Thus, speeding distribution $x^*(\theta)$ will be smooth due to the smooth type distribution $G(\theta)$.

Next, let us introduce speeding tickets. When fines are imposed, a continuum of individuals who value speeding enough such that their interior optimum x^* is higher than the fine notch just above x^h, will slow down and bunch to the corner solution where speed is

$$x^c = x^h \quad \text{such that} \quad \partial u(x, \theta)/\partial x \neq 0$$

Because of bunching, we should see excess mass at x^h. However, some drivers do not react: individuals whose interior optimum speed is lower than the speed limit or whose taste for speed is high enough remain in their interior optimum.

Adjustment Costs Finally, adjustment costs may attenuate the driver's response to fines

[37] As demonstrated by Chetty (2012), Kleven and Waseem (2013) and Kostøl and Myhre (2021), optimization frictions may attenuate individuals' reactions to discontinuous changes in marginal or average prices. In my case, the adjustment costs may arise, for example, from the fact that if individuals cannot drive at the speed defined by their type, they have to monitor their speed constantly, which decreases utility.

Figure 10: Expected Utility, Speed, and Speeding Distribution

(a) Expected utility and Speed (b) Speeding Distribution

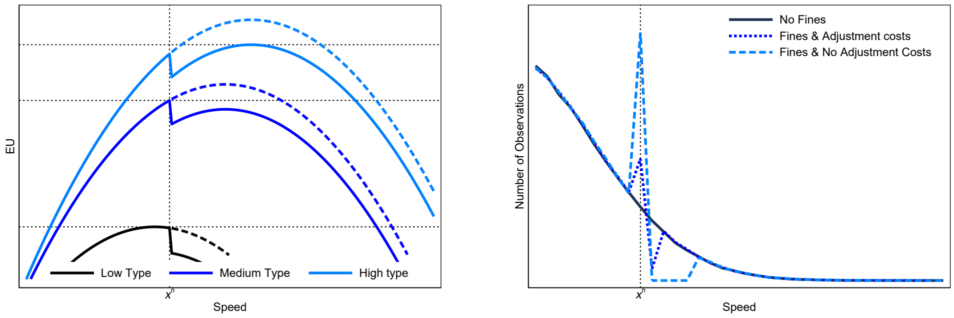

and decrease the mass at the corner solution. Under adjustment costs, the driving speed is given by

$$x = \begin{cases} x^c, & \text{if } [u(x^c, \theta) - p \cdot f^l] - [u(x^*, \theta) - p \cdot (f^l + f^h)] \geq \psi \\ x^*, & \text{if } [u(x^c, \theta) - p \cdot f^l] - [u(x^*, \theta) - p \cdot (f^l + f^h)] < \psi. \end{cases} \tag{4}$$

The equation 4 demonstrates that only individuals whose gain from reallocating from the interior solution to the corner solution is larger than the adjustment cost change their behavior. In other words, fewer individuals react to fines when adjustment costs are present, implying that we should observe less mass at the fine cutoff compared to a frictionless world.

Figure 10 illustrates the results stated above. Figure 10(a) shows the expected utility function for three different types. The dashed (solid) line shows the utility from speed without (with) fines. The medium-type individual may increase utility by moving to the corner solutions x^c at 80 km/h. However, low- and high-type individuals do not have an incentive to react to speeding tickets. Low-type's interior optimum is less than the speed limit, whereas the expected punishment does not deter the high-type enough.

Figure 10(b) demonstrates the impact of speeding tickets on speed distribution. Solid blue lines show the distribution without tickets. Black dots illustrate that the notch in expected punishment moves some people to the corner solution, creating excess mass just below the speeding limit and a hole above it. Finally, the dashed blue line plots the distribution when adjustment costs impact drivers' behavior. The mass concentration below the cutoff is smaller with the adjustment costs.

The model predicts that we should see sharp bunching below when frictions do not exist. However, adjustment costs attenuate individuals' responses and bunching. In extreme cases, the speed distribution may be smooth if individuals consider the expected punishment small compared to adjustment costs.

5.2 A Model with Misspecification and Learning

Next, I present an alternative version of the model in which individuals make speed decisions using a misspecified model and learn from their past choices. Extensive literature shows that individuals do not necessarily adjust their behavior to marginal price changes as a simple economic theory would suggest. One explanation for the muted responses is that people use the "ironing" heuristic when facing complex non-linear price schedules. If people find the mental cost of understanding a non-linear price scheme to be high, they replace the complex version with a linear approximation. In the case of speeding, the "ironing" heuristic could arise from the fact that people do not remember the exact details of the system, or they are just concentrating on many things while driving. (Rees-Jones and Taubinsky, 2019)

Motivated by literature suggesting that people respond to average prices, I assume that people replace the complex non-linear penalty scheme with linear approximation. They start with a prior about the slope of the penalty function that they use to solve for optimal speed. Since the misspecified penalty function is linear, they cannot end up with a corner solution, ruling out bunching. The optimal speed determines the signal the individuals receive. After observing the signal, individuals construct a posterior, which works as the next period's prior. Individuals learn myopically, meaning that they do not experiment.

The optimization problem with misspecification and learning takes the form

$$\max_{x_t} EU_t = u(x_t, \theta) - p \cdot \hat{f}_t(x),$$

where $\hat{f}_t(x) = \beta x_t + f^l \mathbb{1}[x > 0]$ stands for individuals' perceived penalty function at time t. The slope β is an unknown random variable and according to an individual's prior beliefs $\beta \sim \mathcal{N}[\mu_\beta, \tau_\beta^{-1}]$. The known constant equals the low fine f^l.

Each period, the individual chooses a speed such that $\partial u_t(x)/\partial x - p \cdot \partial \hat{f}_t(x)/\partial x = 0$. If they are caught speeding, they receive a fine that they interpret as a signal. The signal takes form $s_t = \beta_t^f + \eta_t, \eta \sim \mathcal{N}[\mu_s, \tau_s^{-1}]$. If speed is larger than x^h, the individual receives a large fine (β_t^h). Otherwise, β_t^f is β_t^l. After seeing the signal, the individual forms a posterior for β using Bayes' rule. Since the prior and signal are normal random variables, the posterior is a weighted average of prior and signal: $\hat{\beta} = \frac{\tau_\beta \mu_\beta + \tau_s s}{\tau_\beta + \tau_s}$ (Baley and Veldkamp, 2021). The relative precision of the prior and signal determines the weights.

The model with misspecification generates the following results related closely to my empirical exercise.

Proposition 1. *Drivers do not bunch below the fine cutoff at time $t = 1$, generating a smooth speeding distribution.*

For a proof, see Appendix Section C.

The proposition follows directly from the fact that drivers ignore the discontinuity because they use a misspecified model. In other words, all the drivers locate at their interior optimum where $\partial u(x, \theta)/\partial x = p \cdot \hat{\beta}$, leading to zero excess mass at the fine cutoff.

Proposition 2. *Under the assumptions of the model, a driver with type θ would drive slower at time $t = 2$ if she were located just above the cutoff x^h versus below it at time $t = 1$. Furthermore, the size of the reaction increases with the size of the discontinuity.*

For a proof, see Appendix Section C.

This result means that if a person with type θ drives faster than the cutoff x^h, she will receive a higher signal compared to a situation in which her speed is lower than x^h. As a result, the perceived slope of the penalty function is steeper in the next period if the individual crosses the threshold. Thus, in the next period, she will slow down. The larger the discontinuity, the steeper the updated slope, and the larger the individual's reaction.

Proposition 3. *There exists a continuum of individuals $[\theta^l, \theta^h]$ whose beliefs and actions never converge but oscillate around the speeding ticket discontinuity.*

For a proof, see Appendix Section C.

Figure 11 illustrates some of the results stated above. Consider an individual at the time t who holds a prior represented by a dark blue line in Figure 11(a). Using this prior, the individual chooses a speed such that $\partial u(x, \theta)/\partial x = p \cdot \hat{\beta}$. Assume that the speed is x^*, and hence, she receives a low signal. Let us also suppose that the prior and signal are equally precise, implying that the posterior is just an average between the signal and the prior. The dashed light blue line represents the posterior in Figure 11(a).

Next, consider another individual who happens to choose a speed equal to $80+\epsilon$ where ϵ is some tiny number. Due to this minimal difference, the second individual receives a larger signal, shown by a solid red line. The dashed red line plots the posterior for the second individual. Next, individuals choose new speeds using posteriors as their priors. The individual who received the higher signal due to discontinuity will drive at a slower speed than the individual who received the lower signal.

Interestingly, if individuals make decisions using a misspecified model and learn from their actions, there will be some individuals whose beliefs and actions do not converge but follow cycles (Nyarko, 1991). Figures 11(b) and 11(c) illustrate the result and show how the individual's beliefs and speed evolve. The individual starts with the belief that the slope of the penalty is low, chooses a speed over 80 km/h, receives a large fine, and forms a posterior higher than the prior. The individual chooses a lower speed and receives a lower signal in the next period. Over time, individual's beliefs decrease back to a low starting level. However, when the individual again thinks the slope is low, she decides to speed again, leading to a new cycle.

Figure 11: The Evolution of Speed and Beliefs in the Model with Misspecification and Learning

(a) Prior, Signal, and Posterior

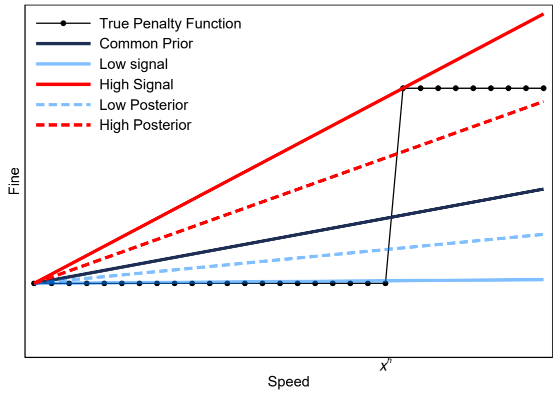

(b) The Evolution of Beliefs

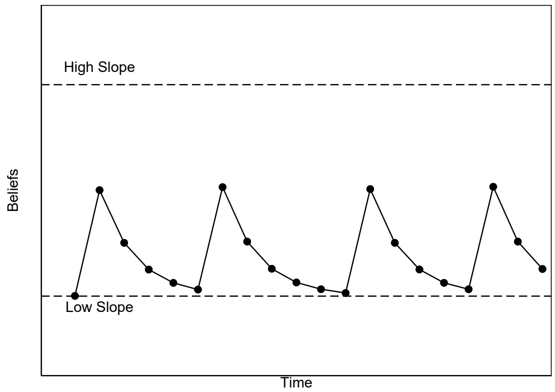

(c) The Evolution of Speed

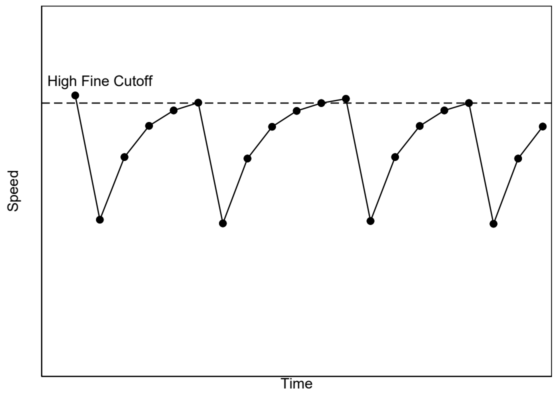

5.3 Predictions of the Models

Table 4 summarizes the predictions from both models. Both models can generate the zero-bunching result. However, only the model with misspecification and learning predicts that a larger speeding ticket (vs. smaller) reduces recidivism. Furthermore, under the misspecification and learning, we may see a fade-out in effect due to cycling beliefs.

The evidence presented in Section 4.2 supports the story that drivers make decisions using a misspecified model, leading to a smooth speeding distribution. In addition to continuous speeding distributions, I find that the size of the speeding ticket reduces the probability of reoffending, but the effect fades out over time. The model with misspecification and learning can explain both of these results.

Table 4: Predictions from the Models

	Model with:	
Prediction	**Adjustment Costs**	**Misspecification**
Bunching	Possibly	No
Reactions to larger fine	No	Yes
Cycle	No	Possibly

The table lists the predictions from the models presented in the Section 5.

5.4 Alternative Interpretations

There are other possible interpretations of the findings, particularly regarding the potential fade-out. Specifically, some of the findings are consistent with associative recall or limited memory models. Associative recall refers to the tendency to retrieve past experiences that resemble the current context when making decisions. In the model by Bordalo *et al.* (2020), individuals make decisions based on price and quality norms shaped by cues and memories that combine hedonic and contextual attributes. When facing new decisions, individuals give greater weight to memories whose cues closely match the current context.

In my setting, fixed fine and income-based fine groups face similar contexts, but the latter receives a larger fine. This makes the large fine the salient cue, leaving a strong memory trace. When deciding whether to speed again, that high-cost episode is more likely to be recalled, leading to greater deterrence. However, this effect may fade through social learning. Each penalty-free trip or encounter with someone who received a modest fine reinforces low-penalty signals. Since fixed fines are more common, drivers are likely to encounter these cues more frequently, diluting the original memory and reducing deterrence.

A related interpretation draws on the literature on selective recall. Zimmermann (2020) shows that negative and positive feedback shift beliefs in the short term, but only positive feedback has lasting effects. This aligns with models by Bénabou and Tirole (2002; 2004), in which individuals selectively forget information that contradicts their self-image. In a speeding context, a large income-based fine may conflict with a driver's belief that they are responsible or law-abiding. To preserve that self-image, the driver may suppress the memory of the fine, weakening its long-term deterrent effect. In contrast, a lower fixed fine may be seen as less threatening and easier to retain.

While associative or selective recall can explain both the initial response and the fade-out, they cannot account for the absence of bunching in the speeding distribution. To explain that pattern, the model must incorporate misperception. Once misperception is introduced, associative recall becomes unnecessary to explain the observed dynamics. In that sense, the misperception model offers a more parsimonious explanation for the findings.

I also provide evidence that the fade-out is unlikely to result from nonlinearities in the control group's behavior or asymmetric exposure to income-based fines. Figure 7(b) shows that the control group's outcome increases linearly over time, leaving no evidence of behavioral shifts that could explain convergence. Second, Appendix Figure A.25 shows that crossing the income-based fine cutoff does not raise the likelihood of later receiving such a fine, ruling out this channel as a driver of the decay.

6 Distributional Implications of Income-Based Versus Fixed Fines at the Cutoff

What would be the distributional consequences of moving from an income-based fine increase to a fixed fine increase that holds constant the aggregate reduction in reoffending at the 20 km/h cutoff? To answer this question, this section uses the heterogeneous IV estimates from Section 4.3 to quantify the income-invariant fine increase required to match the aggregate reduction in reoffending induced by the current income-based discontinuity. The implied safety-neutral fixed fine increase is €215. Although this is close to the average observed increase (€204), the distributional implications are pronounced: relative to the current income-based schedule, the fine increase would be more than 300% larger for the bottom income quartile and roughly 50% smaller for the top quartile. Safety-neutral fines would also shift reoffending toward higher-income individuals.[38]

A Safety-Neutral Fixed Fine I calculate the safety-neutral fixed fine increase as follows. Let g index income quartiles. Define s_g as the share of observations in quartile g at the

[38]My exercise takes into account only the changes at the cutoff, which is why it should be thought of as a local counterfactual exercise. Further, I assume that the income composition of drivers at the cutoff remains unchanged under the counterfactual.

threshold, $\hat{\beta}_g$ as the quartile-specific IV estimate (effect on reoffending per €1), ΔF_g^{IB} as the income-specific discontinuity in fines (first stage), and ΔF^{flat} as the hypothetical safety-neutral fixed fine increase.

The observed threshold-crossing effect on reoffending due to the income-based fine is equal to

$$\Delta S^{IB} = \sum_g s_g \, \hat{\beta}_g \, \Delta F_g^{IB}.$$

Under a fixed fine, the corresponding effect would be:

$$\Delta S^{flat} = \sum_g s_g \, \hat{\beta}_g \, \Delta F^{flat} = \Delta F^{flat} \sum_g s_g \, \hat{\beta}_g.$$

Equating $\Delta S^{IB} = \Delta S^{flat}$ and solving for ΔF^{flat} yields the safety-neutral fixed fine increase

$$\Delta F^{flat} = \frac{\sum_g s_g \, \hat{\beta}_g \, \Delta F_g^{IB}}{\sum_g s_g \, \hat{\beta}_g}. \tag{5}$$

Figure 12(a) compares the safety-neutral fixed fine increase (red bars) to the observed income-based increases (blue bars). Applying equation (5), I find that the safety-neutral fixed fine increase would be € 215. This is close to the average observed increase (€ 204, Table 3), but its distributional implications differ substantially. Relative to the observed schedule, the increase would more than quadruple for the bottom income quartile (€ 52 to € 215) and roughly double for the second quartile (€ 103 to € 215). It would rise modestly for the third quartile (€ 187 to € 215; about 15%), and it would fall sharply for the top quartile (€ 499 to € 215; about 57%).[39]

These changes are sizable relative to income levels. Under the safety-neutral fixed schedule, the fine for offenses just above the 20 km/h cutoff would be €415 (€200 baseline plus a €215 increase). For drivers in the bottom income quartile, this corresponds to roughly 41% of average monthly net income, compared with 25% under the current income-based schedule. For the second quartile, the share would rise from 16% to 21%. By contrast, the burden would fall from 15% to 9% for drivers in the top quartile.[40]

Although I construct the two systems to produce the same local reduction in reoffending at the cutoff, they imply different patterns of behavioral adjustment across income groups. Figure 12(b) reports the implied threshold-crossing effects on reoffending by income quartile

[39]The fine jump for the bottom quartile is small (about €52), so the treatment contrast for this group is modest. This makes it harder to pin down the deterrence effect for the bottom quartile. As sensitivity checks, I set the bottom-quartile IV estimate equal to (i) zero and (ii) the second quartile's IV estimate. These exercises imply safety-neutral fixed fines of €207 and €153, respectively. In both cases, the qualitative distributional conclusions remain unchanged.

[40]Figure 8 presents average monthly net income by income quartile. The average monthly net income is €1011 for the first quartile, €1951 for the second quartile, €2642 for the third quartile, and €4745 for the fourth quartile.

Figure 12: Safety-neutral Fines and Reoffending

(a) Safety-neutral Fines (b) Reoffending Under Safety-neutral Fines

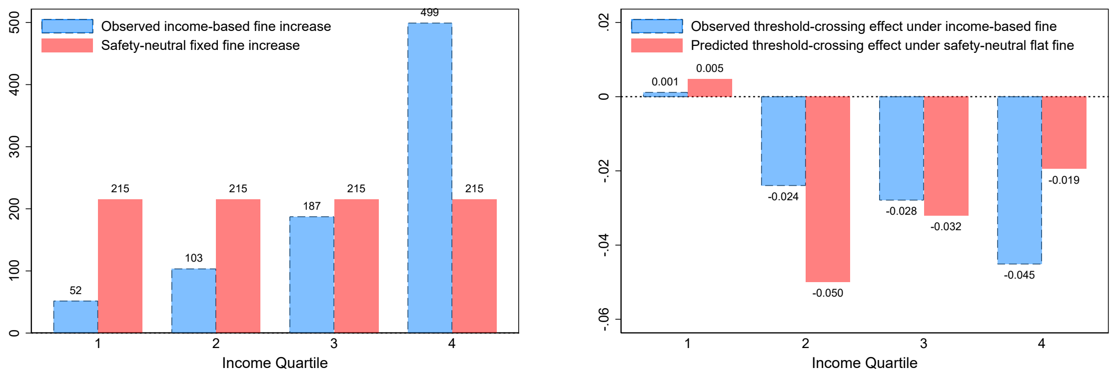

Figure plots the safety-neutral fixed fine increase and its implied effect on reoffending at the income-based cutoff. Panel (a) shows the observed income-based fine increase (blue) and the estimated safety-neutral fixed fine increase (red) by income quartile. The safety-neutral fixed fine increase is calculated using equation 5. By construction, the safety-neutral fixed fine increase yields the same aggregate specific deterrence effect as the observed income-based increase. Panel (b) reports the specific deterrence effects by quartile under the income-based system (blue bars) and the counterfactual effects implied by the safety-neutral fixed fine increase. Counterfactual effects are computed as the product of the fixed fine increase and the quartile-specific IV estimate (Table 3).

under each system. To obtain the counterfactual effects under the fixed-fine schedule, I multiply the safety-neutral fixed-fine increase by the quartile-specific IV estimate. Relative to the current income-based schedule, the fixed-fine counterfactual reduces reoffending in the lower income quartiles and increases it in the top quartile.

Discussion The safety-neutral fixed fine counterfactual illustrates the trade-off between income-based and fixed fine systems. Fixed fines have practical advantages, such as ease of administration. Moreover, from a classical Beckerian perspective, optimal fines need not depend on income unless the harm from the offense varies with income (Becker, 1968).

However, if policymakers seek to minimize safety risks from speeding through fixed fines alone, the implied fine level may impose substantial financial strain on low-income households. Mello (2024) finds that a speeding fine of \$200 can push especially illiquid low-income households to miss bill payments, with persistent downstream consequences such as lower credit scores and borrowing limits, weaker employment, and reduced homeownership. Although the institutional setting differs, these findings suggest that the €415 fine implied by my safety-neutral fixed fine counterfactual — roughly 41% of monthly net income for bottom-quartile drivers — could plausibly generate financial distress among low-income households.

The counterfactual exercise also suggests that moving to a safety-neutral fixed-fine system may have revenue implications for governments. Under this counterfactual, fine revenue may decline through two channels: high-income drivers would face substantially lower fines than under the current income-based schedule, mechanically reducing revenue per offense among those drivers; and the higher burden on low-income drivers may reduce collection rates if fines exceed what some drivers can realistically pay (Norris and Rose, 2024). This concern is relevant to ongoing debates in the United States, where local governments increasingly rely on fixed fines and court fees to finance expenditures (ALCPIL, 2020).[41]

Finally, the safety-neutral counterfactual highlights how income-based and fixed fines differ with respect to the proportionality of punishment when overall safety is held fixed. Retributive justice holds that punishment should reflect the wrongfulness of the act. A central element of modern retributive theory is the principle of proportionality: identical offenses should receive equivalent sanctions. Under the income-based system, the fine equals 25% of monthly net income for drivers in the bottom income quartile and 15% for those in the top quartile. Under the safety-neutral fixed fine, these shares are 41% and 9%, respectively.[42] By this metric, replacing income-based fines with fixed fines would widen the gap in effective punishment across income groups, moving further away from proportionality and from the equal-burden rationale underlying Finland's day-fine system.

7 Conclusions

This paper studied drivers' reactions to discontinuous changes in the criminal sanctions created by the Finnish income-based fine system. First, I find that drivers do not bunch below the income-based fine cutoff, leading to unexpectedly smooth speeding distributions. This result also holds for high-income drivers experiencing the largest jumps in the size of the fine. However, I show that especially high-income drivers are less likely to recidivate in the short term when they are assigned a larger ticket vs. a smaller one. Finally, I calculate that replacing the income-based fine increase with a safety-neutral fixed fine would require a fixed fine increase of about €215.

To reconcile these results, I present a theoretical framework in which individuals make decisions and learn using a misspecified model of the penalty schedule. This model rationalizes the smooth speeding distributions and the specific deterrence response.

My results have important implications for the design of criminal sanctions. First, they suggest that punishments may have their intended effects only if the penalty schedule is sufficiently salient. Second, the counterfactual exercise shows that replacing the income-

[41] In Appendix Section C.3, I apply optimal income and commodity taxation theory to derive the optimal income-based fine when the government aims to redistribute using fines. I show that the optimal income-based fine equals the size of the externality plus the net social value of redistribution, which is divided by the semi-elasticity of speeding with respect to fines.

[42] Appendix Section C.2 shows that under risk aversion, a fixed fine implies larger utility losses for low-income than for high-income drivers.

based fine increase with a safety-neutral fixed fine would substantially redistribute the size of fines across the income distribution.

Data Availability Statement

The analyses in this paper rely on a combination of publicly available data and confidential data, which cannot be shared. All publicly available data and all programs used to produce the results in this paper are provided in a Zenodo repository at https://doi.org/10.5281/zenodo.19662388. The README file included in this repository provides instructions for obtaining the confidential data used in this paper.

References

Aizer, A. and Doyle, J., Joseph J. (2015). Juvenile Incarceration, Human Capital, and Future Crime: Evidence from Randomly Assigned Judges *. *The Quarterly Journal of Economics*, **130** (2), 759–803.

ALCPIL (2020). Fees, Fines, and the Funding of Public Services: A Curriculum for Reform. Contributors: Mitali Nagrecha, Anna VanCleave, Stephanie Garlock, Judith Resnik, Lisa Foster, Jeff Selbin.

Armstrong, T. B. and Kolesár, M. (2020). Simple and Honest Confidence Intervals in Nonparametric Regression. *Quantitative Economics*, **11** (1), 1–39.

Baley, I. and Veldkamp, L. (2021). *Bayesian Learning*. Working Paper 29338, National Bureau of Economic Research.

Banerjee, A., Duflo, E., Keniston, D. and Singh, N. (2019). *The Efficient Deployment of Police Resources: Theory and New Evidence from a Randomized Drunk Driving Crackdown in India*. Working Paper 26224, National Bureau of Economic Research.

Bastani, S. and Selin, H. (2014). Bunching and Non-bunching at Kink Points of the Swedish Tax Schedule. *Journal of Public Economics*, **109**, 36–49.

Bauernschuster, S. and Rekers, R. (2022). Speed Limit Enforcement and Road Safety. *Journal of Public Economics*, **210**, 104663.

Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, **76** (2), 169–217.

Bénabou, R. and Tirole, J. (2002). Self-confidence and Personal Motivation. *The Quarterly Journal of Economics*, **117** (3), 871–915.

— and Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, **112** (4), 848–886.

Berk, R. H. (1966). Limiting Behavior of Posterior Distributions when the Model is Incorrect. *The Annals of Mathematical Statistics*, **37** (1), 51 – 58.

Best, M. C., Cloyne, J. S., Ilzetzki, E. and Kleven, H. J. (2019). Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches. *The Review of Economic Studies*, **87** (2), 656–690.

Bhuller, M., Dahl, G. B., Løken, K. V. and Mogstad, M. (2020). Incarceration, Recidivism, and Employment. *Journal of Political Economy*, **128** (4), 1269–1324.

Bohren, J. A. and Hauser, D. N. (2021). Learning With Heterogeneous Misspecified Models: Characterization and Robustness. *Econometrica*, **89** (6), 3025–3077.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2020). Memory, Attention, and Choice*. *The Quarterly Journal of Economics*, **135** (3), 1399–1442.

Calonico, S., Cattaneo, M. D. and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica*, **82** (6), 2295–2326.

Chalfin, A. and McCrary, J. (2017). Criminal Deterrence: A Review of the Literature. *Journal of Economic Literature*, **55** (1), 5–48.

Chetty, R. (2012). Bounds on Elasticities With Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica*, **80** (3), 969–1018.

—, Friedman, J. N., Olsen, T. and Pistaferri, L. (2011). Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records *. *The Quarterly Journal of Economics*, **126** (2), 749–804.

—, — and Saez, E. (2013). Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings. *American Economic Review*, **103** (7), 2683–2721.

DeAngelo, G. and Hansen, B. (2014). Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities. *American Economic Journal: Economic Policy*, **6** (2), 231–57.

Diamond, P. (1975). A Many-person Ramsey Tax Rule. *Journal of Public Economics*, **4** (4), 335–342.

Doleac, J. L. (2023). Encouraging Desistance from Crime. *Journal of Economic Literature*, **61** (2), 383–427.

Drago, F., Galbiati, R. and Vertova, P. (2009). The Deterrent Effects of Prison: Evidence from a Natural Experiment. *Journal of Political Economy*, **117** (2), 257–280.

—, Mengel, F. and Traxler, C. (2020). Compliance Behavior in Networks: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*, **12** (2), 96–133.

Dušek, L. and Traxler, C. (2022). Learning from Law Enforcement. *Journal of the European Economic Association*, jvab037.

Einav, L., Finkelstein, A. and Schrimpf, P. (2019). Bunching at the kink: Implications for spending responses to health insurance contracts. *Journal of Public Economics*, **171**, 117–130, trans-Atlantic Public Economics Seminar 2016.

Elminejad, A., Havranek, T. and Irsova, Z. (2025). Relative Risk Aversion: A Meta-Analysis. *Journal of Economic Surveys*.

Esponda, I. and Pouzo, D. (2016). Berk–Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models. *Econometrica*, **84** (3), 1093–1130.

Finlay, K., Gross, M., Lieberman, C., Luh, E. and Mueller-Smith, M. (2024). The Impact of Criminal Financial Sanctions: A Multistate Analysis of Survey and Administrative Data. *American Economic Review: Insights*, **6** (4), 490–508.

Fintraffic (2020). Traffic Monitoring System (TMS) Speed Data, 2020.

Ganong, P. and Jäger, S. (2018). A Permutation Test for the Regression Kink Design. *Journal of the American Statistical Association*, **113** (522), 494–504.

Gehrsitz, M. (2017). Speeding, Punishment, and Recidivism: Evidence from a Regression Discontinuity Design. *The Journal of Law and Economics*, **60** (3), 497–528.

Goncalves, F. and Mello, S. (2021). A Few Bad Apples? Racial Bias in Policing. *American Economic Review*, **111** (5), 1406–41.

— and — (2023). *Police Discretion and Public Safety*. Tech. rep.

Hansen, B. (2015). Punishment and Deterrence: Evidence from Drunk Driving. *American Economic Review*, **105** (4), 1581–1617.

Harju, J., Matikka, T. and Rauhanen, T. (2019). Compliance costs vs. tax incentives: Why do entrepreneurs respond to size-based regulations? *Journal of Public Economics*, **173**, 139–164.

HE (1920). Hallituksen esitys 36/1920.

Ito, K. (2014). Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing. *American Economic Review*, **104** (2), 537–63.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, **8** (1), 435–464.

— and Waseem, M. (2013). Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan *. *The Quarterly Journal of Economics*, **128** (2), 669–723.

Kostøl, A. R. and Myhre, A. S. (2021). Labor Supply Responses to Learning the Tax and Benefit Schedule. *American Economic Review*, **111** (11), 3733–66.

Lee, D. and McCrary, J. (2017). The Deterrence Effect of Prison: Dynamic Theory and Evidence. *Advances in Econometrics*, **38**, 73–146.

Lochner, L. (2007). Individual perceptions of the criminal justice system. *American Economic Review*, **97** (1), 444–460.

Londoño-Vélez, J., Rodríguez, C. and Sánchez, F. (2020). Upstream and downstream impacts of college merit-based financial aid for low-income students: Ser pilo paga in colombia. *American Economic Journal: Economic Policy*, **12** (2), 193–227.

Mello, S. (2024). Fines and Financial Wellbeing. *The Review of Economic Studies*.

Mueller-Smith, M. (2014). The Criminal and Labor Market Impacts of Incarceration. *Working Paper*.

National Police Board of Finland (2020). Speeding Ticket Data, 2018–2020.

Norris, S. and Rose, E. K. (2024). Laffer's Day in Court: The Revenue Effects of Criminal Justice Fees and Fines. *Journal of Public Economics*, **240**, 105249.

Nyarko, Y. (1991). Learning in Mis-specified Models and the Possibility of Cycles. *Journal of Economic Theory*, **55** (2), 416–427.

Our World in Data (2022). Causes of Death.

Philippe, A. (2024). Learning by Offending: How Do Criminals Learn about Criminal Law? *American Economic Journal: Economic Policy*, **16** (3), 27–60.

Rees-Jones, A. and Taubinsky, D. (2019). Measuring "Schmeduling". *The Review of Economic Studies*, **87** (5), 2399–2438.

Rincke, J. and Traxler, C. (2011). Enforcement Spillovers. *The Review of Economics and Statistics*, **93** (4), 1224–1234.

Rose, E. K. and Shem-Tov, Y. (2021). How Does Incarceration Affect Reoffending? Estimating the Dose-Response Function. *Journal of Political Economy*, **129** (12), 3302–3356.

Saez, E. (2001). Using Elasticities to Derive Optimal Income Tax Rates. *The Review of Economic Studies*, **68** (1), 205–229.

Statistics Finland (2022a). Finnish Police Reports.

Statistics Finland (2022b). FOLK Basic Data, 1987–2020.

Statistics Finland (2022c). Statistics on Offences and Coercive Measures.

Statistics Finland (2023). Road Accident Data, 1989–2022.

Traxler, C., Westermaier, F. G. and Wohlschlegel, A. (2018). Bunching on the Autobahn? Speeding responses to a 'notched' penalty scheme. *Journal of Public Economics*, **157**, 78–94.

van Benthem, A. (2015). What is the Optimal Speed Limit on Freeways? *Journal of Public Economics*, **124**, 44–62.

WHO (2017). *Managing Speed*. Report, World Health Organization.

Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review*, **110** (2), 337–61.