Income-Based Fines: Evaluating Rationales for Finnish Speeding Penalties

Aaron Payne* *(Job Market Paper)* & Martti Kaila†

January 6, 2026

[Click here for most recent draft]

Abstract

We evaluate rationales for income-based speeding fines. Such fines are famously applied in Finland, where a $250 offense for a low-income speeder can cost $100,000 for the rich. We consider four potential policy goals—externality mitigation, redistribution, equal compliance, and proportional punishment—using linked Finnish administrative data and an original survey. First, using a "job-loss" design, we find that marginal speeding costs *decrease* with offender income, leading us to reject externality mitigation as a rationale. Next, we assess the redistributive rationale. In standard models, indirect taxes aid redistribution only if the taxed action serves as a signal for earnings ability. We recover this signal by differencing speeding behavior with respect to causal effects of income on speeding, estimated from within-individual earnings variation and inheritance shocks. The estimated signal is *negatively* correlated with earnings, implying that Mirrleesian redistribution rationalizes a *lower* fine on the rich. Finally, we use our survey to assess whether preferences for equal compliance or proportional punishment are motivating rationales. Respondents trade off fixed and income-based speeding fine policies, where the latter vary in the induced across-income distribution of behavior and schedule "steepness." Net of redistribution, respondents are on-average willing to forgo €144 million in government revenue to implement income-based policies. However, these valuations are insensitive to induced behavior or schedule-steepness, ruling out our candidate rationales. Given our findings, speeding fines should incorporate *some* income-dependence but need not feature the extremes in Finland. Residual support for income-dependence suggests that current policy reflects either unmodeled instrumental considerations or direct, non-instrumental preferences.

*The Wharton School at the University of Pennsylvania; E-mail: `aapayne@wharton.upenn.edu`; corresponding author.

†Adam Smith Business School at the University of Glasgow; E-mail: `martti.kaila@glasgow.ac.uk`

We thank Benjamin Lockwood, Alex Rees-Jones, Judd Kessler, and David Abrams for their invaluable feedback and support. We also want to thank Kathy Bi, Jessie Handbury, Kristiina Huttunen, Matti Liski, Thomas Lloyd, Peter Lugthart, Juuso Toikka, Ben Sprung-Keyser, Eva Wang, and participants at the National Tax Association Annual Conference, Wharton Applied Economics Workshop, and International Institute of Public Finance Annual Congress for helpful comments and discussions. We acknowledge financial support from the Yrjö Jahnsson Foundation and the Wharton Behavioral Lab. The survey in this project is pre-registered at AsPredicted.org

1 Introduction

Governments use punishments to deter anti-social behavior, be it tax evasion, littering, or murder. In developed countries, the majority of offenses committed are punished with a fine. However, despite their ubiquity, there is substantial heterogeneity in how fines are determined. While fines commonly scale with the severity of the offense, several countries (e.g. Germany, Switzerland, Denmark, Sweden, and Finland) also scale some fines with income. In Finland, the oldest and most comprehensive income-based system, the rich can receive fines in excess of €100,000 for speeding offenses that would cost a few hundred euros for low-income drivers.[1]

How can we rationalize income-based fine systems? More broadly, are they effective in achieving their economic and/or policy goals? In this paper, we present the leading economic and fairness-based rationales that have been put forth for these policies and test whether any of these motives can explain the Finnish speeding-fine system.

On the economic side, two frameworks are readily applicable to speeding fines. First, we consider the optimal penalty framework in Becker (1968), which prescribes that fines be set equal to the marginal social cost of crime. In this framework, income-based speeding fines can be rationalized if the marginal social cost of speeding increases with income. Second, we apply a Mirrleesian optimal redistribution framework to our setting. In this framework, the planner's objective is to redistribute resources from high- to low-ability types, but the un-observability of type means the planner must turn to distortionary taxation. In the canonical Atkinson-Stiglitz result (Atkinson and Stiglitz, 1976), it is most efficient for the planner to use direct income taxes, rather than indirect commodity taxes, to redistribute resources. However, subsequent work (Saez, 2002a; Allcott et al., 2019; Ferey et al., 2024) shows that indirect taxes can lower the efficiency costs of redis-

[1]For example, Andres Wiklöf was fined €121,000 for speeding in 2023; in 2002, Nokia executive Anssi Vanjoki received a speeding fine of €116,000.

tribution if the taxed behavior is a signal of earnings ability. Building on these results, we show that speeding fines should increase with income if the signaling value of speeding behavior increases with income, giving us a testable benchmark against which to assess redistributive rationales for income-based fines.

Fairness rationales have also featured prominently in the debate surrounding income-based fines. For example, official argumentation in the lead up to the implementation of Finland's policy cited a desire to equalize the "impact" of fines on people of varying means (Lahti, 2021). This evidence suggests a deterrence objective, consistent with scholarly work citing general deterrence across incomes as an attractive theoretical basis for income-based systems (Kantorowicz-Reznichenko, 2013); however, it is unclear what formal notion of deterrence underlies these motivations. As we show in section 3, with log-utility over consumption, speeding fines that are proportional to net-of-tax income (as in Finland) generate equal *marginal deterrence* across incomes, which we define as across-income equality in the schedule of marginal speeding costs. Therefore, if Finland's income-based fine policy is an instrument for achieving marginal deterrence across incomes, then policy preferences should load on a particular fine schedule (e.g. proportional to income). However, under the same log-utility model, proportional fines *also* generate equal speeding rates across incomes; as it so happens, income-based speeding fines in Finland are empirically associated with across-income equality in speeding rates. We conclude that cited rationales are ambiguous: they may be motivated by either a particular fine schedule (e.g. as an instrument to achieve marginal deterrence) or a particular behavioral outcome (e.g. equal speeding rates across incomes). We therefore test whether the equal distribution of behavior (hereafter, "equal compliance") or the proportional schedule itself is the instrumental objective of income-based fines in Finland.

To empirically evaluate each of these four rationales, we draw on linked Finnish administrative data and construct an original survey that we administer to a representative sample of Finnish residents.

To evaluate the Beckerian rationale, we measure the slope of marginal social speeding costs with respect to income using the injuries that result from accidents, as collected in accident report data. Linking this information to data from income-tax returns and crime histories, we leverage cross-sectional variation in rates of speeding-ticket receipt conditional on income to measure the marginal impact of speeding on accident-related injuries across income levels. With respect to injuries, we find the slope of marginal speeding costs is *negative*: the marginal impact of speeding on accident-related injuries is *lower* at high incomes. Drawing on additional data, we show that this result is robust to controlling for non-speeding traffic-crime rates, car ownership status, annual driving distance (measured using odometer data), and vehicle mass. Evidence from the accident reports suggests that across-income differences in traffic density on the road may explain our results. To translate the marginal speeding cost gradient into a money-metric, we measure the cost of each injury using a "job-loss" design (von Wachter and Bender, 2006); this approach allows us to capture across-victim heterogeneity in costs and to separate fiscal externalities from net-income losses, which may be weighed differently in the planner's objective. In particular, we compare the earnings trajectories of individuals injured in an accident to those of demographically-matched individuals from the general population to estimate the net-present-value (NPV) cost of each injury. We apply welfare weights (Saez, 2002b) to NPV net-income losses to account for the possibility that the planner weighs a given income loss differently depending on the income of the victim. With our NPV costs in-hand, we re-estimate the slope of marginal speeding costs with respect to income. Regardless of the discount rate we apply to future injury costs or how we welfare-weight net-income losses, we find a consistently *negative* slope. In our preferred estimates, we find that the marginal social cost of speeding decreases by €7.6 (se=4.4) for every €10,000 of income, well below our estimate of the overall slope of the Finnish speeding fine schedule (+12.4 €). We therefore reject the externality-based rationale.

Turning to the redistributive rationale, we measure the extent to which speeding behavior is a signal of earnings ability. As clarified by Allcott et al. (2019) and Ferey et al. (2024), this signal is identified by differencing observed across-income variation in speeding rates with respect to causal effects of earnings on speeding, where the latter are identified from variation in *labor supply*. With weak separability between labor supply and consumption (Allcott et al., 2019; Ferey et al., 2024), the earnings effect can be identified using any exogenous earnings variation. We reject this assumption in our setting, implying that the earnings effect is the sum of a labor supply effect (holding earnings fixed) *and* an income effect (holding labor supply fixed). To measure the labor supply effect, we use a yearly panel of income and speeding ticket receipt and leverage within-individual variation in months-employed (holding earnings fixed). We estimate an average labor-supply effect that is about 25% as large as the across-income gradient of speeding ticket receipt. Due to mis-measured labor supply, the coefficient on earnings in this regression is likely upward-biased, so we turn to an inheritance-shock approach to estimate the income effect. Building on Nekoei and Seim (2022), we exploit the differential timing of first parental deaths as an exogenous income shock; to address violations of the exclusion restriction, we also difference with respect to inheritance size, yielding a triple-differences approach. From this approach, we estimate an income effect that is roughly equal to the average labor supply effect. We then difference observed speeding behavior with respect to our earnings effects and find that the resulting earnings ability signal is *negatively* correlated with income, implying that Mirrleesian redistributive motives rationalize a *lower* fine for the rich. In our preferred estimate, fines should *decrease* by €12 for every €10,000 of income. Across various assumptions about the planner's redistributive preferences, we find consistently negative fine gradients, leading us to reject the redistribution-based rationale for income-based fines.

Finally, to evaluate preferences for equal compliance and proportional punishment, we turn to our survey. In particular, we ask Finnish residents to make trade-offs between hypothetical income-dependent and fixed-fine policies using government revenue (our willingness-to-pay medium). We vary three key elements of the policy comparison: the distribution of behavior under the income-dependent policy, the "steepness" of the income-dependent fine schedule, and the behavior change induced by income-dependent fines. We hold total speeding fixed across policies to address safety-related confounds; to address redistributive motives, we difference WTP for income-dependent fines with respect to WTP for income-dependent transfers. Net of redistribution, 56% of respondents have a positive WTP for income-dependent fines, with respondents on-average willing to forgo €144 million in government revenue to implement an income-based policy. However, we find relatively small and insignificant responses to the three randomized fine-policy characteristics described above; if anything respondents express a slight *distaste* for equal behavior and proportional fine schedules. To assess the robustness of our findings, we explore heterogeneity with respect to other randomized characteristics of the policy comparison and respondents' background characteristics; across sub-groups, respondents consistently value income-based fines but remain unresponsive to randomized features of the comparison.

In summary, our results suggest that commonly cited economic and fairness-motivated rationales for income-based fines are inadequate to account for observed policy. Nonetheless, we find that Finnish residents have a strong residual willingness-to-pay for income-based fines that is orders of magnitude larger than the inefficiencies resulting from Beckerian and redistributive motives. Given the absence of a preference for steeper fine schedules in our survey, optimal fines should therefore incorporate *some* income-dependence, but need not exhibit the extremes of the Finnish system. Our results also suggest that accounting for the popularity of income-based policies requires either an appeal to alternative fairness consideration (e.g. procedural fairness) or a richer economic model that incorporates, e.g., liquidity concerns or income-correlated mis-perceptions.

After discussing related literature and contributions, our paper proceeds as follows. First, section 2 discusses our setting: speeding fines in Finland. We then turn to a theoretical framework (section 3), which clarifies the empirical objects we plan to measure for assessing each of the four rationales for income-dependent fines. We briefly discuss data (section 4) before testing the externality- (5) and redistribution-based (6) rationales. We then turn to our survey (7), which we use to assess the fairness-based rationales of equal compliance and proportional punishment. Section 8 assesses our results jointly, and section 9 concludes.

1.1 Related Literature

This projects speaks to three main literatures. First, there is an extensive literature that explores the redistributive potential of commodity taxes in Mirrleesian models (Mirrlees, 1971) where income taxes are present. The benchmark result, due to Atkinson and Stiglitz (1976), shows that, under preference homogeneity, commodity taxes are inefficient (relative to income taxation) for redistributing resources. Saez (2002a) clarified conditions under which this result breaks down, and more recent work (Allcott et al., 2019; Ferey et al., 2024) shows how to identify the redistributive value of commodity taxes for optimal tax design using the difference between observed across-income consumption variation and causal effects of earnings on consumption. While we build methodologically on this literature, our paper differs in two key ways. First, empirically, recent work that identifies earnings effects (Allcott et al., 2019; Ferey et al., 2024) plausibly relies on weak separability assumptions between labor supply and consumption, implying that earnings effects can be identified using any exogenous income variation. However, we test for and reject weak separability in our setting; as a result, our empirical approach involves estimating both labor supply and income effects on consumption. Second, rather than seeking to identify optimal tax rates, as Allcott et al. (2019) does for soda and Ferey et al. (2024) does for savings, we invert the approach, using the optimal redistributive speeding fine as a benchmark against which to test rationales for a given policy.

In relation to this inversion approach, our paper also relates to the "inverse optimum" literature (Christiansen and Jansen, 1978; Bourguignon and Spadaro, 2012; Bargain et al., 2014; Hendren, 2020), which uses analytical results from optimal tax theory and estimates of behavioral elasticities to infer society's revealed redistributive preferences. This approach has the benefit of removing the difficulty of selecting normative criteria, treating this as an empirical inference exercise. However, such an approach leaves the source of society's preferences opaque and can lead to results that contradict conventionally held assumptions (Lockwood and Weinzierl, 2016) with little ability to diagnose the source of the inconsistency. Our paper follows this literature in spirit but seeks to open the black box of society's revealed preferences, enabling us to assess their normative consistency.

Finally, this work relates to a theoretical and experimental literature on domain-specific fairness. On the theoretical side, specific egalitarianism (Tobin, 1970) argues that specific goods should be equally allocated; recent experimental evidence provides support for this view in the domains of health care and legal aid (Caspi et al., 2024). In our setting, the fact that high-income individuals consume more speeding (a good allocated by the justice system) may be viewed as a violation of specific egalitarianism.[2] Furthermore, in the context of real-stakes games, a wealth of lab evidence has examined under what conditions individuals favor equal allocations, findings that have given rise to numerous fairness theories (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Charness and Rabin, 2002; Engelmann and Strobel, 2004). However, work that explicitly links such fairness concerns to observed policy remains limited. Our paper makes progress on this front by taking as its object a policy (income-based speeding fines) that is associated with equality of behavior and eliciting preferences for such equality among individuals

[2]Repeat offender penalties can also be viewed as part of this objective.

subjected to the policy. This approach bears similarities to Rafkin and Soltas (2025), which uses experimentally elicited social preferences among landlords and tenants to conduct counterfactual analysis of rental assistance policies.

2 Setting: Speeding Fines in Finland

We now emphasize key features of our empirical setting. We use such features to inform our model and empirical approach.

2.1 Day Fine System in Finland

The first iteration of Finland's income-based fine system was implemented in 1921. Official argumentation in the lead-up to the implementation cited a desire to reduce the number of crimes punished by prison and equalize the "impact" of fines on people with varying means (Lahti, 2021). As a result, except for minor offenses (e.g. littering), all financial punishments are income-based. Each crime is punished with a fixed number of "day units." The value of each day-unit I depends on an individual's net-of-tax income according to the following formula:

$$I = \frac{y_m - 255}{60} - 3 \times D \tag{2.1}$$

where y_m is monthly net-of-tax income and D is an individual's number of dependents. Other than a €255 basic-living deduction, fines are proportional to half-a-day's net income, hence the "day fine" terminology.

Speeding is partially subjected to the income-based system: drivers are given fixed fines of 140-200 euros for exceeding the limit by 7-20 km/h (minor traffic infractions) and receive income-dependent fines for exceeding the limit by more than 20 km/h.[3] To visualize this schedule, figure 2.1 displays hypothetical speeding fines for a 100 km/h limit as

[3]Drivers are given grace for exceeding the limit by no more than 6 km/h.

a function of net-of-tax income for different levels of "excess speed" (*i.e.* the amount by which a person exceeds the speed limit).[4] For "minor" speeding crime (7-20 km/h over the limit), fines are constant across incomes; however, for day-fine eligible speeds, fines are essentially proportional to net-of-tax income except for a fine-floor that holds at low incomes.[5] The proportionality is shown explicitly in Appendix figure 10.1: above the fine floor, income-dependent fines constitute a fixed share of monthly net-of-tax income. We also note that there is no ceiling for speeding fines in Finland: individuals can receive arbitrarily large fines at arbitrarily high incomes.

Figure 2.1: Hypothetical Speeding Fine Schedule (100 km/h limit)

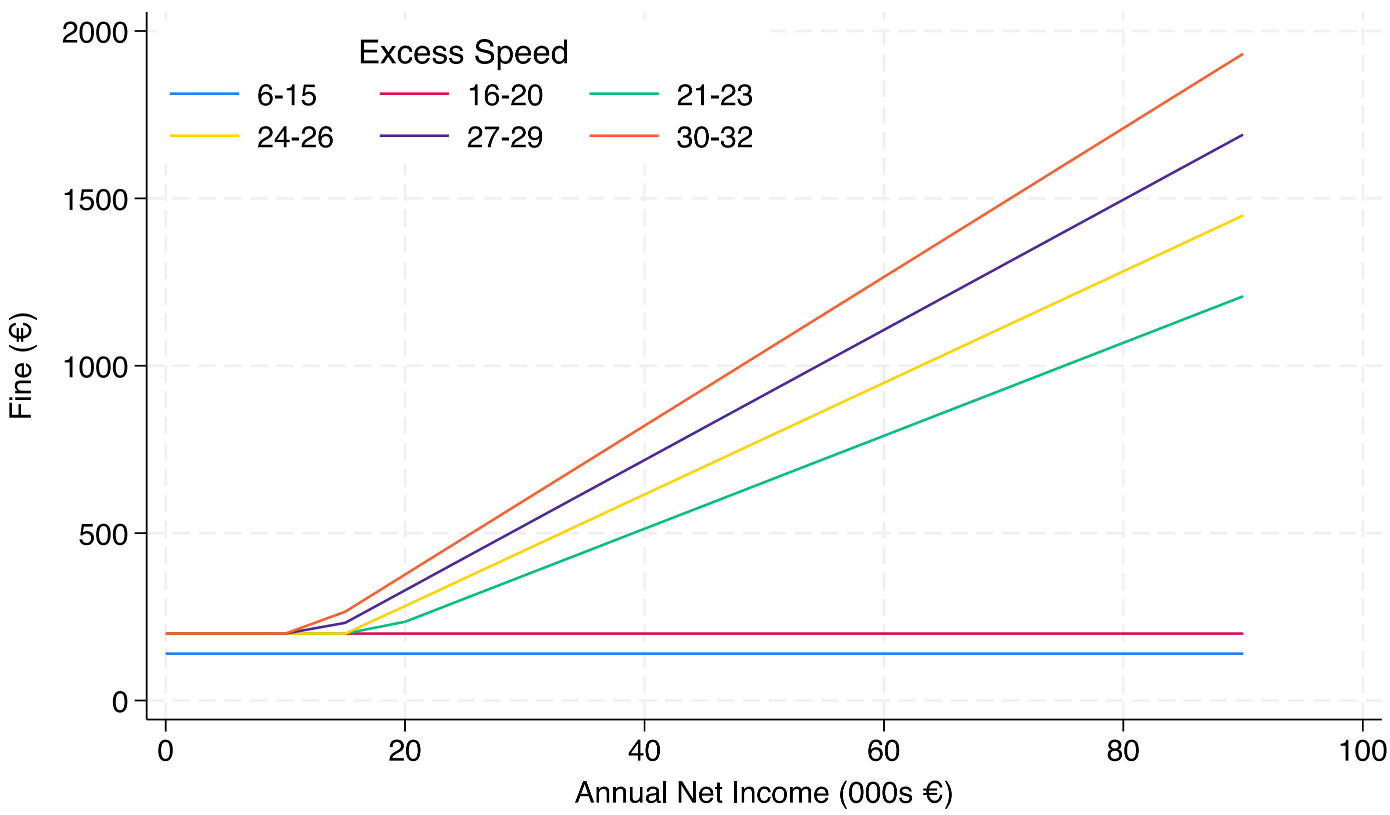

Hypothetical speeding-fine schedule, as a function of net income, separately by the amount by which a driver exceeds the limit (in kilometers per hour). Net income consists of earned, capital income, and transfers received, net of tax payments. To construct this schedule, we use equation 2.1 (assuming no dependents) for the monetary value of a day unit, together with information from the Finnish police on the typical number of day units assigned to a given speeding crime.

[4]While the income-dependent cutoff (21 km/h over the limit) is constant across limits, a given excess speed will have larger fines at lower speed limits.

[5]This floor is inherited from the fixed-fine region of the fine schedule.

Because of the income-dependent cutoff, there is a fine discontinuity at 121 km/h, the size of which grows with income.[6] As shown in Kaila (2024), despite large discontinuities, individuals do not bunch, either in ticket data collected from speeding cameras or in measurements from traffic monitoring stations. Furthermore, while Kaila (2024) shows using a regression-discontinuity-design that individuals who receive larger fines respond by *ex-post* reducing their likelihood of speeding (the extensive margin), there is no evidence that they respond on the *intensive-margin* (i.e. by lowering their speed to avoid the higher fine). We use the latter finding to inform our empirical and model decisions discussed below.

2.2 Relationship Between Income and Speeding

To motivate the potential to use speeding fines as a means for redistribution, figure 2.2 plots the relationship between income and the probability of receiving a speeding ticket in a given year using our crime-income panel described in section 4. We plot the relationship separately for the fixed- and day-fine regions of the fine schedule. We see a strong positive gradient for fixed-fine speeding crime: the fixed-fine speeding rate moves from around 2.5pp to 9pp as we move from the bottom to top earnings decile. While this relationship ostensibly provides a basis to use income-dependent speeding fines as a way to redistribute resources from high- to low-earners, such a possibility needs to be juxtaposed against the option of simply increasing income taxes on high earners, as we discuss in section 3.

We also use figure 2.2 to motivate key features of our survey and elicitation of fairness preferences. In particular, the day-fine region of the speeding-fine schedule is associated with fines that are proportional to net-of-tax income (equation 2.1) and a significantly compressed across-income distribution of behavior. While the latter is not necessarily causal, such an outcome could be an intended "impact" of income-based fines over which

[6]In particular, holding income fixed in figure 2.1, the discontinuity is the difference between the red ("16-20") line and the green ("21-23") line.

Figure 2.2: Binscatter Relationship Between Income and Speeding-Crime Rates

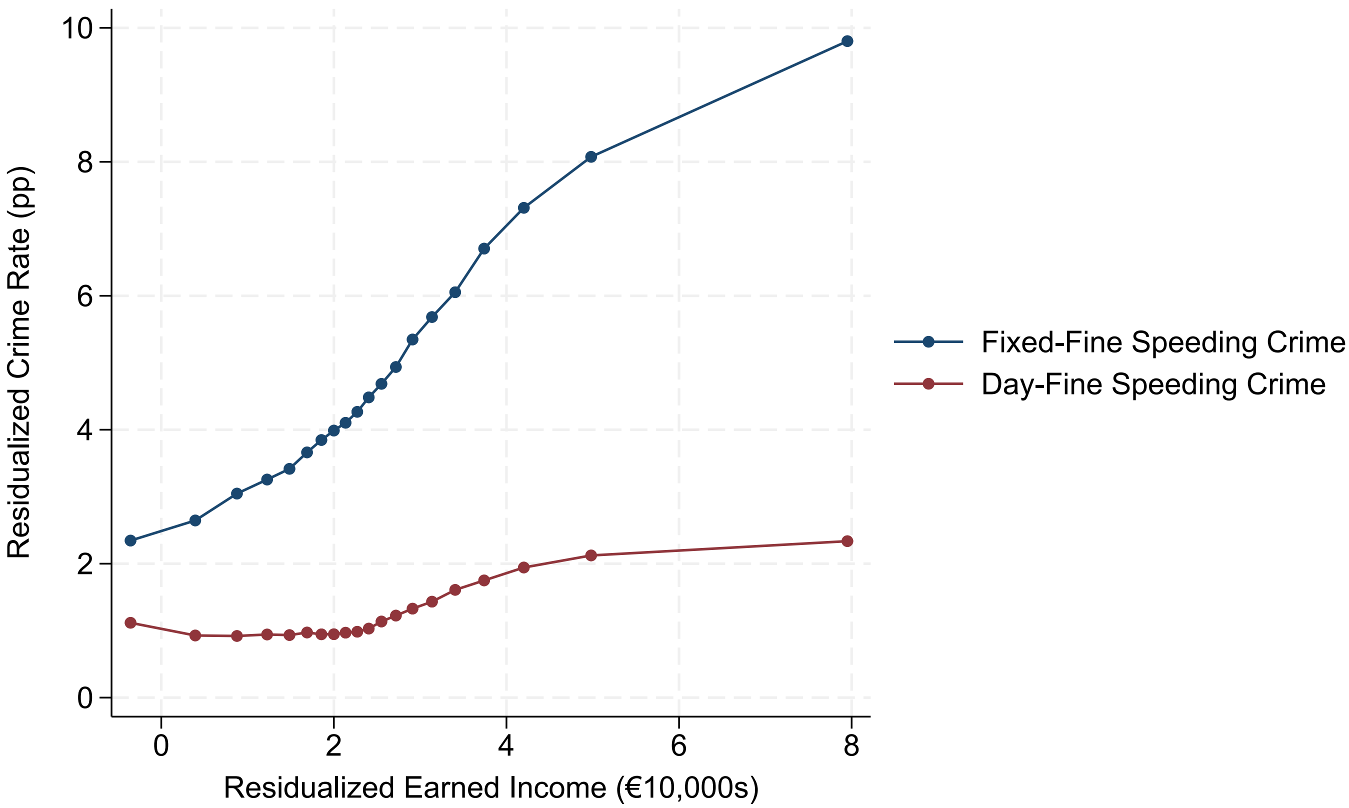

Residualized binscatter relationship between earned income and speeding-crime rates for the population of Finnish adults (*i.e.* individuals who are 18 years or older) from 2006-2019. "Earned income" is defined as taxable earned income, which is the sum of wage/salary income, unemployment insurance payments, and pension income. We say that an individual i committed a "speeding crime" in a given year y if i received a speeding ticket in year y, according to crime-history reports. Fixed-fine speeding crime corresponds to speeding tickets received for driving no more than 20 km/h above the limit, while day-fine speeding crime corresponds to all other speeding crime. We residualize earned income and speeding-crime rates on age-by-year fixed effects.

society may have fairness preferences. We disentangle the proportionality and equal-compliance channels using our survey (section 7).

2.3 Typical Speeding Crime

In our theoretical and empirical analysis, we use extensive-margin speeding-ticket receipt as our measure of (observed) speeding behavior, collapsing all speeding crime into a single "typical" speeding offense. We justify this simplification in two ways. First, as discussed above, there are no discernible intensive-margin responses to fine size; instead, individuals appear to respond on the extensive-margin (Kaila, 2024). Second, the "typical" speeding crime is similar across incomes. In particular, Appendix figure 10.2 shows, using police speeding-ticket data, that the average speed limit for a given speeding crime is uncorrelated with income (panel 10.2a). While there is a negative correlation between income and excess speed (panel 10.2b), Appendix table 10.1 shows that most of this relationship can be explained by differential police-officer discretion across municipalities (column 3 vs. column 5). We therefore view the typical speeding crime as approximately equivalent across incomes conditional on age and municipality.

2.4 Observed Income-Gradient of Speeding Fines

Given that our goal is to test whether the rationales we consider can justify the speeding-fine schedule in Finland, we need a benchmark fine schedule that accounts for our simplified extensive-margin measure of speeding behavior. To obtain this benchmark, we estimate the observed relationship between speeding fine amounts and income using the speeding offenses in our police ticketing data. In particular, Appendix figure 10.3 shows a binscatter of earnings levels and speeding fine amounts (combining both day- and fixed-fine speeding crimes); while the relationship exhibits a kink at low-incomes,[7] it is linear for most of the income distribution, with a slope of €12.4 per €10,000.[8] We take this

[7]The kink arises from the fine floor for income-based fines in figure 2.1.

[8]Under this gradient, the top income decile pays around €100 more on-average than the bottom income decile for a typical speeding crime.

slope as a benchmark against which to test rationales for Finland's speeding-fine policy; in particular, for each of the various rationales we consider, we test whether the optimal implied slope of the fine-schedule is positive and large enough to account for the observed fine-schedule gradient.

Appendix figure 10.4 re-estimates the slope between earnings and speeding fine amounts using *only* day-fine speeding crime; we find a gradient of €88 per €10,000 of income. Since this gradient encodes Finland's income-based speeding-fine policy, we discuss our results in the context of both this slope and the flatter €12.4 slope. However, as we show, even against the straw-man gradient of €12.4, we cannot rationalize the Finnish speeding-fine system.

3 Theoretical Framework

In this section, we build a model of optimal fine-setting in the spirit of an optimal commodity tax framework (Atkinson and Stiglitz, 1976; Saez, 2002a; Kaplow, 2006; Allcott et al., 2019), which embeds an optimal commodity tax decision within the redistributive income-tax framework of Mirrlees (1971). Our goal is to clarify the role of each rationale that we consider for income-dependence and set the stage for our empirical approach. In section 3.3, we also make our argument that cited rationales for income-based fines are theoretically ambiguous.

3.1 Setting

Individuals in the economy are indexed by a two-dimensional type $(\theta, \epsilon) \in \big(\Theta, E\big)$; θ indexes an individual's "earnings ability," *i.e.* their latent wage $w(\theta)$, while ϵ is a shock that governs the benefit of speeding.

There are two periods: in period 0, individuals make a labor supply choice $l(\theta)$ given their type and *ex-ante* uncertainty over ϵ; this choice generates an earnings level $z(\theta) := w(\theta) * l(\theta)$ for each type. In period 1, ϵ is realized and individuals make an *extensive-margin* speeding decision given their period-0 earnings choice; see section 2.3 for a justification of our decision to focus exclusively on the extensive-margin behavioral margin.

In period 1, an individual with crime decision $s \in \{0, 1\}$, earnings z, earnings ability θ, disposable income c, and shock ϵ receives *ex-post* utility $U\big(c, s * \epsilon, z; \theta\big)$. We assume that U is strictly increasing and concave in its first two arguments and strictly decreasing and concave in z. This assumption gives rise to a threshold $\epsilon^*(c, z; \theta)$ such that all type-θ z-earners with $\epsilon \geq \epsilon^*$ will choose to speed. We define $s(z; \theta)$ as the speeding rate for type θ (given earnings z). Given period 1 choices, individuals choose earnings $z(\theta)$ to maximize expected utility.

Taxes and Fines In the spirit of the Mirrlees model, the planner cannot observe individual types θ and therefore must levy taxes on choices z and s, giving rise to distortions.

We ignore any risk associated with speeding ticket receipt and assume that individuals are assessed fines based on their expected speeding rate $s(z; \theta)$. We justify this by noting that, unlike our one-shot static model, individuals in fact make many speeding decisions (e.g. every second they are driving). By central-limit-theorem arguments, as the number of choice occasions grows large, the variance of fine payments (and therefore utility) will be negligible.

We consider a restricted system of taxes and fines that reflects our empirical setting: an individual earning z with speeding rate $s(z)$ pays income taxes $T(z)$ and speeding fines $\tau(z) * s(z) := (a + b * z) * s(z)$. When $b = 0$, fines are "fixed," whereas $b > 0$ implies that fines scale with income. We refer to b as the "income-gradient" of the fine schedule.

Externalities Each individual's speeding behavior imposes a cost on others; this cost can vary with the identity of the offender, e.g. because they interact with certain types of victims. We denote the dis-utility experienced by victims of type θ from others' speeding as $E(\theta)$ and denote the imposed by offenders of type θ as $\overline{e}(\theta)$.[9]

Standard Planner's Problem The planner chooses $T(\cdot)$, a, and b to solve:

$$\begin{aligned} &\max_{T(\cdot),a,b} \int_{\Theta} \alpha(\theta) G\Big(V(\theta) - E(\theta)\Big) dF_{\theta}(\theta) \\ &\text{s.t.} \quad \int_{\Theta} \big[T\big(z(\theta)\big) + \big(a + b * z(\theta)\big) * s\big(z(\theta);\theta\big)\big] dF_{\theta}(\theta) \geq R \end{aligned} \tag{3.1}$$

where $G(\cdot)$ is a concave social welfare function, $\alpha(\cdot)$ is a Pareto-weighting function, and R is an exogenous revenue requirement. We make the following assumption.

Assumption 3.1. *At the optimum: (i) the nonlinear income tax $T(\cdot)$ is smooth; (ii) $z(\theta)$ is smooth and strictly increasing, and $s(z;\theta)$ is a smooth function of z and θ; and (iii) individuals' optima are unique and their first- and second-order conditions strictly hold.*

Parts (i)-(iii) of assumption 3.1 allow us to use perturbation-based techniques to derive the optimal tax and fine schedule. Furthermore, this assumption allows us to express all heterogeneity with respect to z, rather than θ, at the optimum; in other words, the tax-and-fine system perfectly separates types. Therefore, in what follows, we suppress dependence on type.

3.2 Optimal Fines

Before discussing optimal fine schedules, we define several terms that appear in our formulas. First, the *social marginal welfare weight* for a z-earner $g(z)$ captures the social

[9]We use overline notation to denote that the imposed cost is aggregated over victims.

value of a marginal transfer to a z-earner at the optimum (relative to the value of public funds); these weights embed the planner's preferences for redistribution.[10] We also define the following behavioral responses:

- Fine (semi-)elasticities: $s'_{\tau_s}(z) = \frac{\partial s(z)}{\partial \tau}$
- Causal earnings effects on speeding: $s'_{inc}(z) := \frac{ds(z;\theta)}{dz}\Big|_{\theta=\theta^{-1}(z)}$

Causal earnings effects measure the change in speeding caused by a change in earnings *holding type fixed*; this is distinct from the *cross-sectional* slope $s'(z) := \frac{ds(z)}{dz}$, which measures the difference in observed speeding across individuals with different incomes. These earnings effects are crucial for recovering the earnings ability signal embedded in speeding behavior. The intuition is given in figure 3.1, which shows two adjacent earnings levels, z_0 and z_1, and associated hypothetical speeding rates, $s(z_0)$ and $s(z_1)$. Under assumption 3.1, those earning z_1 have a higher type (earnings ability) than those earning z_0; we index their respective types by θ_1 and θ_0. The key experiment we use to determine whether speeding behavior is (locally) informative about earnings ability is by inducing those earning z_0 to increase their labor supply and earn z_1. Despite having the same earnings z_1, the two types speed at different rates, reflecting differences in underlying earnings ability and/or preferences; the planner can exploit this information for redistributive purposes by targeting fines based on earnings ability. We refer to the earnings ability tag from figure 3.1 as $s'_{het}(z) = s'(z) - s'_{inc}(z)$, following the notation in Ferey et al. (2024).

With this intuition in-hand, we now present the first result regarding the size of the fine gradient b, which is the key focus of our paper. Expressions for the optimal fixed component a and marginal income-tax rates $T'(\cdot)$ are shown in Appendix section 10.1.

[10] Formally:

$$g(z) := \frac{\alpha(z) G'\Big(V(z) - E(z)\Big) V'_c(z)}{\lambda}$$

where λ is the multiplier on the government budget constraint.

Figure 3.1: Latent Earnings Ability Tag

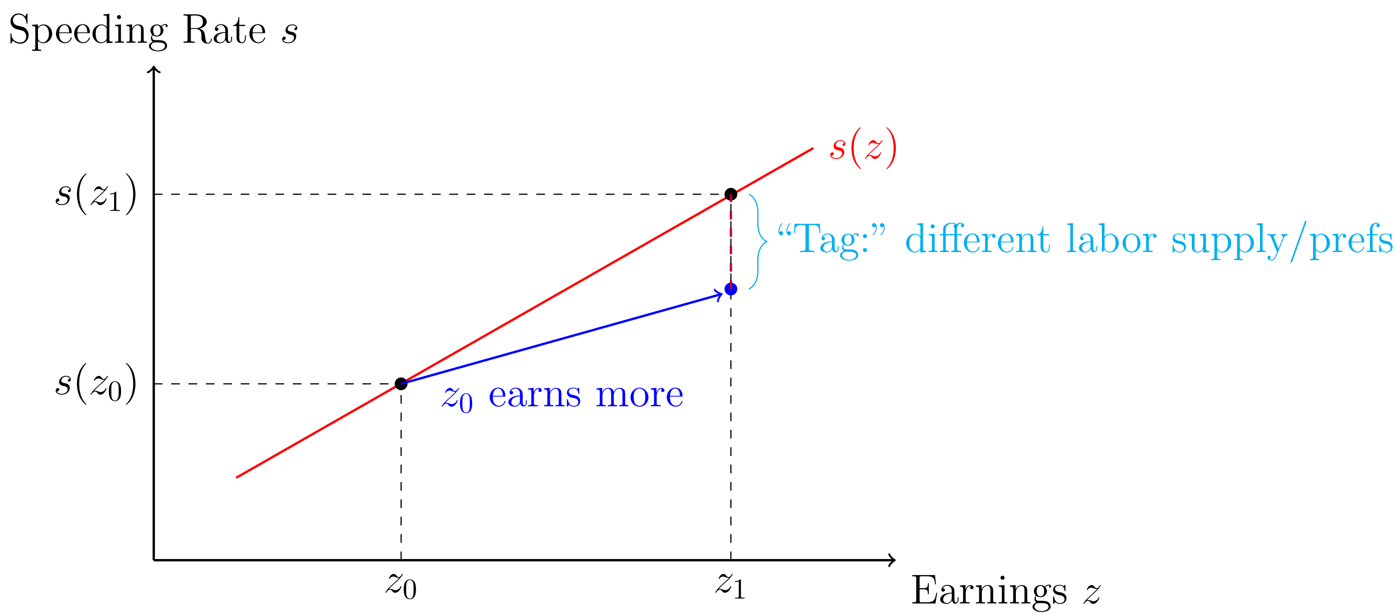

Theorem 1. *Suppose assumption 3.1 holds, the income tax schedule $T(\cdot)$ is optimal, the fixed component $a^*(b)$ is optimal, and $s'_{\tau_s}(z) \perp z$. Then, at the optimum, the income-gradient of the fine schedule, b, satisfies the following:*

$$b^* = \frac{Cov[\overline{e}'(z)s'_{\tau_s}(z), z] + Cov\left[\hat{g}(z), \int_0^z (y - \overline{z})s'_{het}(y)dy\right]}{\overline{s}'_{\tau_s}\mathbb{V}[z]}$$

where $\hat{g}(z)$ are augmented social welfare weights.[11]

Under conventional welfare weights that are monotonically declining in z, the optimal income gradient of the fine schedule depends essentially on how deterrence effects $\overline{e}(z) * s'_{\tau_s}(z)$ and earnings ability tags $s'_{het}(z) \coloneqq s'(z) - s'_{inc}(z)$ co-vary with income. With respect to deterrence, we consider the special case where the externality $\overline{e}(z) = \alpha_e + \beta_e z$ is affine in z. This implies:

$$b^* = \beta_e + \frac{\text{Cov}\left[\hat{g}(z), \int_0^z (y - \overline{z})s'_{het}(y)dy\right]}{\overline{s}'_{\tau_s}\mathbb{V}[z]}$$

[11]The augmented weight $\hat{g}(z)$ captures the full social value (including fiscal externalities) of a marginal transfer to an individual with a given earnings level z.

In words, an optimally corrective fine gradient (ignoring redistribution) would simply set the fine gradient equal to the slope of marginal speeding costs, $b^* = \beta_e$.[12] In section 5, we estimate the slope of marginal speeding costs β_e in order to test the externality-driven rationale for income-based speeding fines.

With regards to the redistributive term $\frac{\text{Cov}\left[\hat{g}(z), \int_0^z (y-\overline{z}) s'_{het}(y)dy\right]}{\overline{s}'_{\tau_s}\mathbb{V}[z]}$, we refer to the numerator as the "redistributive value:" it captures the social value of using the ability tags $s'_{het}(z)$ to move resources between individuals with different welfare weights $\hat{g}(z)$ through changes in the fine slope b. However, since speeding behavior is a *mutable* tag, the redistributive value is dissipated by a fiscal externality term (*i.e.* a consumption distortion) that is proportional to the average fine semi-elasticity $\overline{s}'_{\tau_s}$.

It can be shown that, when $s'_{het}(z) = \alpha_{het} + \beta_{het} * z$ is affine with $\beta_{het} > 0$, the redistributive term will be positive. In words, if the earnings ability tag $s'_{het}(z)$ is *more* informative at higher income levels, then fines that scale with income ($b^* > 0$) will have an instrumental redistributive value. The relationship between earnings z and the tag $s'_{het}(z)$ is therefore the key metric that informs our test of the redistributive rationale; in section 6, we estimate the earnings effects $s'_{inc}(z)$ that are used to recover the tag $s'_{het}(z)$. To quantify the consumption distortion, we use estimates of extensive-margin fine elasticities from Kaila (2024).

3.3 Incorporating Fairness

We now formalize the fairness rationales we consider and show how they enter our expression for the optimal fine gradient. As discussed in section 2, official argumentation for income-based fines in Finland cites a desire to equalize the "impact" of fines across incomes (Lahti, 2021). One interpretation of this rationale, consistent with recent legal theory (Kantorowicz-Reznichenko, 2013), is that income-based fines are intended to

[12] The intercept term would be captured by the fixed component of the fine schedule: $a^* = \alpha_e$.

achieve "equal marginal deterrence" across incomes, which we define as across-income equality in the schedule of marginal speeding costs (under some utility function). In particular, with log-utility over consumption, fines that are proportional to net-of-tax income (such as day-fines in Finland) generate equal marginal deterrence.

To see this, let $y = z - T(z)$ be net income and $u(c) = \ln\big(y - \tau(y) * s\big)$ be the sub-utility of income; suppose $\tau(y) = b * y$, *i.e.* fines are proportional to net income. Then, the marginal cost of speeding is given by:

$$u'_s\big(y - \tau(y) * s\big) = \frac{-b}{1 - b * s}$$

In words, when $\tau(y)$ is proportional to net-income, the marginal cost of speeding is independent of income y at all speeding levels s. Thus, under log utility, Finland's income-proportional fine policy equalizes marginal speeding costs across incomes; if the latter is the instrumental objective of income-based fines, then preferences over different fine policies will depend on the shape of the fine schedule *per se*

However, with log-utility of consumption, proportional fines *also* generate across-income equality in speeding rates under any monotonic and increasing sub-utility of speeding $v(s)$. As we showed in figure 2.2, the income-based speeding fine system is associated with rough equality of speeding rates across incomes, in contrast to the inequality within the fixed-fine region of the schedule. Thus, income-based fines may instead be an instrument for achieving a particular behavioral outcome (equal speeding rates across incomes), rather than marginal deterrence. Given this ambiguity, and motivated by theories of fairness that emphasize equal allocations/behavior,[13] we consider two separate fairness-driven effects of a given fine policy: "behavioral fairness," which depends on the distribution of behavior induced by a policy, and "fine fairness," which depends only on

[13]For example, specific egalitarianism emphasizes equality in the allocation of certain goods (such as health insurance); it may be natural to extend this to recent empirical evidence (Caspi et al., 2024) provides support for this theory in certain domains.

the fine schedule itself.

To formalize these preferences within our optimal tax framework, we model each rationale in a reduced-form (sufficient-statistic) manner. In particular, suppose there is a prevailing fine schedule $\tau(z) = a^*(b) + b * z$, where $a^*(b)$ is optimal given a slope parameter b. Given this schedule, we define the fine fairness term $\overline{e}_{\tau_s}(z)$ as the social value of a marginal speeding fine decrease for z-earners; similarly, we define the behavioral fairness term $\overline{e}_s(z)$ as the social value of a marginal decrease in the speeding rate among z-earners. The term $\overline{e}_s(z)$ captures local preferences over induced behavior, while $\overline{e}_{\tau_s}(z)$ captures local preferences over the shape of the fine schedule.

Because the behavioral fairness term $\overline{e}_s(z)$ is proportional to behavioral responses $s'_{\tau_s}(z)$, it enters the expression for the optimal fine gradient in precisely the same manner as a standard externality. By contrast, $\overline{e}_{\tau_s}(z)$ produces effects proportional to the size of the fine change; as a result, it enters the optimal fine gradient in a manner similar to the redistributive value. To see this most clearly, suppose that both $\overline{e}_s(z) = \alpha_s + \beta_s * z$ and $\overline{e}_\tau(z) = \alpha_\tau - \beta_\tau * z$ are affine.[14] Then, if $s'_{\tau_s}(z) \perp z$, our expression for the optimal fine gradient becomes:

$$b^* = \beta_e + \beta_s + \frac{-\beta_\tau \mathbb{V}[z] + \mathrm{Cov}\left[\hat{g}(z), \int_0^z (y - \overline{z}) s'_{het}(y) dy\right]}{\overline{s}'_{\tau_s} \mathbb{V}[z]} \tag{3.2}$$

Considering the fairness mechanisms in isolation, if society has a *stronger* desire to decrease speeding rates among the rich ($\beta_s > 0$), the optimal gradient b^* will be positive. Furthermore, if society has a *weaker* desire to lower fines among the rich ($\beta_\tau > 0$), then the optimal gradient will be positive.[15] In our survey, we test whether $\beta_s > 0$ by eliciting subjects' value for policies that lead to more equality in across-income speeding rates (relative to a distribution where high-income individuals speed more); similarly, we test

[14]We normalize both β slope terms so that positive values correspond to a positive optimal fine gradient.

[15]This follows from the fact that $\overline{e}_\tau(z) = \alpha_\tau - \beta_\tau * z$ and fine-fairness enters the optimal gradient as $\frac{-\beta_\tau}{\overline{s}'_{\tau_s}}$. Because $\overline{s}'_{\tau_s}$ is negative, $\beta_\tau > 0$ implies the result.

whether $\beta_\tau > 0$ by varying the "steepness" of the fine schedule (*i.e.* how much larger the fine is at high incomes relative to low incomes).

Appendix equation 10.3 presents a generalized version of equation 3.2 when, rather than assuming $s'_{\tau_s}(z) \perp z$, we assume a constant fine elasticity ζ_τ; this is the equation we use to derive the optimal redistributive fine slope in section 6. There are two differences relative to equation 3.2. First, $\overline{z}$ is replaced by a consumption-weighted average earnings level $\overline{z}^{rw}$; second, the fine-fairness intercept α_τ appears in the optimal gradient. While our survey design described in section 7 only allows us to estimate the slope term β_τ, this will be sufficient to determine the sign that fine-fairness contributes to the optimal gradient b^* if α_τ is not *too* large (see assumption 10.1 in the Appendix for details).[16] Since our survey respondents are relatively unresponsive to randomized fine-schedule steepness (and, if anything, slightly dislike steeper schedules), we focus our discussion on the slope term β_τ.

We now discuss the administrative data sources that we use to estimate the marginal speeding cost gradient β_e and causal earnings effects on speeding $s'_{inc}(z)$.

4 Administrative Data

We now briefly describe the administrative data we use to assess the Beckerian and redistributive rationales for income-dependent speeding fines.

Our primary dataset is an annual panel of earnings, speeding-ticket receipt, and accident costs from 2006-2019. We construct this panel using a population register, income-tax register, crime reports, and accident reports. We use taxable earned income as our earnings measure, though we also make use of information on transfers received and taxes paid from the tax returns. Using the crime reports, we construct a 0-1 indicator for

[16] The claim that β_τ is determinative under assumption 10.1 is dependent on the joint distribution of fines, speeding rates, and earnings. Essentially, if $s(z)/\tau(z)$ is weakly increasing in z (as is the case in Finland when we consider overall speeding behavior and its associated fines), then this claim will hold under any constant fine elasticity.

speeding-ticket receipt in a given year as our measure of speeding behavior; from this data, we also construct an indicator for receipt of a non-speeding traffic ticket. We restrict this panel to individuals who are 18 years or older in a given year.

The accident report data is crucial for testing whether the marginal cost of speeding varies with income. To measure accident costs for a given "offender" i in year t, we count the number of *injuries* (both fatal and non-fatal) incurred by "externality-relevant" participants in accidents involving i in year t. In our baseline results, we assign all participants *outside* of i's vehicle for a given accident to the externality-relevant group; implicitly, we assume that the driver i internalizes the cost to everyone in their own vehicle. We also consider an alternative definition where all other participants (except the driver i) constitute the externality-relevant group; qualitatively, this does not change our results.

We supplement this primary panel dataset using several additional registers. First, we make use of the vehicle register (2013-2023), which we use to measure each individual's annual driving distance from vehicle-inspection odometer information (see Appendix section 10.2.1 for details); this measure serves as an important control in our estimate of the marginal speeding cost gradient. We also draw on the inheritance tax register (2013-2022), the universe of parent-child links, and the cause-of-death register (2013-2022) to construct inheritance-shock events around first parental deaths, which we use to measure income effects on speeding.

5 Externalities

We now discuss our approach for estimating the income-gradient of marginal speeding externalities, β_e, in Finland; in the absence of other concerns, value of β_e allows us to directly evaluate the externality-based rationale via equation 3.2. Our approach to measuring β_e has two key features. First, using linked data from accident reports, crime reports, and income tax returns, we exploit cross-sectional variation in speeding ticket receipt-rates (conditional on a set of controls) to measure the marginal cost of speeding with respect to injuries. Second, we leverage a "job-loss" design to calculate the monetary cost of injuries (both fatal and non-fatal) that result from accidents. We begin with the cross-sectional approach.

5.1 Cross-Sectional Approach

We measure the cost of speeding using accident-related injuries from our accident reports data. To measure marginal speeding costs, we use cross-sectional variation in *observed* speeding behavior conditional on income.[17] In particular, we collapse our panel of earnings, speeding-ticket receipt, and externality-relevant injuries (see section 4) into an average individual-level dataset[18] with the following key variables:

- $\overline{z}_i$: average taxable earned income
- $\overline{s_i}$: average speeding rate, *i.e.* (# of years with speeding ticket/# of panel years)
- $\overline{I}_i^{ext}$: average number of injuries incurred by externality-relevant parties in accidents involving driver i

As discussed in section 4, we take as our baseline externality-relevant group individuals who are outside the focal driver's vehicle (hereafter, "other-party" victims); our results are qualitatively similar if we instead designate all other accident participants besides the driver (hereafter, "other-person" victims) as the externality-relevant group. For individuals who are never in an accident that causes an externality-relevant injury, $\overline{I}_i^{ext}$ will be equal to 0. Summary statistics for this dataset are available in Appendix table 10.2.

[17] Note that, because we implicitly take enforcement as given, we require behavior-driven variation in *observed* speeding behavior (*i.e.* ticket receipt) to measure the appropriate externality. Since we do not observe enforcement probabilities $p(z)$, it is difficult to make use of outside estimates for speeding externalities (van Benthem, 2015).

[18] We do not balance the panel before collapsing since missing status is endogenous (e.g. if an individual dies in an accident). We therefore weight our regressions by the number of years in which i appears in the population register from 2006-2018.

We distinguish between β_e^I, the gradient of marginal speeding costs measured in *injuries*, and β_e^M, the *money-metric* gradient that we need to test the externality rationale. Since we measure the cost of each injury using a job-loss design, we first establish results for the injury gradient β_e^I before turning to money-metric estimates. To estimate β_e^I, we use the following regression:

$$\overline{I}_i^{ext} = \rho_e^I * \overline{z}_i + \alpha_e^I * \overline{s}_i + \beta_e^I * \overline{z}_i * \overline{s}_i + \kappa * \overline{X}_i + \gamma_{a(i)} + \epsilon_i \tag{5.1}$$

where $\gamma_{a(i)}$ is a set of birth-cohort fixed-effects and $\overline{X}_i$ is a set of controls. The coefficient ρ_e^I measures the injury cost-gradient for non-speeders, while α_e^I measures the intercept of marginal social costs (corresponding to the constant in the optimal fine slope); our primary object of interest is β_e^I, which measures how the marginal effect of speeding varies with income.

We employ this cross-sectional approach for two primary reasons. First, due to the stochastic nature of ticket receipt, the average speeding rate $\overline{s}_i$ is a better representation of underlying behavior than ticket receipt $s_{i,t}$ in a given year, thereby mitigating attenuation concerns. Second, since accidents are rare, cross-sectional variation provides more powerful tests than behavioral variation induced by price instruments. Our approach relies on the following identification assumption:

$$\mathbb{E}[\overline{s}_i * \epsilon_i | \overline{z}_i, \overline{X}_i] = 0$$

In words, our measure of observed speeding behavior $\overline{s}_i$ must be uncorrelated with other determinants of accident costs conditional on income and covariates $\overline{X}_i$. However, we do allow for across-income differences in characteristics that affect marginal speeding costs. For example, if income is positively correlated with vehicle mass, the marginal cost of speeding may be positive correlated with income because speeding with a heavier car can lead to more severe accidents. In other words, our goal is to measure a "second-best" gradient of marginal speeding externalities, allowing some unpriced margins to affect the income-gradient of marginal speeding costs.

In our preferred specification, we include the following controls in $\overline{X}_i$:

- Average car ownership rates (for both an individual and their spouse), collected from the population register
- Average annual driving distance, measured using successive odometer readings from vehicle inspections
- Average non-speeding traffic-crime rates, constructed from crime reports
- Own accident-related injuries and involvement in non-injurious accidents

Car ownership serves as a baseline proxy of driving distance (and therefore accident exposure); we supplement this broadly available variable with an often-missing measure of annual driving distance constructed using odometer readings. We control for non-speeding traffic-crime rates to absorb a driver's "general riskiness," which may be correlated with their speeding behavior. Given the cross-sectional nature of our main dataset, we also control for each individual i's *own* accident-related injuries and their involvement in non-injurious accidents, both of which address reverse causality concerns.[19] As an additional safeguard against reverse causality concerns, we drop from our sample any individuals involved in an accident from 2003-2005, prior to the 2006-2018 period we use to create our cross-sectional dataset.

5.2 Marginal Speeding Gradient (Injuries)

In this section, we present our estimates of the marginal speeding cost gradient in terms of injuries (equation 5.1). We also discuss potential mechanisms for our baseline findings.

[19] As we show with our job-loss design, accident-related injuries lead to earnings losses (and may have an impact on how likely someone is to subsequently speed). This can confound the relationship between income and marginal speeding costs.

Figure 5.1: Expected Accident-Related Injuries, by Income and Speeding History

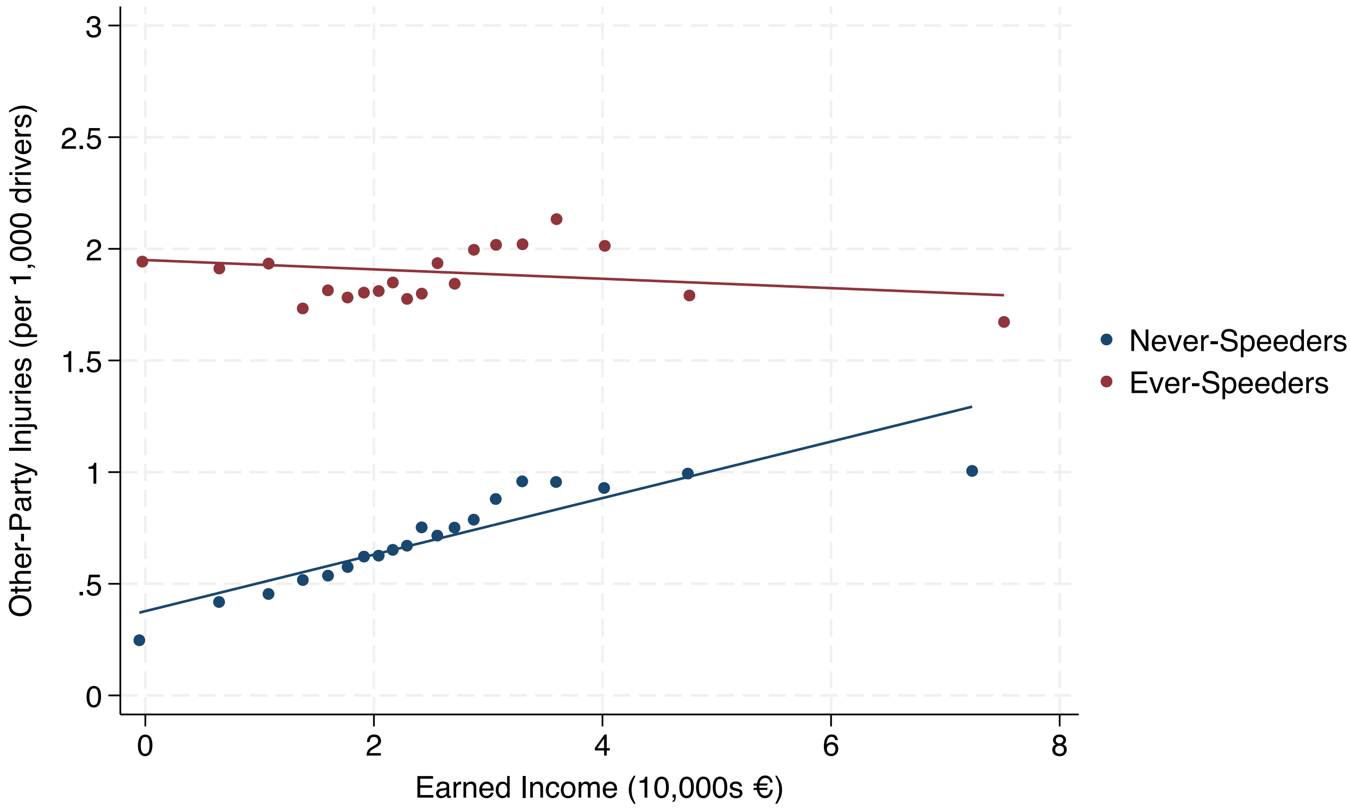

Binscatter relationships between average annual earned income and the average number of accident-related injuries incurred by other-party victims in a given year, using the cross-sectional dataset described in section 5.1. Other-party injuries are injuries incurred by accident participants outside of the focal individual's vehicle; we multiply other-party injuries by 1,000 for interpretability. "Never-Speeders" are those who never receive a speeding ticket from 2006-2018, according to crime-history reports; "ever-speeders" are those who receive at least one speeding ticket from 2006-2018. Earned income and other-party injuries are residualized on birth-cohort fixed effects; furthermore, we weight each individual i by the number of years in which i appears in the population register from 2006-2018.

To motivate our main results, figure 5.1 shows binscatter relationships between average earnings $\overline{z}_i$ and expected other-party injuries $\overline{I}_i^{ext}$, separately for those who never receive a speeding ticket from 2006-2018 ("never-speeders") and those who receive at least one ticket in that time-frame ("ever-speeders"). Despite a positive income-gradient of expected costs among non-speeders, the gradient among speeders is nearly flat, implying that the marginal cost of speeding *decreases* with earnings.

We now turn to our estimates of equation 5.1, which are displayed in table 5.1. To begin, in column 1, we include only the reverse-causality controls. Reassuringly, the baseline speeding effect on injuries (row 1) is positive; consistent with figure 5.1, the gradient (row 2) is negative, implying that the marginal effect of speeding on the likelihoood of an external injury decreases by 0.07 percentage points for every €10,000 increase in driver earnings. Column 2 conditions just on individuals who *ever* receive a speeding ticket; while the marginal speeding gradient decreases somewhat, it is clear that the negative gradient in column 1 is not entirely driven by comparisons between never- and ever-speeders. In column 3, motivated by concerns that differential across-income enforcement might explain our results, we include fixed effects for the municipality of residence;[20] however, compared to column 1, our results are nearly unchanged. Column 4 adds our car ownership and non-speeding traffic-crime controls (as well as interactions between these variables and earnings).[21] The gradient attenuates by around 36% relative to column 1, but remains negative.

Column 5 of table 5.1 adds a control for average annual driving distance, which may be correlated with speeding rates and itself affects accident exposure. To construct annual driving distance, we use odometer data from successive vehicle inspections that occur between 2013 and 2018.[22] Due to the infrequency of inspections, this variable is often missing for a given driver-year; therefore, in column 5, we weight each individual by the number of years for which annual driving distance is non-missing. While some of the difference is selection-driven, the distance control attenuates the marginal cost gradient relative to column 4 by more than 50%, suggesting that driving distance is a consequential omitted variable. Therefore, column 5 represents our preferred estimate of the marginal speeding cost gradient, implying that the marginal effect of speeding on the likelihoood of an external injury decreases by 0.02 percentage points for every €10,000 increase in driver earnings. In Appendix table 10.3, we replicate table 5.1 using other-person victims (*i.e.* including individuals in the driver's vehicle) as the externality-relevant group; we find somewhat more negative marginal speeding gradients, consistent with a correlation

[20] Specifically, since this is a cross-sectional dataset, we include fixed-effects for the first and last municipality in which an individual resides during the 2006-2018 period.

[21] Coefficient estimates for these variables are available upon request.

[22] See Appendix section 10.2.1 for details.

Table 5.1: Income Gradient of Marginal Social Costs (Injuries)

	(1) Other-Party Injuries (per 1,000 drivers)	(2)	(3)	(4)	(5)
Speeding Rate (MSC Intercept)	8.484***	7.680***	8.374***	2.883***	1.370***
	(0.163)	(0.279)	(0.164)	(0.168)	(0.217)
Speeding X Earnings (MSC Gradient)	-0.741***	-0.590***	-0.742***	-0.476***	-0.219***
	(0.0380)	(0.0605)	(0.0380)	(0.0391)	(0.0513)
Earned Income (10,000s €)	0.122***	0.0849***	0.131***	0.106***	0.138***
	(0.00411)	(0.0113)	(0.00422)	(0.00608)	(0.0242)
Constant	0.304***	0.588***	0.289***	-0.0841***	-0.0373
	(0.0104)	(0.0455)	(0.0105)	(0.0135)	(0.0783)
N	5033597	1588308	5033597	5033597	2223285
Cohort FE	X	X	X	X	X
Reverse-Causality Controls	X	X	X	X	X
Ever-Speeders Sample		X			
Municipality FE			X		
Car Ownership & Other Traffic Crime Controls				X	X
Driving Distance Control					X

Estimates of equation 5.1 using the cross-sectional dataset described in section 5.1. Other-party injuries are injuries incurred by accident participants outside of the focal individual's vehicle; we multiply other-party injuries by 1,000 for interpretability. An individual i's "speeding rate" is defined as the number of years from 2006-2018 in which i receives a speeding ticket, divided by 13; earned income refers to an individual's average taxable earned income from 2006-2018. Column 2 restricts to individuals who receive at least one speeding ticket from 2006-2018; column 3 includes fixed-effects for municipalities of residence in 2006 and 2018. For details on the construction of control variables, see section 5.1. The car-ownership rate, non-speeding traffic-crime rate, and driving-distance controls are also interacted with earnings, when included. In columns 1-4, we weight each individual by the number of years in which they appear in the population register from 2006-2018; in column 5, we use a compound weight which is the product of our baseline weight and the number of years for which driving distance information is available from 2013-2018. Heteroskedasticity-robust standard errors are in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

between income and car safety.

Appendix table 10.4 uses information from accident reports to explore the mechanisms underlying our results. In particular, we show that, *conditional on an accident*, the negative marginal speeding cost gradient dramatically attenuates and becomes insignificant when we control for the number of participants involved in an accident.[23] While we caution that this conditional analysis cannot be used to explain the full unconditional cost gradient in table 5.1, our results suggest that high-income drivers tend to drive on busier roads (e.g. during rush hour), where *other* drivers' behavior plays a larger relative role in the severity of an accident.[24]

Taking stock, our results suggest that optimal externality-motivated fines may actually be *decreasing* in income, under the presumption that the cost of an external injury is unrelated to the driver's income. However, to motivate why such costs may not be constant, Appendix figure 10.5 uses accident-level data from the accident reports to plot the relationship between driver income and victim income, both for "other-party" victims (panel 10.5a) and "other-person" victims (panel 10.5b). We see modest positive relationships in both figures, implying that higher-income drivers are involved in accidents with higher-income victims; to the extent that injury costs are correlated with income (e.g. due to lost labor time), this could weaken or even flip the negative marginal speeding cost gradient. More generally, victims' injury costs could be correlated with driver income for a number of reasons, e.g. if the injuries resulting from accidents with high-income drivers are more severe. We therefore turn to a job-loss design (von Wachter and Bender, 2006) that enables us to measure the cost of individual injuries. Importantly, this ap-

[23]Compare column 1 to column 2 of table 10.4. In unreported results, we examine a number of alternative controls in this conditional sample, including municipality of the accident, prevailing speed limit, collision type (E.g. head-to-head collision, side-to-side collision, etc.), and vehicle type. We find a persistently negative and significant gradient across these controls; among these, only collision-type attenuates the baseline conditional gradient from column 1 of table 10.4 (by around 42%).

[24]In Appendix table 10.4, we also see that the positive relationship between earnings and external injuries among non-speeders (row 3) effectively disappears when we control for the number of accident participants. This substantiates our claim that traffic density explains our results: high income drivers *in general* (regardless of their speeding behavior) seem to drive on busier roads, leading to accidents that involve more participants.

proach allows us to separately estimate fiscal externalities (the injury costs borne by the government) and net-income losses (the costs borne by victims); this separation allows us to apply income-dependent weights to net-income losses.

5.3 Money-Metric Costs of Injuries

Our approach to measuring the monetary costs of injuries relies on comparisons between individuals who are injured or killed in an accident and matched individuals from the general population. In particular, for each individual injured or killed in an accident involving a motor vehicle (the treatment group), we assign two random control individuals from the population register with the same birth year, sex-at-birth, and education level in the year prior to the accident.[25] We then construct an event-time panel for each match-group around the accident event year, following individuals until 2023 (the last year in which we have detailed income data).[26]

To illustrate our approach, we begin by estimating event-study models of the form:

$$Y_{i,t,k} = \beta_0 * \text{Treat}_{i,t} + \gamma_k + \sum_{k=-4,k\neq-1}^{2023-t} \beta_k * \text{Treat}_i * I(t+k \geq t) + \epsilon_{i,t,k} \tag{5.2}$$

where $Y_{i,t,k}$ is an outcome for individual i in year $t+k$, $\text{Treat}_{i,t}$ is indicator for whether or not individual i is treated in year t, $I(t+k \geq t)$ is a post-period indicator, and γ_k is a set of event-time dummies.[27] Figure 5.2 shows event-study coefficients from equation 5.2, using disposable income (*i.e.* income net of taxes and transfers) as the outcome. We see no discernible pre-trends and an immediate income loss in the year of the accident, followed by a gradual divergence between the treatment and control groups. To assess how fatal injuries influence our results, Appendix figure 10.9 displays event-study esti-

[25]We use four education categories (based on highest completed qualification): less than upper-secondary, upper secondary or post-secondary non-tertiary, short-cycle tertiary, and bachelor's or above.

[26]We have reliable accident report data going back to 2003.

[27]Due to the nature of our data, in which we lose one accident year in each successive post-period, the event-time dummies are equivalent to including accident-year fixed-effects for the post-period.

Figure 5.2: Event Study Effects of Accident-Related Injury/Death on Disposable Income

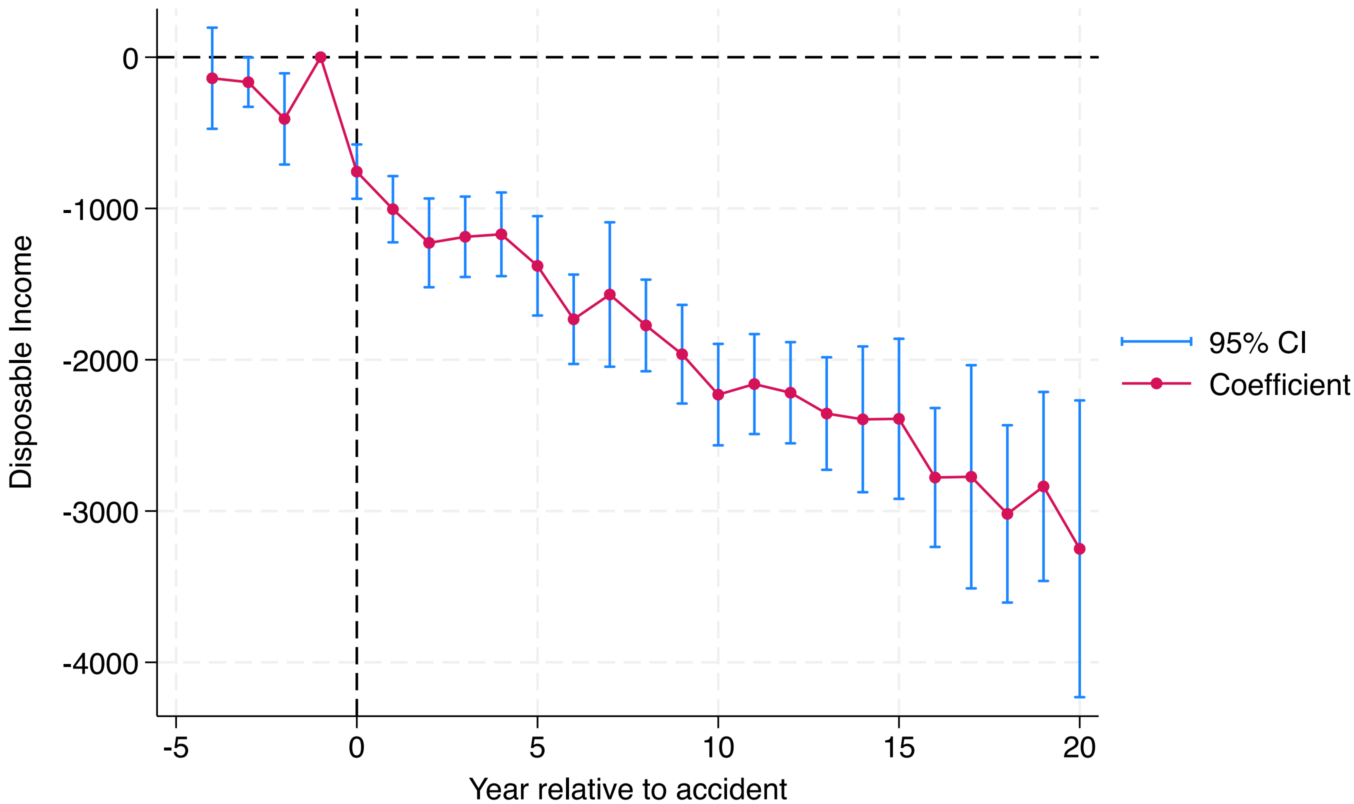

Estimates of β_k from equation 5.2, using disposable income as the outcome variable; disposable income is the sum of earned income, capital income, and transfers received, net of taxes paid. For details on sample construction, see section 5.3. We restrict to accident events that involve a motor vehicle (cars, trucks, buses, vans, mopeds, motor bikes, and motorcycles); injuries from accidents involving a snow mobile, trailer, tram, train, tractor, or other unidentified motor vehicles are excluded. We also restrict to individuals who are at least 18 years old in the year prior to the accident.

mates using only non-fatal injuries; we see little-to-no decline in disposable income for several years following the accident, after which we see the same divergence pattern as figure 5.2 emerge.[28]

To account for the gradual divergence between treatment and control groups in figure 5.2, we show in Appendix figure 10.7 that around 4% of individuals who are non-fatally injured drop out of the labor force within a few years of the accident and never return.[29]

[28]In unreported results, we show that, for non-fatal injuries, income before transfers and taxes declines immediately in the year of the accident, implying that the initial lack of decline in *net* income is driven by insurance provided by the tax-and-transfer system.

[29]We define labor-force non-participation as being employed for zero months during the year.

Since the average injured individual is in their prime earnings-growth years (around 45 years old), it is unsurprising that income gradually diverges between the treatment and control group.

While our findings on disposable income effects in Figure 5.2 are indicative of the consumption loss that victims experience, the government (through increased transfer payments and reduced tax receipts) also bears a cost from accident-related injuries. To compare these two cost margins, figure 5.3 plots event-study treatment effects from equation 5.2 separately for pre- and post-tax income; the red-shaded region corresponds to the losses in 5.2, while the blue-shaded region measures the replacement income (fiscal externality) paid out by the government.[30] In our main sample that combines fatal and non-fatal injuries, disposable income losses dominate (especially in later years); in Appendix figure 10.10, we show that, for non-fatal injuries, fiscal externalities are dominant until well after the treatment event.

Motivated by these findings, we calculate the cost of each accident-related injury in our data. To do so, we start by calculating net-present-value (NPV) fiscal externalities and disposable income losses. In particular, for each injury and post-period, we calculate difference-in-differences effects relative to the two matched control individuals.[31] We then aggregate period-specific effects using a discount rate δ (equal to 3% in our baseline) to obtain NPV costs. The results of this aggregation are shown in Appendix table 10.5, separately for injuries and deaths. Under our baseline discount rate of 3%, non-fatal injuries generate around €10,600 in NPV costs to the government and €6,500 in NPV net-income losses to individuals, based on a 13-year treatment horizon. For the small number of fatal injuries, we find cost *savings* for the government, though these are swamped in magnitude by nearly €200,000 of NPV net-income losses through 13 years. Unsurprisingly, given that injuries cause permanent declines in labor supply (Appendix

[30] For a decomposition of this fiscal externality into transfer payments and tax receipts, see Appendix figure 10.8.

[31] For additional details, see Appendix section 10.2.2.

Figure 5.3: Event Study Effects of Accident-Related Injury/Death on Components of Social Costs

Estimates of β_k from equation 5.2, using disposable income and pre-tax income as outcomes. Disposable income is the sum of earned income, capital income, and transfers received, net of taxes paid; pre-tax income is the sum of capital and earned income. For details on sample construction, see section 5.3. We restrict to accident events that involve a motor vehicle (cars, trucks, buses, vans, mopeds, motor bikes, and motorcycles); injuries from accidents involving a snow mobile, trailer, tram, train, tractor, or other unidentified motor vehicles are excluded. We also restrict to individuals who are at least 18 years old in the year prior to the accident.

figure 10.7), we find in Appendix figure 10.11 that estimated fiscal externalities and disposable income losses are increasing with the pre-period income of victims.

To aggregate victim-specific cost components, we apply welfare weights g_i (Saez, 2002b) to each individual's NPV net-income losses:

$$g_i = \frac{\kappa}{y_i^{-\nu}}$$

where y_i is i's disposable income in the year prior to the accident and ν is a CRRA parameter that encodes the planner's distaste for inequality. Under this specification of welfare weights, the planner weighs money-metric losses to each individual in proportion to the marginal utility from a CRRA utility function; the larger is ν, the larger is the weight placed on a given monetary loss to low-income individuals relative to high-income individuals. We set κ such that $\overline{g} = 1$, meaning that the planner is indifferent between a €1 increase in revenue and distributing €1 equally among the population. The weights g are therefore measured relative to government revenue, implying that we can leave fiscal externalities unweighted; as a result, we calculate the total cost of a given injury as the sum of NPV fiscal externalities and welfare-weighted NPV net-income losses. In our baseline results, we set $\nu = 1$, though we assess robustness to higher and lower degrees of inequality aversion.

In section 5.4 below, we import these individual-specific cost estimates into our cross-sectional dataset from section 5.1 to estimate the money-metric gradient of marginal speeding costs. For this context, we note two important features of our cost estimates. First, we do not project injury-related costs forward using our job-loss design. Within our cross-sectional design, this choice will not bias our estimate of the money-metric gradient β_e^M as long as accident year (which determines the treatment horizon in our job-loss design) is not correlated with offender income;[32] indeed, we find that the corre-

[32]In particular, while we understate the marginal social cost of speeding by failing to project the cost of injuries, our

lation between accident year and earnings is quantitatively negligible. Second, we elect to use our job-loss estimates for the cost of fatal injuries, rather than relying on outside estimates for the value of a statistical life (VSL). We favor this approach so that the costs of fatal and non-fatal injuries are determined on a similar basis: as long as our approach appropriately captures the *relative* severity of fatal and non-fatal injuries, our results are unlikely to be biased.[33]

Robustness To claim our findings on the cost of accident-related injuries as causal, we rely on the standard parallel-trends assumption, *i.e.* that, in the absence of an accident-related injury, the treatment group would have progressed on the same trajectory as the control group. While we do not detect any significant pre-trends across various outcomes estimated using equation 5.2, one concern is that negative shocks *specific to the year of the accident* may simultaneously cause both an accident and an earnings decline. For example, an individual who loses their job might experience an earnings decline and, as a result, begin to drive more recklessly. To assuage this concern, we estimate equation 5.2 separately for drivers and passengers: if an earnings decline causes drivers to become reckless, we would expect to see stronger effects for drivers than passengers. Appendix figure 10.6 shows the results for drivers in panel 10.6a and passengers in panel 10.6b: overall, the two trajectories are extremely similar. Based on this evidence, simultaneous negative shocks are unlikely to confound our estimates.

5.4 Marginal Speeding Gradient (Money-Metric)

To estimate the money-metric gradient of marginal speeding costs, β_e^M, we merge our estimates for the welfare-weighted NPV monetary costs of each injury from section 5.3 into our cross-sectional dataset from section 5.1, replacing external accident-related in-

estimate of the gradient is unaffected if the level of understatement is uncorrelated with offender income.

[33]VSL estimates likely capture much more than mere earnings losses; therefore, if we use VSL estimates for fatal injuries, our cost estimates for non-fatal injuries would likely be too small in relative terms.

juries $\overline{I}_i^{ext}$ with its money-metric equivalent $\overline{M}_i^{ext}$.[34] We then estimate the money-metric gradient of marginal speeding costs, β_e^M, using the following regression:

$$\overline{M}_i^{ext} = \rho_e^M * \overline{z}_i + \alpha_e^M * \overline{s}_i + \beta_e^M * \overline{z}_i * \overline{s}_i + \kappa * \overline{X}_i + \gamma_{a(i)} + \epsilon_i \tag{5.3}$$

Table 5.2 displays estimates of equation 5.3 using a 3% discount rate and no welfare-weighting (*i.e.*, $\nu = 0$) to aggregate individual-specific NPV costs; we display results separately for other-person victims (columns 1 and 2) and other-party victims (columns 3 and 4). Columns 1 and 3 present estimates with only our reverse-causality and birth-cohort controls; in these specifications, we estimate baseline marginal speeding costs (MSC intercept α_e^M) of €169 using other-person victims and €118 using other-party victims. We note that the former is nearly identical to the average fixed speeding fine in our police ticket data (approximately €167), suggesting that our empirical approach does a reasonable job of capturing the marginal cost of speeding. Columns 2 and 4 of table 5.2, add the rest of our preferred controls: car ownership rate, non-speeding traffic-crime rate, and average annual driving distance, as well as interactions between these variables and earnings. These additional controls attenuate the MSC intercept term (row 1) but do little to change the marginal speeding cost gradient (row 2). Therefore, for other-party victims (our preferred externality-relevant group), our estimates in column 4 indicate that fines should *decrease* by almost €19 for every €10,000 of income, well below our estimated fine-schedule gradients for both overall speeding crime (+12.4 €) and day-fine speeding crime (+88 €) from section 2.

To see how different discount rates and profiles of welfare-weights affect this baseline result, we re-estimate the marginal social cost gradient β_e^M with our preferred set of controls under varying discount rates $\delta \in \{0.03, 0.05, 0.07\}$ and inequality aversion pa-

[34] In practice, the individual-specific $\overline{M}_i^{ext}$ are estimated. However, beyond adjusting standard errors for heteroskedasticity, we do not account for this estimation uncertainty when estimating equation 5.3; this introduces additional noise and inflates our standard errors. However, as we will see, across nearly all specifications, our standard errors are precise enough to rule out meaningfully positive estimates for β_e^M.

Table 5.2: Income Gradient of Money-Metric Marginal Social Costs: 3% Discount Rate

	Other-Person Costs		Other-Party Costs	
	(1)	(2)	(3)	(4)
Speeding Rate (MSC Intercept)	168.8***	84.39*	118.1***	75.43*
	(25.49)	(34.35)	(22.54)	(31.15)
Speeding X Earnings (MSC Gradient)	-27.19***	-21.45*	-17.54**	-18.60*
	(5.978)	(8.425)	(5.372)	(7.712)
Earned Income (10,000s €)	0.925	0.160	2.100***	2.929
	(0.681)	(4.222)	(0.617)	(3.743)
Constant	6.338***	-10.06	2.189	-21.38
	(1.698)	(13.08)	(1.507)	(11.44)
N	5033597	2223285	5033597	2223285
Cohort FE	X	X	X	X
Reverse Causality Controls	X	X	X	X
Car Ownership & Other Traffic Crime Controls		X		X
Driving Distance Control		X		X

Estimates of equation 5.3 using the cross-sectional dataset described in section 5.1 and estimates of money-metric injury costs described in sections 5.3 and 5.4. Columns 1 and 2 use the cost of accident-related injuries incurred by "other-person" victims as the outcome; columns 3 and 4 use costs incurred by "other-party" victims as the outcome. Other-person injuries are injuries incurred by all accident participants excluding the focal individual; other-party injuries are injuries incurred by accident participants outside of the focal individual's vehicle. An individual i's "speeding rate" is defined as the number of years from 2006-2018 in which i receives a speeding ticket, divided by 13; earned income refers to an individual's average taxable earned income from 2006-2018. For details on the construction of control variables, see section 5.1. The car-ownership rate, non-speeding traffic-crime rate, and driving-distance controls are also interacted with earnings, when included. In columns 1 and 3, we weight each individual by the number of years in which they appear in the population register from 2006-2018; in columns 2 and 4, we use a compound weight which is the product of our baseline weight and the number of years for which driving distance information is available from 2013-2018. Heteroskedasticity-robust standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

rameters $\nu \in \{0, 0.25, 1, 4\}$. Table 5.3 displays the results. First, we note that, regardless of the welfare-weighting scheme or the discount rate, we find consistently negative estimates and can rule out marginal social cost gradients larger than +1 € per €10,000 of earnings. Appendix Table 10.7 shows nearly identical quantitative results when we consider other-person victims. For our preferred estimate (which we use in section 8), we set $\nu = 1$ and $\delta = 0.03$,[35] implying from table 5.3 that marginal speeding costs decrease by €7.6 (se=4.4) per €10,000 of income.

[35] This discount rate reflects the yield on Finnish 10-year government bonds.

Comparing across estimates, we also see in table 5.3 that higher discount rates and stronger inequality aversion (higher ν) lead to more attenuation of the gradient. While the attenuation from the discount rate is mechanical, the role of inequality aversion is consistent with our findings in figure 10.5: since the victims of high-income drivers also have higher incomes and therefore larger money-metric injury costs, the negative injury gradient we measure in section 5.2 is "stretched" in money-metric terms, relative to a case where money-metric injury costs are constant across driver income. Therefore, when the planner down-weights the relative cost of injuries to high-income victims, the gradient increases. We do not consider inequality aversion parameters above $\nu = 4$ because, as shown by Appendix table 10.6, $1 < \nu < 4$ effectively implies a constant weighted cost of injuries across victim incomes (in contrast to the upward-sloping relationship for unweighted injury costs). We view constant across-income costs as a normatively appealing stopping point for consideration.

To summarize, regardless of how we specify the externality-relevant group (other-party or other-person), we cannot identify a positive marginal speeding cost gradient across a range of discount rates and levels of inequality-aversion; furthermore, our estimates are precise enough to rule out meaningfully positive gradients. We therefore reject externality-based motives as a rationale for the Finnish speeding fine system.

6 Redistribution

As elaborated in Section 3, we built our test of the redistributive rationale for income-based fines on the logic of "tagging:" fines serve a redistributive role if speeding behavior is a signal of earnings ability. This signal can be recovered by differencing observed across-income variation in speeding behavior with respect to causal earnings effects on speeding.[36] Our goal is to measure causal earnings effects as an instrument for evaluating

[36]By *causal*, we mean that, for an individual with a given income z, we can construct counterfactual speeding consumption at earnings level $z+dz$. In other words, the variation in earnings dz should be exogenous to other determinants of consumption.

Table 5.3: Income Gradient of Money-Metric Marginal Social Costs (Other-Party Victims), by Discount Rate and Degree of Inequality Aversion

	Discount Rate		
	3%	5%	7%
Degree of Inequality Aversion			
None ($\nu = 0$)	-18.6*	-15.89*	-13.67*
	(7.712)	(6.694)	(5.875)
Low ($\nu = 0.25$)	-16.07*	-13.73*	-11.84*
	(7.477)	(6.504)	(5.717)
Middle ($\nu = 1$)	-7.617	-6.486	-5.568
	(4.392)	(3.839)	(3.393)
High ($\nu = 4$)	-7.55	-6.421	-5.507
	(4.371)	(3.821)	(3.376)
N	5033597	5033597	5033597
Cohort FE	X	X	X
Reverse Causality Controls	X	X	X
Car Ownership Control	X	X	X
Other Traffic Crime Control	X	X	X
Driving Distance Control	X	X	X

Estimates of β_e^M from equation 5.3 using the cross-sectional dataset described in section 5.1 and estimates of money-metric injury costs described in sections 5.3 and 5.4; for details on the inequality-aversion parameter and the discounting procedure, see section 5.3. Other-party victims are accident participants outside of the focal individual's vehicle who incur an injury. For details on the construction of control variables, see section 5.1; the car-ownership rate, non-speeding traffic-crime rate, and driving-distance controls are also interacted with earnings. We weight each individual by the number of years in which they appear in the population register from 2006-2018, multiplied by the number of years for which an individual's driving distance information is available from 2013-2018. Heteroskedasticity-robust standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the tagging potential of speeding behavior.

We begin by providing a decomposition of the causal earnings effect into a labor-supply and income effect, before discussing our empirical approaches for estimating these objects.

6.1 Earnings Effects and Weak Separability

Recall from figure 3.1 that the earnings effect is identified from earnings changes induced by variation in *labor supply*; intuitively, labor-supply induced variation in speeding behavior carries additional information about latent wages that pure income effects (holding labor supply fixed) do not.

To simplify notation, suppose individuals are indexed only by their implicit wage w; given some tax-and-fine system, an individual with wage w chooses a speeding rate $s(z,l)$ as a function of pre-tax income $z = w * l$ and labor supply l. Under weak separability, this function would reduce to $s(z)$: there is no dependence on labor supply except through pre-tax income. However, in the more general case, the derivative with respect to labor supply l, holding the wage w fixed, is given by:

$$\left.\frac{ds(z,l)}{dl}\right|_w = \frac{\partial s(z,l)}{\partial z} * w + \underbrace{\frac{\partial s(z,l)}{\partial l}}_{=0 \text{ under weak sep.}} \tag{6.1}$$

To recover the earnings effect, we re-scale this expression by $\frac{1}{w}$ to obtain:

$$\left.\frac{ds(z,l)}{dz}\right|_w := s'_{inc}(z) = \underbrace{\frac{\partial s}{\partial z}}_{\text{Pre-tax income effect}} + \underbrace{\frac{\partial s}{\partial l} * \frac{1}{w}}_{\text{Re-scaled labor supply effect}}$$

In general, without a weak separability assumption, the earnings effect on speeding is composed of a pre-tax income effect and a wage-scaled (money-metric) labor supply effect. In what follows, we identify the labor supply effect using within-individual variation in labor supply, holding earnings fixed; to estimate the income effect, we turn to an inheritance shock design.

6.2 Labor Supply Effects

Our approach for measuring the labor supply effect is motivated by a test of weak separability: weak separability is violated if, for a given individual, variation in labor supply, holding earnings fixed, has an effect on speeding behavior. To demonstrate this approach, in this section we estimate a single labor supply effect for the general population; to construct local earnings effects $s'_{inc}(z)$, we estimate the labor supply effect separately by income in section 6.4.

To implement our approach, we use our annual population panel from 2006-2019 (see section 4) and model ticket receipt $s_{i,t}$ for individual i in year t as:

$$s_{i,t} = \beta_1 * z_{i,t} + \beta_2 * \overline{w} * l^m_{i,t} + \gamma_i + \kappa_{a(i),t} + \epsilon_{i,t} \tag{6.2}$$

where $z_{i,t}$ is taxable earned income;[37] $\overline{w} * l_{i,t}$ is months employed during the year, re-scaled into a money-metric using the average monthly wage;[38] γ_i is an individual fixed-effect; and $\kappa_{a(i),t}$ are year-by-cohort fixed effects. For demonstrating our approach, we set $\overline{w}$ equal to the 2012 average wage in order to facilitate comparability with our inheritance sample in section 6.4. Under the null hypothesis of weak separability, $\beta_2 = 0$: labor supply has no independent effect on speeding behavior conditional on earnings. Our identification approach for the labor-supply effect β_2 relies on the assumption that, conditional on earnings and cohort-year effects, within-individual variation in labor supply is uncorrelated with other time-varying determinants of speeding behavior. We expect

[37] This measure includes some taxable transfers like unemployment benefits and pensions. Since these transfers insure against within-lifetime income risk, rather than across-individual income risk, we view this measure as corresponding to the static z in our model framework.

[38] Due to co-linearity, we cannot use an individual's imputed wage in this regression. We elect for a cross-sectional average wage since it is well-measured (as opposed to an individual's average wage) and less likely to be endogenous. When we estimate labor supply effects separately by income group, we calculate separate average wages for each group.

that, due to mis-measurement of *true* labor supply, our estimate of β_2 will be attenuated; we view this as unproblematic since it leads to underestimation of earnings effects, thereby biasing us in favor of finding a redistributive role for income-based fines.

Note that the coefficient β_1 in equation 6.2 ostensibly yields an estimate of the income effect, *i.e.* the effect of income on speeding rates holding labor supply fixed. However, due to mis-measured labor supply, variation in earnings $z_{i,t}$ likely includes variation labor supply, leading to bias in the estimated income effect. We discuss the likely direction of this bias below.

Column 2 of table 6.1 displays our estimates from equation 6.2. For comparison, column 1 shows the cross-sectional relationship between earnings and speeding rates: each €10,000 increase in earnings is associated with a 1.2 percentage point increase in the probability of receiving a speeding ticket in a given year. In column 2, we find a significant wage-scaled labor supply effect, leading us to reject weak separability; in particular, a labor supply increase that generates €10,000 of earnings causes a 0.29 percentage point increase in the likelihood of receiving a speeding ticket. Based on this estimate, we can say that the wage-scaled labor supply effect explains around 25% of the overall cross-sectional gradient in column 1.

Given our findings, $\hat{\beta}_1$, the "naive" income effect estimate, is likely upward-biased. Intuitively, some of the earnings variation used to identify β_1 will be driven by labor supply changes that are not perfectly captured by variation in employment months $l^m_{i,t}$; since the labor supply effect is positive ($\beta_2 > 0$), this mis-measurement leads to upward bias. As a result, we turn to an alternative approach for estimating the income effect on speeding behavior.

Table 6.1: Income Gradient of Speeding Behavior & Labor Supply Effect on Speeding

	Speeding Rate (pp)	
	(1)	(2)
Earned Income (€10,000s)	1.166***	0.499***
	(0.00609)	(0.0106)
Re-scaled Labor Supply (€10,000s)		0.288***
		(0.00768)
N	15207808	15172401
R^2	0.019	0.196
Pop. Sample	25%	25%
Cohort-Year FE	X	X
Individual FE		X

Estimates of equation 6.2 using a 25% sample of Finnish adults (*i.e.* individuals 18 years of age or older) that appear in the population register from 2006-2019. The "speeding rate" refers to the share of individuals who receive a speeding ticket in a given year, according to crime-history reports. "Earned income" is defined as taxable earned income, which is the sum of wage/salary income, unemployment insurance payments, and pension income. "Re-scaled labor supply" refers to an individual's months employed during the year, multiplied by the average imputed monthly wage (taxable earned income in 2012 divided by months employed in 2012) among our sample. Standard errors (clustered at the individual level) are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

6.3 Income Effects

Our identification approach for estimating income effects builds on Nekoei and Seim (2022), which leverages the differential timing of parental deaths to explore how inheritances contribute to wealth inequality. Taking this approach off-the-shelf poses a problem: parental deaths may have concomitant effects on speeding that are unrelated to the associated wealth shock, *e.g.* due to grieving or caregiving. As we show, individuals who receive little-to-no inheritance indeed exhibit speeding trends around parental death. We therefore modify the approach in Nekoei and Seim (2022) and difference across inheritance sizes to address non-wealth effects within a triple-difference framework.

First Difference: Timing of Parental Deaths As in Nekoei and Seim (2022), the treatment event in our paper is the first parental death that a child experiences. We identify all such deaths from 2013-2022 (the years for which we have inheritance tax returns) using the cause-of-death register and parent-child links. We compare heirs whose first parental death occurs in year t (treated) to heirs whose first parental death occurs in $t+\delta$ (control). To address observable differences between parental-death cohorts, Nekoei and Seim (2022) apply DFL-reweighting (DiNardo et al., 1996) to homogenize the distribution of education and age across cohorts. We follow a similar approach, re-weighting each cohort with respect to sex, age, and education; since we are interested in population-level estimates, we normalize using the distribution of demographics in the 2012 population register.

In our baseline results, we set $\delta = 4$ and follow individuals for four years ($t = 0, 1, 2, 3$) post-treatment; we examine robustness to using 5 post-periods below. This "fixed-control" method (one control cohort per treatment cohort) generates a clean comparison that is balanced across periods and avoids any contamination from using early-treated observations as controls for late-treated observations. To create our event-time dataset, we construct treatment and control groups for each treatment year $y \in \{2013, 2014, 2015, 2016\}$; we then stack observations across treatment years. Observations are therefore indexed by an individual i, a treatment year y, and an event time t (defined relative to y). We cluster all standard errors at the individual level and apply the weights described above.

Second Difference: Inheritance Size Using a comparison of early- vs. late-treated heirs to measure the income effect on speeding behavior is potentially confounded by non-wealth effects of parental death. In particular, let β^{wealth} be the effect of a given wealth shock on speeding behavior; similarly, let α^{death} be the effect of a first parental death on speeding behavior. A first-difference across parental death years measures the sum of both effects. However, assuming that the death effect α^{death} is common across inheritance size, a second comparison between those who receive *positive* inheritances

$\Delta w > 0$ and those who receive *no* inheritance $\Delta w = 0$ will remove the death effect:

$$\underbrace{\underbrace{(\beta^{\text{wealth}} + \alpha^{\text{death}})}_{\text{DID}|\Delta w>0} - \underbrace{\alpha^{\text{death}}}_{\text{DID}|\Delta w=0}}_{\text{Triple Difference}}$$

To motivate our construction of the $\Delta w > 0$ and $\Delta w = 0$ groups, figure 6.1 shows quintiles of cumulative net-of-tax inheritance flows around the year of first parental death. We see that there are essentially no inheritance flows prior to the first parental death, with the majority of flows occurring within a year of the treatment event.[39] After $t + 1$, we see some additional flows, generated by a combination of above-average tax-processing times and second parental deaths. On average, heirs receive a total €38,000 in the four years following a first parental death. However, the distribution is skewed: the bottom two quintiles receive around €2,300 on-average, while the top quintile receives close to €140,000. To construct the inheritance-size groups in our second difference, we combine the bottom two quintiles (the "bottom" group) and the top three quintiles (the "top" group): we use the bottom group to difference out the common death effect α^{death}. We include the second quintile, which receives around €5,000 on-average, in the bottom group in order to improve power. While, in theory, this can lead to bias if income effects differ across inheritance size, we expect this bias to be vanishingly small.[40]

Estimating Income Effects To motivate our triple-difference approach, we begin by showing how speeding behavior evolves around parental death for each of our inheritance-size groups. In particular, figure 6.2 plots (re-weighted) event study estimates from the following regression:

[39]The estate inventory must be filed with the tax authority within 3 months of the parent's death, after which the tax authority assesses the inventory and makes a tax decision (which usually takes 6-12 months). Therefore, most inheritance decisions are made in the year following the parental death.

[40]Even if the income effects differ between groups, we show in Appendix section 10.2.3 that the bias disappears as $\frac{\Delta W^{\text{top}} - \Delta W^{\text{bottom}}}{\Delta W^{\text{top}}} \to 1$; given that, on-average, the bottom group receives around €2,300 and the top group around €63,000, we expect this bias is negligible.

Figure 6.1: Cumulative Inheritance Flows Around First Parental Death

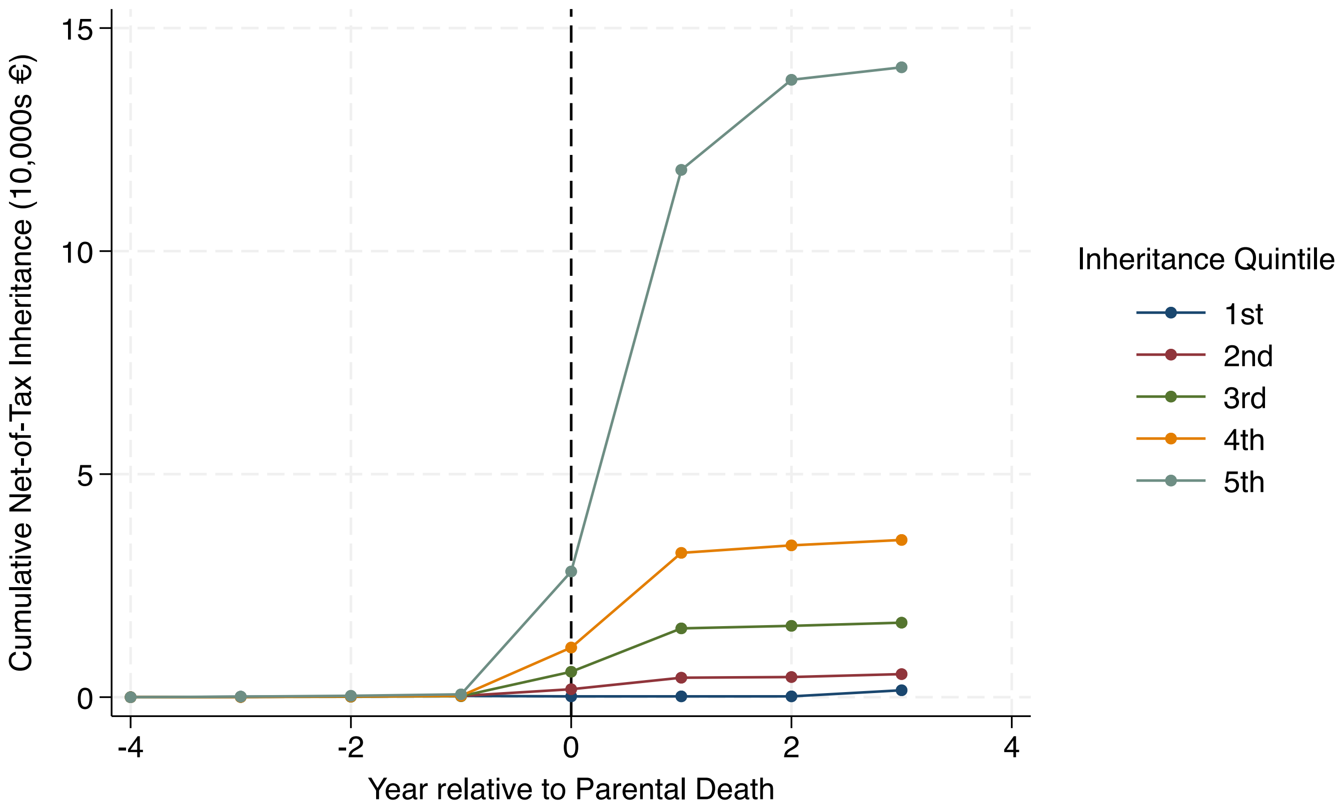

Cumulative net-of-tax inheritance flows around the year of first parental death for children whose first parent dies between 2013 and 2020, separately by quintile of flows; cumulative flows are normalized to 0 in $t-5$. Net-of-tax inheritances refer to the monetary value of the inheritance received—according to the inheritance-tax return—net of inheritance taxes due. Heirs are allocated to inheritance-flow quintiles within each year-of-death cohort d based on the cumulative inheritances received in $[d, d+2]$ (*i.e.* using 3 years of post-death data). We restrict to children who are 18 years of age or older in the year of the parental death event.

$$s_{i,y,t} = \gamma_t + \rho * I_{YOD(i)=y} + \sum_{k=-3}^{4} \beta^k * I_{YOD(i)=y} * I_{t=k} + \epsilon_{i,y,t} \tag{6.3}$$

where $s_{i,y,t}$ is an indicator for whether i received a speeding ticket in year $y + t$, γ_t are event-time fixed effects, $I_{YOD(i)=y}$ is an indicator for whether i's first parental death occurred in year y, and β^k is the treatment coefficient for event year k. We use $t - 4$ as the omitted year in order to show anticipation effects, which appear for both inheritance-size groups in year $t - 1$ of figure 6.2. Following this anticipation effect, the two groups diverge, with the top group exhibiting higher speeding rates.

Given that the common anticipation effect in $t - 1$ is present in the bottom inheritance group, which receives little-to-no-inheritance, this spike is unlikely to be driven by anticipated wealth shocks. To bolster this claim, Appendix figures 10.12 and 10.13 show that there are no anticipatory effects on earnings or car ownership. Therefore, this common anticipation effect likely reflects non-wealth effects (e.g. travel for care-giving purposes) that we want to difference out of our income-effect estimates. To this end, figure 6.3 plots the difference in event-study coefficients, $\beta^t_{\text{top}} - \beta^t_{\text{bottom}}$, from figure 6.2 for each event year t. While the series is somewhat noisy, we see a clear level difference between the pre- and post-treatment periods, as well as the absence of statistically significant pre-trends.

In our estimate of the income effect, we use the post-treatment effects in figure 6.3 as the reduced-form in a triple-difference IV approach. To measure the denominator (first-stage), we use the cumulative inheritances received in the post-treatment period. We view this approach as conservative: implicitly, it assumes that the extent to which the inheritance wealth shock will be "spent" on speeding fines is confined to our four-year treatment horizon. However, if consumption responses are positive and persistent, our estimate of the income effect would be biased downwards. While we are limited in our ability to explore longer treatment horizons, we do find that our income effect is somewhat larger with a 5-year horizon. However, the 4-year horizon yields more power for

Figure 6.2: Event Study Effects of First Parental Death on Speeding Rates, by Inheritance-Size Group

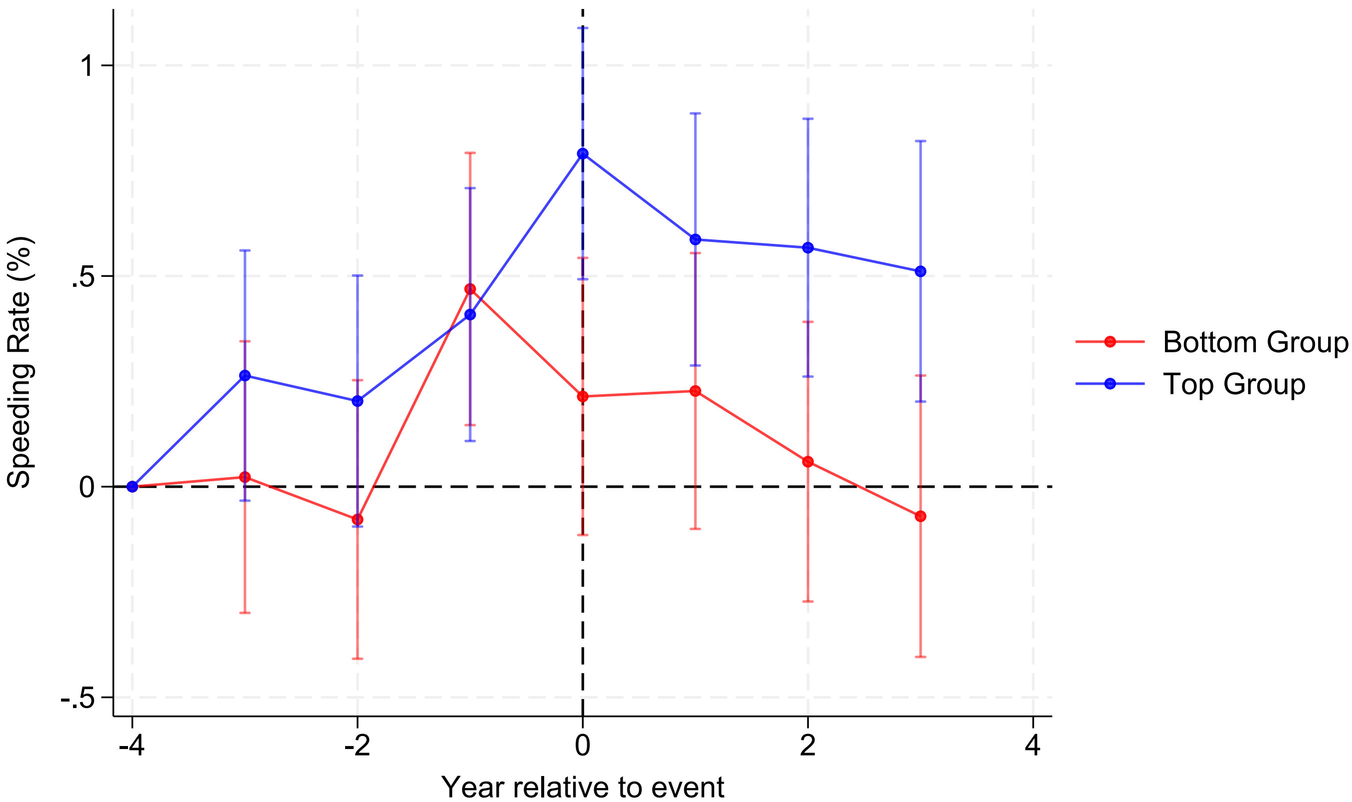

Estimates of β^k from equation 6.3, separately by inheritance-size group. The treatment group consists of heirs whose first parent dies in [2013,2016]; the control group consists of heirs whose first parent dies in [2017,2020]. The bottom inheritance group, for a given parental death-cohort d, consists of heirs in the bottom two quintiles of total inheritance flows received in $[d, d+2]$; the top inheritance group for a given death-cohort consists of heirs in the top three quintiles of total inheritance flows received in $[d, d+2]$. We restrict to heirs who are 18 years-of-age or older in the year of the parental death event. Each death-cohort-by-inheritance-size group is reweighted to match the 2012 distribution of age, sex, and education in the population register. Standard errors are clustered at the individual level.

Figure 6.3: Event Study Effects of First Parental Death on Speeding Rates: Triple Difference

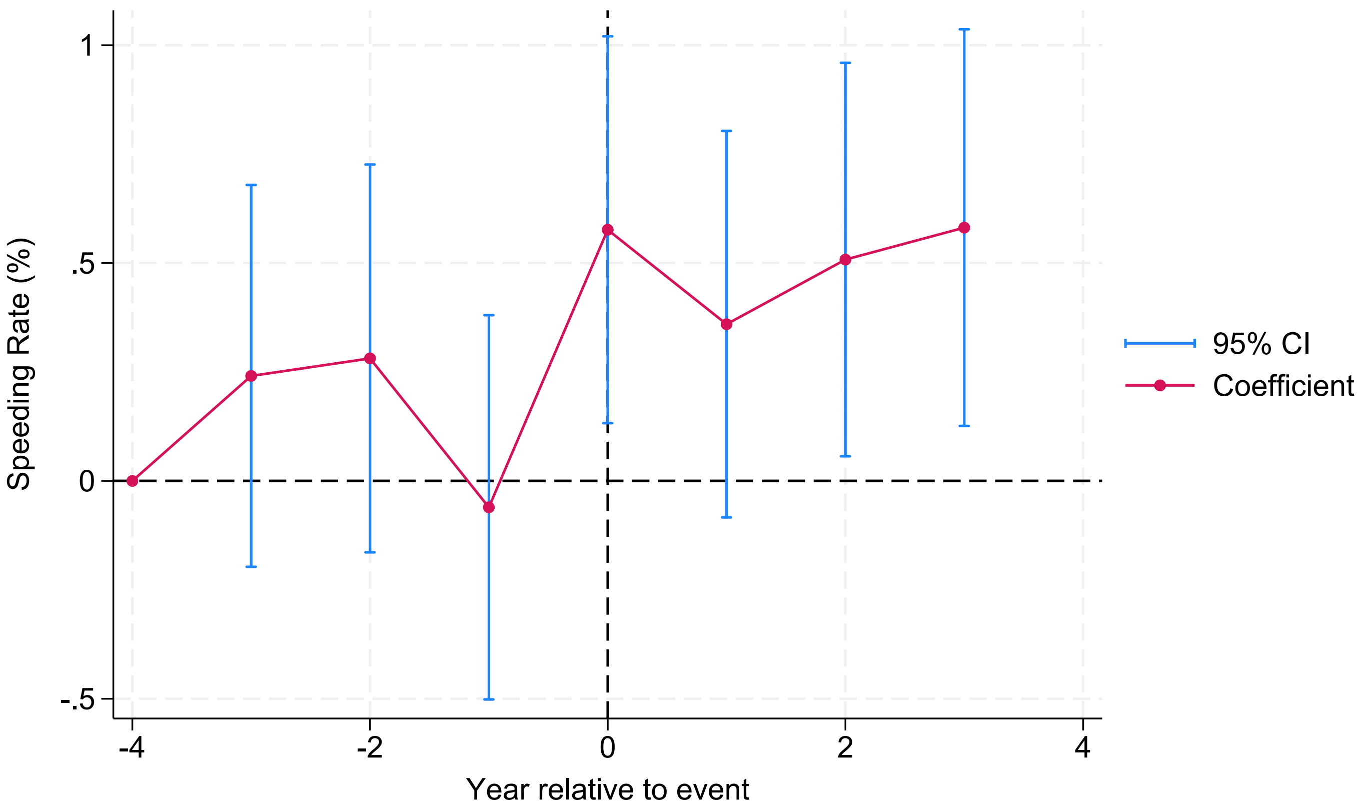

Differences in event-study coefficients, $\beta^t_{\text{top}} - \beta^t_{\text{bottom}}$, between top and bottom inheritance-size groups from equation 6.3. See figure 6.2 for separate estimates of β^t_{top} and β^t_{bottom}, as well as details on sample construction. Each death-cohort-by-inheritance-size group is reweighted to match the 2012 distribution of age, sex, and education in the population register. Standard errors are clustered at the individual level.

estimating income-effect heterogeneity, which is important for evaluating the redistributive potential of income-based fines; we therefore use a 4-year horizon in our main results.

To demonstrate our DDD-IV design, we estimate a single average income effect before estimating income-effect heterogeneity in section 6.4. In particular, we specify the following system:

$$\begin{aligned}\Delta W_{i,y,t} = \theta_0 + &\theta_1 * I_{t\geq 0} + \theta_2 * I_{YOD(i)=y} + \theta_3 * \text{Top}(i) + \theta_4 * I_{t\geq 0} * I_{YOD(i)=y} + \theta_5 * I_{t\geq 0} * \text{Top}(i) \\ &+ \theta_6 * I_{YOD(i)=y} * \text{Top}(i) + \theta * I_{t\geq 0} * I_{YOD(i)=y} * \text{Top}(i) + \epsilon_{i,y,t}\end{aligned} \tag{6.4}$$

$$\begin{aligned}s_{i,y,t} = \beta_0 + &\beta_1 * I_{t\geq 0} + \beta_2 * I_{YOD(i)=y} + \beta_3 * \text{Top}(i) + \beta_4 * I_{t\geq 0} * I_{YOD(i)=y} + \beta_5 * I_{t\geq 0} * \text{Top}(i) \\ &+ \beta_6 * I_{YOD(i)=y} * \text{Top}(i) + \beta * I_{t\geq 0} * I_{YOD(i)=y} * \text{Top}(i) + \nu_{i,y,t}\end{aligned} \tag{6.5}$$

where $\Delta W_{i,y,t}$ is the wealth shock for individual i in year $y+t$, $\text{Top}(i)$ is an indicator for whether i is in the top inheritance group (quintiles 3-5), and $I_{t\geq 0}$ is a post-event indicator. We estimate the income effect $\hat{\beta}/\hat{\theta}$ via two-stage least-squares. Table 6.2 shows the results, with the reduced-form (equation 6.5) in column 1, the first stage (equation 6.4) in column 2, and the DDD-IV estimate in column 3. Our estimate in column 3 implies that a €10,000 inheritance shock causes a 0.26 pp increase in the likelihood of receiving a speeding ticket in the post-period; this effect is similar in magnitude to the average labor supply effect we estimate in section 6.2. As a result, the average earnings effect (labor supply effect + income effect) explains nearly one-half of the cross-sectional relationship between income and speeding in column 1 of table 6.1.

Overall, we view the income effect we estimate in column 3 of table 6.2 to be somewhat conservative for several reasons: (1) we assume that heirs "spend" the full inheritance

Table 6.2: DDD-IV Estimates of the Income Effect

	(1) Reduced-Form	(2) First-Stage	(3) IV
DDD Coefficient	0.391***	1.511***	
	(0.115)	(0.0344)	
Inherit. Shock (10,000s €)			0.259***
			(0.0765)
Constant	7.316***	0.00701***	7.314***
	(0.0540)	(0.000927)	(0.0541)
N	4912424	4912424	4912424

Estimates of β from equation 6.5 (column 1), θ from equation 6.4 (column 2), and β/θ (column 3). The treatment group consists of heirs whose first parent dies in [2013,2016]; the control group consists of heirs whose first parent dies in [2017,2020]. The bottom inheritance group, for a given parental death-cohort d, consists of heirs in the bottom two quintiles of total inheritance flows received in $[d, d+2]$; the top inheritance group for a given death-cohort consists of heirs in the top three quintiles of total inheritance flows received in $[d, d+2]$. We restrict to heirs who are 18 years-of-age or older in the year of the parental death event. Each death-cohort-by-inheritance-size group is reweighted to match the 2012 distribution of age, sex, and education in the population register. Standard errors (clustered at the individual level) are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

on speeding tickets within our 4-year treatment horizon, (2) some inherited assets are less liquid than cash,[41] and (3) inheritances are, to some extent, anticipated. As with our attenuated labor supply effects (section 6.2), weaker income effects likely bias us in favor of finding that speeding behavior is a tag for earnings ability. We now discuss several robustness checks before turning to section 6.4, in which we estimate across-income heterogeneity in our earnings effects and calculate the optimal redistribution-motivated fine gradient.

Labor Supply and Earnings Responses Our inheritance-shock approach is meant to measure the effect of income on speeding, holding labor supply fixed. However, if labor supply responds to inheritance receipt, our income-effect estimates may be contaminated. Similarly, to the extent that earnings respond to inheritance receipt, we may be mis-measuring the denominator in our income effect. Appendix figures 10.14 and 10.15

[41] Recall that we use the assessed value of the estate, regardless of its asset composition, as the denominator in our income effect. Since it is costly to liquidate certain asset types, using the assessed value is likely to be conservative. Unfortunately, we have no information on asset composition.

show triple-difference event studies around the parental death event for months employed and earnings respectively. We see a decrease in earnings following inheritance receipt; for months employed, we observe an anticipation effect in the year before the parental death. However, the magnitudes of these two effects are quantitatively negligible: in any given period, earnings decrease by no more than €400 and employment months increase by no more than 0.07 months.[42] Given the size of the inheritance shocks we use to estimate the income effect, we view these potential sources of contamination as negligible.

Effect Horizon We elect to maintain a 4-year effect horizon in order to increase our power for estimating across-income heterogeneity in income effects (see section 6.4). However, we examine robustness to including an additional post-treatment year, the results for which can be found in Appendix table 10.10.[43] Our estimates imply that a €10,000 inheritance shock causes a 0.3 pp increase in the likelihood of speeding ticket receipt post-treatment; this effect is about 15% higher than our baseline estimate in column 3 of table 6.2. These results support the contention that our baseline estimates are somewhat conservative, though we also note that the triple-difference event-study with 5 post-periods (Appendix figure 10.17) shows a decline in speeding effects after four years.

Comparability of Death Effects A key assumption we use to identify income effects is that the death effect is common across inheritance-size groups; this assumption allows us to attribute differences in speeding behavior between the top and bottom inheritance groups to differences in the size of inheritance shocks.

One reason this assumption may fail is if the likelihood of a transport-related cause-of-death (e.g. a car crash) is systematically correlated with inheritance size: heirs whose parent dies in a transport-related accident may exhibit differential speeding responses

[42]If anything, ignoring the earnings effect causes us to overstate the total income shock from inheritance receipt, leading to *downward* bias in the income effect.

[43]For this analysis, we use individuals whose first parent dies in $[y+1, y+5]$ as a control group for individuals whose first parent dies in y in order to increase power.

to parental death, potentially leading to bias in our DDD-IV approach. Indeed, using the cause-of-death register, we find in unreported results that parents of heirs in the top inheritance-size group are more likely to die from a transport-related accident.[44] However, in our inheritance sample transport-related accidents are quite rare, constituting no more than 0.6% of any given death cohort's causes-of-death. Therefore, when we re-estimate the income effect after dropping all heirs whose parent dies in a transport-related accident (Appendix table 10.8), our estimate is unchanged relative to our baseline.

Another reason that death effects may not be comparable across inheritance size is because the amount of effort involved in disbursing the estate varies across inheritance-size groups in a way that systematically affects the likelihood of speeding. For example, larger inheritances may be more complex and therefore require more travel to-and-from the parent's residence. To evaluate the role of inheritance complexity, we examine two features of the estate: (1) the number of (potential) child heirs and (2) the number of tax revisions we observe for an heir's inheritance. Appendix figure 10.16 shows that inheritance size is negatively correlated with the number of potential child heirs (panel 10.16a) and positively correlated with the number of tax revisions (panel 10.16b). To examine the relevance of these differences, Appendix table 10.9 shows income-effect estimates from our DDD-IV model when we include "placebo" controls for the number of child heirs (column 1) or the number of tax revisions (column 2);[45] our income effect estimates are virtually unchanged, even increasing slightly after controlling for the number of child heirs. We therefore do not view estate complexity as a meaningful confound in our estimates.

[44] Transport-related accidents correspond to Chapter V deaths in the cause-of-death code.

[45] In particular, this set of controls consists of all possible interactions between the estate complexity measure, the treatment indicator $I_{YOD(i)=y}$, and the post-indicator $I_{t\geq 0}$. In other words, we add a set of controls that substitutes the top-inheritance group indicator for the estate complexity variable, allowing us to test whether estate complexity is the true underlying treatment dimension that explains our results.

6.4 Optimal Redistributive Fine Gradient

We now use our estimates of the income and labor-supply effects to identify the optimal redistribution-driven gradient of speeding fines. We assume a constant fine elasticity ζ_τ and use the generalized expression for the optimal fine gradient from Appendix equation 10.3:

$$b^*_{redist} := \frac{\mathrm{Cov}\left[\hat{g}(z), \int_0^z (y - \overline{z}^{rw}) s'_{het}(y) dy\right]}{\zeta_\tau \int_Z z * (z - \overline{z}^{rw}) \frac{s(z)}{\tau(z)} dH(z)} \tag{6.6}$$

where $\overline{z}^{rw} := \frac{\mathbb{E}[\frac{s(z)}{\tau(z)} * z]}{\overline{s}/\overline{\tau}}$ is behavior-weighted average earnings. The numerator captures the redistributive value of fines and depends on the covariance between welfare weights $\hat{g}(z)$ and the earnings ability tag $s'_{het}(z) = s'(z) - s'_{inc}(z)$. The denominator, which is proportional to the fine elasticity ζ_τ, captures the consumption distortion induced by increasing the fine gradient. We take our estimate of ζ_τ from Kaila (2024), which uses a regression-discontinuity-design around the income-based fine cutoff to measure the effect of larger speeding fines on future recidivism.[46] We also explore variation around this estimate, though we note that, in our setting, the sign of the denominator in equation 6.6 will always negative if the elasticity is constant.[47]

Therefore, the sign of b^*_{redist} hinges on the relationship between the ability tag $s'_{het}(z)$ and earnings z: under conventional welfare weights that decrease with z, the numerator in equation 6.6 will be negative if $s'_{het}(z)$ is increasing in z. Thus, if the earnings ability signal $s'_{het}(z)$ is *positively* correlated with earnings, it is optimal to have speeding fines scale with income.

To obtain measures of $s'_{het}(z) = s'(z) - s'_{inc}(z)$ across earnings levels, we estimate heterogeneity in our labor supply and income effects. To make these two sets of estimates

[46]Recall that Kaila (2024) shows that individuals do not bunch at fine discontinuities, a finding which motivates the RDD approach.

[47]In particular, it is sufficient that $s(z)/\tau(z)$ is increasing in income. For overall speeding crime (fixed and day-fine crime combined), this condition holds in our setting.

internally consistent, we estimate labor supply effects using the individuals from our inheritance-shock sample who belong to the top inheritance-size group, *i.e.* the group for which we identify income effects in our DDD-IV approach (section 6.3); to obtain representative estimates, we also apply the demographic weights from our inheritance-shock design.[48] In particular, we assign individuals in the top-inheritance sample into 8 equal-sized income groups based on 2012 age-adjusted earnings;[49] we then estimate the wage-scaled labor supply effect (equation 6.2) separately for each group.[50] For income-effect heterogeneity, we are more limited by power considerations and therefore estimate separate effects only for below- and above-median earnings groups;[51] however, despite a large earnings difference between these two groups, we find only a small difference (0.076 pp) for the income effect, suggesting that heterogeneity is minimal. Given these two sets of estimates, we measure the earnings effect $s'_{inc}(z_j)$ for each income group $j \in \{1, 2, ..., 8\}$ as the sum of group-specific labor-supply and income effects; for the full set of estimates, see Appendix table 10.12.

To obtain the earnings ability tags $s'_{het}(z)$, we use the 2012 cross-sectional profile of speeding rates $s(z)$ to calculate across-income variation $s'(z)$, from which we subtract our earnings effects $s'_{inc}(z)$. Figure 6.4 displays our results. Cross-sectional variation $s'(z)$ (blue dashed line) initially increases with earnings before declining, reflecting the s-shaped cross-sectional profile $s(z)$ in Appendix figure 10.18. By contrast, the earnings effects $s'_{inc}(z)$ (red-dashed line) are relatively homogeneous; as a result, the profile of ability tags $s'_{het}(z) = s'(z) - s'_{inc}(z)$ (solid purple line) inherits the shape of the cross-sectional profile.

[48]Following this approach, Appendix table 10.11 estimates the overall cross-sectional gradient and average labor-supply effect for the top-inheritance sample using our within-individual design from section 6.2; the results are extremely similar to those in table 6.1 for the general population, suggesting that selection into the inheritance-shock sample is largely driven by observable demographics.

[49]We use 2012 as our benchmark year since it is one year prior to the first wave of parental deaths in our sample.

[50]For these estimates, we use the group-specific average wage to scale labor supply.

[51]To increase power, we use the set of parental death cohorts $[y + 1, y + 4]$ as a control group for cohort y, in contrast to our baseline approach which uses only the $y + 4$ cohort as a control. In unreported results, we find that expanding the control group and re-estimating the average income effect yields a result nearly identical to that of table 6.2.

Figure 6.4: Relationship between Ability Tags and Earnings

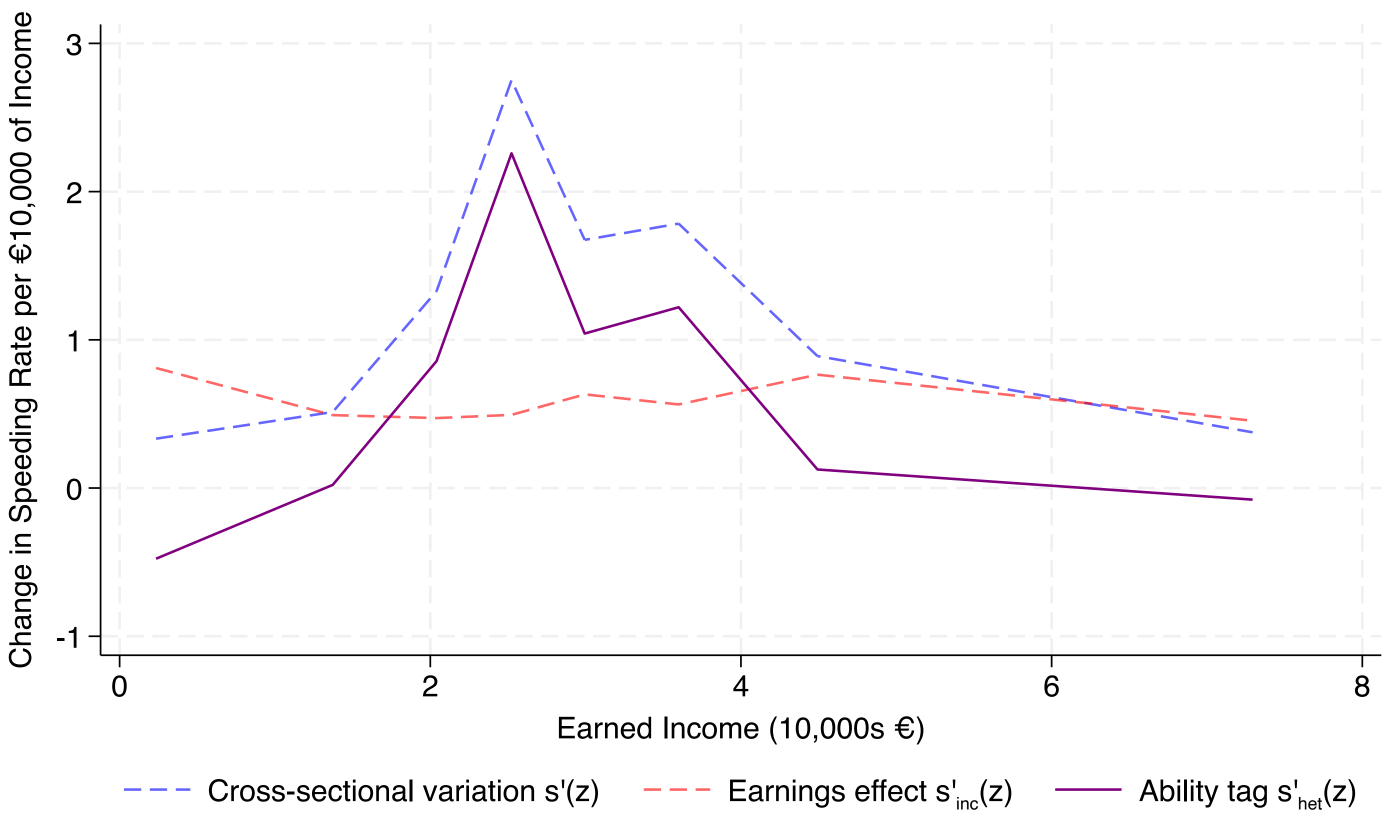

Estimates of $s'(z_j)$, $s'_{\text{inc}}(z_j)$, and $s'_{\text{het}}(z_j)$ for each income group $j \in \{1, 2, ..., 8\}$ among the sample of heirs in the top inheritance-size group. Individuals are reweighted to match the 2012 distribution of age, sex, and education in the population register. Income groups are assigned based on 2012 age-adjusted earnings. To calculate $s'(z_j)$, we estimate age-adjusted speeding rates in 2012 for each income group $s(z_j)$ (Appendix figure 10.18); then, for each income group $j < 8$, $s'(z_j) = \frac{s(z_{j+1}) - s(z_j)}{z_{j+1} - z_j}$. For the top income group $j = 8$, we estimate $s'(z_8)$ using a regression of speeding ticket receipt on earnings, as in column 1 of table 6.1. To estimate $s'_{\text{inc}}(z_j)$, we combine estimates of β_2 from equation 6.2 (the labor-supply effect) and β/θ from equations 6.4 and 6.5 (the income effect); for details on sample construction and the implementation of these equations, see section 6.4. Appendix table 10.12 shows the estimates that underlie this graph.

Recall that an increasing ability tag profile $s'_{het}(z)$ is sufficient for a positive fine gradient b^*_{redist}; however, the profile in figure 6.4 exhibits an inverted u-shape. With a more general non-linear speeding-fine system, as in Ferey et al. (2024), this result implies that the optimal fine schedule should *also* follow an inverted u-shape. Under the affine fine schedule we consider, the covariance between welfare weights $\hat{g}(z)$ and ability tags $s'_{het}(z)$ is determinative: the redistributive gradient b^*_{redist} will be positive if this covariance is negative. For our construction of welfare weights, we take the same approach detailed in section 5.3, using welfare weights $g(z) \propto y^{-\nu}$, where y is net income (Saez, 2002b).

Figure 6.5 plots the earnings ability tags alongside a profile of welfare weights derived using $\nu = 1$; the two series move in opposite directions at low incomes but co-move for a large segment of the earnings distribution. The resulting covariance is *positive*, implying that the optimal fine gradient b^*_{redist} is *negative*; using equation 6.6 and the fine elasticity estimate $|\zeta_\tau| = 0.157$ from Kaila (2024), we estimate an optimal redistributive gradient of -12€ per €10,000 of earnings, implying that fines should *decrease* by €12 for every €10,000 of income at the optimum.[52] Using this estimate, we cannot rationalize the steepness of either the day-fine gradient (+88 €) or the combined-fine gradient (+12.4 €) that we estimated in section 2.

Figure 6.5: Covariance between Welfare Weights and Ability Tags

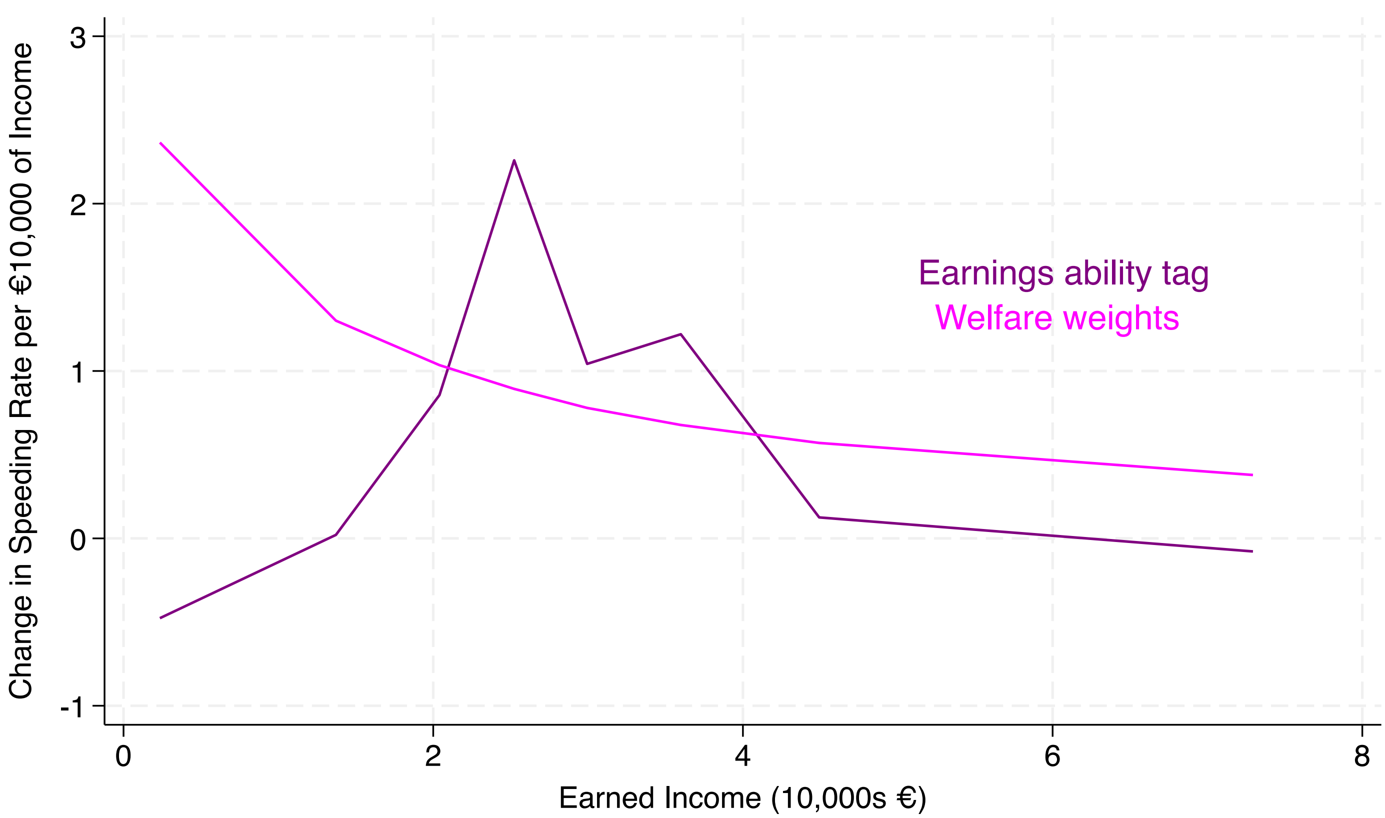

Estimates of welfare weights $g(z_j)$ and earnings ability tags $s'_{\text{het}}(z_j)$ for each income group $j \in \{1, 2, ..., 8\}$ among the sample of heirs in the top inheritance-size group. Individuals are reweighted to match the 2012 distribution of age, sex, and education in the population register. For details on the estimation of $s'_{\text{het}}(z_j)$, see section 6.4. To construct welfare weights, we use $g(z_j) \propto y(z_j)^{-1}$, where $y(z_j)$ is the average net income of the jth income group; see Appendix table 10.12 for the estimates of $y(z_j)$ and $s'_{\text{het}}(z_j)$ underlying this graph.

[52]This estimate also depends on the observed distribution of income, speeding rates, and fine levels; see Appendix table 10.12 for details.

To assess the robustness of our results, Appendix figure 10.19 shows a heatmap of implied optimal fine gradients b^*_{redist} as a function of the fine elasticity $|\zeta| \in \{0.1, 0.157, 0.2\}$ and the degree of inequality aversion $\nu \in \{0.25, 1, 4\}$. Across all considered values, we find a negative gradient, with the strongest (*i.e.* most negative) gradients arising from weaker fine elasticities and stronger redistributive preferences (higher ν). While the former is a mechanical result of equation 6.6, the relationship between the optimal gradient and the degree of inequality aversion ν reflects the logic of tagging. In particular, when ν is large, the planner only places weight on the welfare of the poorest individuals and therefore seeks to maximize revenue collection for the purpose of redistribution; the (overall) negative relationship between ability tags and earnings in figure 6.4 implies that revenue is maximized using (1) lower fines and higher income tax rates for the rich, together with (2) higher fines and lower income tax rates for low-to-middle incomes.

In summary, we find that the latent earnings-ability tag in speeding behavior is *decreasing* with earnings overall, implying that speeding fines are less efficient as a means of raising revenue for redistribution at high incomes. As a result, given our estimates, optimally redistributive speeding fines would *decrease* with income under any conventional parameterization of welfare weights. We therefore reject the redistributive rationale for income-based speeding fines in Finland.

7 Fairness Preferences

We now discuss our survey, which we use to evaluate the two fairness-driven rationales for income-based speeding fines outlined in our theory (section 3). In particular, we posit that income-based fines may be valuable as instruments for achieving "equal compliance" (*i.e.* equal speeding rates across incomes) or proportional punishment, where the latter may reflect e.g. a desire to equalize marginal deterrence (see section 3). We build our survey to quantify both these forces and any residual value that respondents may place on income-based fine policies. The key question in our survey asks respondents to hypothetically compare a given fixed-fine policy to a randomly-varied income-dependent policy; we begin by showing how such comparisons can be used to identify the fairness parameters in equation 3.2.

7.1 Conceptual Approach

Suppose there is a prevailing fine schedule $\tau(z) = a^*(b) + b * z$, where $a^*(b)$ is the optimal fine intercept given a slope parameter b. Suppose also that the behavioral- and fine-fairness terms $\overline{e}_s(z)$ and $\overline{e}_\tau(z)$ are affine (as in section 3). Then, the parameters that determine how fairness rationales enter the optimal fine gradient b^* are β_s and β_τ, the slope terms for behavioral- and fine-fairness respectively (see equation 3.2). The key result informing our survey design is that these slope parameters can be identified from comparisons between simple two-tiered fine policies $\{\tau_{\text{below}}(z), \tau_{\text{above}}(z)\}$ that assign a fine $\tau_{\text{above}}(z)$ for above-median earners and a fine $\tau_{\text{below}}(z)$ for below-median earners. To see this, consider a given two-tiered speeding fine policy with associated speeding rates $\{s_{\text{below}}(z), s_{\text{above}}(z)\}$. Then, randomly varied $\big\{\{-d\tau, d\tau\}, \{ds, -ds\}\big\}$ reforms, *i.e.* reforms with equal-and-opposite fine/behavior changes, are sufficient to identify β_s and β_τ. Intuitively, because the two income groups are equal-sized, the equal-and-opposite nature of the reform eliminates each of the α intercept terms, implying that the total fairness effects depend only on the β slope terms.[53] Therefore, comparisons between simple two-tiered fine policies are adequate to test whether our hypothesized fairness motives can rationalize a day-fine system.

We also introduce a residual preference β_0 that "turns on" when the fine schedule exhibits *some* income-dependence ($b > 0$);[54] this term captures preferences for income-dependence that are unrelated to either induced behavior or the shape of the fine schedule, such as a

[53]For expressions of these effects, see Appendix table 10.13.

[54]Note that, conditional on having a positive fine gradient $b > 0$, the fine- and behavioral-fairness terms inform the optimal "steepness" of the fine schedule, while the residual term only contributes to the extensive-margin choice between $b = 0$ and $b > 0$.

preference for procedural fairness. The residual can be separately identified using comparisons between income-based and fixed-fine policies; we therefore ask each respondent to compare hypothetical fixed and income-based policies, where the latter are randomly varied across respondents with respect to both fine-schedule "steepness" and the induced distribution of behavior. Figure 7.1 demonstrates our identification approach. The policy in the left column levies fixed fines and exhibits across-income inequality in speeding rates; the policies in the right column are all income-dependent, but vary with respect to fine-schedule steepness and the distribution of speeding rates. If equal compliance (*i.e.* equal speeding rates) is the instrumental objective of income-based fines, then respondents should value the income-based policy in the first row ("Proportional Fines") less than the income-based policy in the second row ("Proportional Fines and Equal Speeding"), implying $\beta_s > 0$. Similarly, if proportional punishment is the instrumental objective, then respondents should value the income-based policy in the second row more than the income-based policy in the third row, implying $\beta_\tau > 0$. Any residual value not attributable to these effects is ascribed to β_0.

Note that, since we use equal-and-opposite fine reforms, the overall speeding rate remains constant across the policies in figure 7.1 when the income groups are equal-sized; therefore, unless respondents believe that the cost of speeding depends on the income of the speeder, safety-related costs should not be a confound between policies.[55] However, three other confounds might arise. First, the comparison across income-dependent policies, *i.e.* within the right column of figure 7.1, is confounded by the behavior *change* induced by the income-based policy relative to the fixed-fine policy. We therefore also randomize the size of the equal-and-opposite behavior change induced by a given income-based policy; implicitly, this provides another way to test the behavioral fairness rationale since, for a given income-dependent policy, increasing the behavior change implies more inequality in the fixed fine policy. If respondents value increases in behavioral equality, they should

[55]See section 7.2 for how we make this aspect salient to respondents. We also conduct robustness checks based on free-response data (see section 7.6).

Figure 7.1: Survey Identification

Fixed-Fine Policy

Inc. Group	Avg. Inc.	Fine	Crime Rate
Low	€15,000	€330	50
High	€51,000	€330	90

Proportional Fines

Inc. Group	Avg. Inc.	Fine	Crime Rate
Low	€15,000	€150	60
High	€51,000	€510	80

Prop. Fines & Equal Speeding

Inc. Group	Avg. Inc.	Fine	Crime Rate
Low	€15,000	€150	70
High	€51,000	€510	70

Equal Speeding

Inc. Group	Avg. Inc.	Fine	Crime Rate
Low	€15,000	€240	70
High	€51,000	€420	70

Fixed-Fine Policy $\xrightarrow{\beta_0 - \beta_s}$ Proportional Fines

Fixed-Fine Policy $\xrightarrow{\beta_0}$ Prop. Fines & Equal Speeding

Fixed-Fine Policy $\xrightarrow{\beta_0 - \beta_\tau}$ Equal Speeding

Prop. Fines & Equal Speeding $\xrightarrow{-\beta_s}$ Proportional Fines

Prop. Fines & Equal Speeding $\xrightarrow{-\beta_\tau}$ Equal Speeding

therefore place a higher value on income-dependent policies that induces larger behavior changes.

Second, in comparison to a fixed-fine policy, fine policies that decrease fines for low-income individuals may be desirable due to liquidity concerns, a preference that is arguably separate from the fairness concerns we're interested in.[56] To deal with this concern, we vary the fixed-fine level, which, by construction, also equals the average fine under the income-dependent policy:[57] if liquidity drives demand for income-dependent fines, then demand should decrease as we lower the fixed-fine level.

Finally, the comparison between fixed and income-dependent policies may be confounded by redistributive preferences: respondents may favor income-dependent fines simply because they implicitly transfer money to infra-marginal low-income individuals (at the expense of inframarginal high-income individuals). To first-order, increasing the speed-

[56]E.g., the formalization of the marginal deterrence rationale outlined in section 3 uses a static model with no liquidity concerns.

[57]By average, we mean an unweighted average between the statutory fine levels, *i.e.* not a behavior-weighted average.

ing fine by $d\tau$ for an individual i with a speeding rate s_i is equivalent to taking $d\tau * s_i$ € away from i; our approach to addressing redistributive preferences builds on this first-order approach. In particular, in an auxiliary question, we ask respondents to compare a baseline fixed-transfer policy to an income-dependent transfer policy, where the difference between the two transfer schedules captures the implicit (first-order) redistribution implied by the income-based fine policy.[58] We then difference out respondents' valuations for the income-dependent transfer policy from their valuations of the income-dependent fine policy, yielding a net-of-redistribution value for income-dependent fines.

7.2 Survey Design

We now discuss the key elements of our survey design that facilitate the hypothetical policy comparisons discussed in section 7.1. The fine policy comparison consists of four parts: (1) an explanation of speeding-crime rates + a comprehension question to ensure that respondents understand this concept (see Appendix figure 10.20);[59] (2) a qualitative choice between fixed and income-dependent fine policies; (3) a slider that respondents use to indicate how much government revenue they would be willing to forgo in order to implement their preferred fine policy; and (4) an open-text question that asks respondents to describe the factors that influenced their choice in (2) (see Appendix figure 10.21). Figure 7.2 shows an example of a qualitative fine-policy comparison (part 2); we note several features. First, respondents are told that their own hypothetical fine is unaffected and to make a choice based on what would be better for Finland overall; these instructions are intended to remove own-fine considerations and put respondents in the frame of a planner making trade-offs between social objectives. Second, in addition to the fine schedules and crime rates, we also indicate the total number of speeding tickets that arise under each policy; to control for safety-related concerns, this number is held fixed, a fact that we also emphasize to respondents (see the bottom of Figure 7.2).

[58] For details on how we construct these transfers, see Appendix section 10.2.4.

[59] We force respondents to correctly answer this question before proceeding.

After making a choice between the two fine policies, we tell respondents that the government could hypothetically save on administrative costs by implementing their dis-preferred system (see Appendix figure 10.22 for exact wording). Respondents then use a slider (Figure 7.3) to indicate the minimum annual savings that would compel them to switch to their dis-preferred policy. We display the savings in both total and per-capita terms; the slider is capped at an annual savings of €275 per-capita,[60] which corresponds to a total annual savings of €1.238 billion. Finally, respondents use a text box to describe the factors that influenced their fine-policy choice; we use these responses to evaluate the persistence of confounds in our elicitation (section 7.6).

In an auxiliary set of questions, we also ask respondents to compare fixed and income-dependent transfer policies for the purpose of differencing out redistributive motives from the fine-policy comparison, as described in section 7.1. Analogous to our approach for fine policies, respondents make a qualitative comparison between fixed and income-dependent transfers (Appendix figure 10.23), use a slider to indicate willingness-to-pay for their preferred transfer policy, and explain their reasoning using an open-text box. As in the fine policy comparison, we tell respondents that they would not personally receive a transfer and ask them to choose the better policy for Finland overall.[61] We randomize whether or not the transfer questions appear before the fine-policy questions and present an attention check (described in section 7.3) between the two sets of questions. Each respondent completes only one set of responses (*i.e.* one fine-policy comparison and one transfer comparison); this mitigates survey fatigue and ensures that respondents do not see mutually inconsistent policies.[62] Our identification therefore relies on across-

[60] In our survey, "per-capita" means per Finnish adult, of which there are around 4.5 million.

[61] One difference between the two policy comparisons is that the transfer policy is motivated as a one-time occurrence, which is a more natural story for transfers. Because the government savings being considered in the transfer scenario are also one-time savings—as opposed to annual savings in the fine-policy comparison—the distinction does not matter for comparing valuations across the two policy types.

[62] The survey as-is takes the median respondent around 13 minutes to complete and the fine policy questions in particular are highly involved. Furthermore, additional fine comparisons would also require additional transfer comparisons. On the second point, given that our survey is hypothetical, showing respondents mutually inconsistent policies (e.g. income-dependent

Figure 7.2: Example Fine-Policy Comparison

This question compares two hypothetical speeding ticket systems that could be introduced in Finland:

1. **Fixed fines -** the fines are the same for everyone.
2. **Income-based fines -** the size of the fine depends on whether a person belongs to a low or high income group.

We ask you to choose **which system would be better for Finland overall**; suppose that **the fine YOU would have to pay would not be determined by the systems described above.**

The tables below present these two systems for two different income groups with different income levels. **Each income group covers half of the adult population in Finland. The low-income group includes the half of the population with the lowest income. The high-income group includes the half with the highest income.**

Fixed fines		
Income Group	Speeding Fine	Fines per 1000 people
Low (*Average Income*: 15 000 €)	**330 €**	**50**
High (*Average Income*: 51 000 €)	**330 €**	**90**
Total speeding tickets per year: 315 000		

Income-based fines		
Income Group	Speeding Fine	Fines per 1000 people
Low (*Average Income*: 15 000 €)	**150 €**	**70**
High (*Average Income*: 51 000 €)	**510 €**	**70**
Total speeding tickets per year: 315 000		

Which system do you prefer?

Please note that the systems differ in how much speeding is done in different income groups and how large the fines are. **However, the overall speeding rate is the same in both systems.** The systems are otherwise completely similar.

Figure 7.3: Slider for Fine Policy Comparison

Consider options A and B below.

Option A	Option B
Income-based fines are implemented	**Fixed fines** are implemented and **Savings of Y € per adult (X million € in total) annually**

The tables below present the two fine policies again. What would be the minimum annual savings required for you to choose option B?

I would choose option B if the savings were
at least Y € per adult (X million € in total) annually

(Remember, the ***fine YOU would have to pay is not determined by the policies described below.*** *If the smallest total savings that would make you switch is greater than the values shown on the slider, choose the highest possible value.)*

respondent randomization.

7.3 Survey Recruitment

We collected 3,038 survey responses from a survey participant pool of Finnish residents maintained by Norstat. To create representative survey samples, Norstat monitors the demographic distribution of respondents as surveys are completed and dynamically adjusts its contact rate for different groups. As a result, our survey is approximately representative on income, age, and gender, based on demographic information we collect from respondents at the end of the survey (see Appendix table 10.14). With respect to rates of speeding-ticket receipt and car-ownership, our sample is also consistent with statistics from administrative data.[63]

Consistent with our pre-registration, we drop respondents that fail an attention check[64] (104 respondents) or who choose the maximum slider value in *both* the fine- and transfer-policy questions (210 respondents) in order to construct our main analysis sample.[65] The summary statistics of our main estimation sample are shown Appendix table 10.14; differences between the analysis sample and the full set of respondents are negligible.

policies with the same schedules but different speeding rates) could undermine the perceived legitimacy of a given policy.

[63]In our administrative data on speeding ticket receipt, available through 2019, around 7% of individuals receive a ticket in a given year. In our survey, we find that the share of respondents who report receiving a ticket in the last *2* years is slightly above 7%, though it is well known that enforcement rates have declined following the implementation of a new camera system in 2020. For car ownership rates, a 2016 household survey (see here) found that around 74% of households have at least one car; for our survey respondents, this rate is around 80%.

[64]We use the following attention check, taken from Rees-Jones and Taubinsky (2019):

> *Sometimes, respondents to online surveys do not read all the instructions carefully and proceed too quickly. With this question, we want to make sure that respondents to this survey pay attention to the instructions provided. We now ask you to simply press the "Next" button at the bottom of the page. Do not select any of the options listed below. You can now move forward by pressing "Next".*

Respondents are provided three answer choices, "Never," "Sometimes," and "Always;" any respondent who fills in one of these answers fails the attention check.

[65]For respondents that choose the maximum slider value in both policy questions, we cannot infer the sign of WTP for income-dependent fines after differencing out redistributive motives.

7.4 Estimation

Our main dependent variable of interest is net-of-redistribution WTP for income-dependent fines $\text{WTP}_i^{\text{net}}$ for each respondent i. To calculate this variable, we take $\text{WTP}_i^{\text{inc-fine}}$, respondent i's WTP for the income-based fine policy, and subtract $\text{WTP}_i^{\text{inc-transer}}$, respondent i's WTP for the income-based transfer.[66] Our main regression assesses how net WTP responds to the randomized features of the fine-policy comparison, which are displayed in table 7.1; details on how we choose the average speeding rate, fine levels, and below-/above-median incomes are discussed in Appendix section 10.2.5. Two points are worth emphasizing about table 7.1. First, starting from an induced speeding behavior distribution with $s_{\text{below}}^{\text{inc}} = s_{\text{above}}^{\text{inc}}$ (equal speeding rates), we consider induced distributions with either $s_{\text{above}}^{\text{inc}} < s_{\text{below}}^{\text{inc}}$ (low-income individuals speed *less* often) or $s_{\text{above}}^{\text{inc}} < s_{\text{below}}^{\text{inc}}$ (low-income individuals speed *more* often) in order to test the equal-compliance rationale, which predicts that deviations from equality are undesirable. Second, given the average earnings levels of the below- and above-median income groups, the income-based fine schedules we consider can be grouped into three discrete categories: income-proportional, below-proportional, and above-proportional.[67] In our main regression, we collapse fine-schedule heterogeneity into these categories, a choice that allows us to explicitly test the proportional-punishment rationale.

Echoing our notation elsewhere in the paper, we use b to refer to the fine gradient of the income-based policy and designate b_{prop} as the proportional fine gradient. Letting $\mathrm{I}(\cdot)$ be an indicator function, our main regression is given by:

[66] For individuals that prefer fixed fines or fixed transfers, these WTP measures are negative.

[67] Formally, income-proportional fines satisfy $\tau_{\text{above}}^{\text{inc}}/\tau_{\text{below}}^{\text{inc}} = \overline{z}_{\text{above}}/\overline{z}_{\text{below}}$. Below-proportional fines are more equal than the income ratio ($\tau_{\text{above}}^{\text{inc}}/\tau_{\text{below}}^{\text{inc}} < \overline{z}_{\text{above}}/\overline{z}_{\text{below}}$) and above-proportional fines are more unequal than the income ratio ($\tau_{\text{above}}^{\text{inc}}/\tau_{\text{below}}^{\text{inc}} > \overline{z}_{\text{above}}/\overline{z}_{\text{below}}$).

Table 7.1: Survey Randomization

Characteristic	Values
Behavior (crimes per 1,000)	
Behavior Change $\Delta s := s^{\text{inc}}_{\text{below}} - s^{\text{fixed}}_{\text{below}}$	$\{20, 30\}$
Induced Speeding Distribution $(s^{\text{inc}}_{\text{below}}, s^{\text{inc}}_{\text{above}})$	$\{(60, 80)\ ;\ (70, 70)\ ;\ (80, 60)\}$
Fines (€)	
Fixed-fine $\tau^{\text{fixed}} = 200$	
Income-Based Fines $(\tau^{\text{inc}}_{\text{below}}, \tau^{\text{inc}}_{\text{above}})$	$\{\ \underbrace{(145, 255)}_{\text{Below-Proportional}}\ ;\ \underbrace{(90, 310)}_{\text{Proportional}}\ ;\ \underbrace{(35, 365)}_{\text{Above-Proportional}}\ \}$
Fixed-fine $\tau^{\text{fixed}} = 330$	
Income-Based Fines $(\tau^{\text{inc}}_{\text{below}}, \tau^{\text{inc}}_{\text{above}})$	$\{\ \underbrace{(240, 420)}_{\text{Below-Proportional}}\ ;\ \underbrace{(150, 510)}_{\text{Proportional}}\ ;\ \underbrace{(60, 600)}_{\text{Above-Proportional}}\ \}$

$$\begin{aligned}\text{WTP}^{\text{net}}_i =& \beta_0 - \beta^{\text{equal}}_s * \text{I}(s^{\text{inc}}_{\text{below}} < s^{\text{inc}}_{\text{above}}) + \beta^{\text{unequal}}_s * \text{I}(s^{\text{inc}}_{\text{below}} > s^{\text{inc}}_{\text{above}}) + \beta^{\text{behavior}}_s * \Delta s_i \\ & - \beta^{\text{prop}}_\tau * \text{I}(b < b_{\text{prop}}) + \beta^{\text{above-prop}}_\tau * \text{I}(b > b_{\text{prop}}) + \epsilon_i\end{aligned} \tag{7.1}$$

The constant term β_0 measures net WTP for the omitted policy with proportional fines, equal speeding rates, and a baseline behavior change $\Delta s = 20$ (see table 7.1); the other β parameters therefore measure how net WTP responds to deviations from this baseline policy. In particular, the parameters β^{equal}_s and β^{unequal}_s measure the effects of deviating from an equal distribution of speeding rates; we normalize the sign of these parameters so that β^{equal}_s measures the effect of moving from $s^{\text{inc}}_{\text{below}} < s^{\text{inc}}_{\text{above}} \rightarrow s^{\text{inc}}_{\text{below}} = s^{\text{inc}}_{\text{above}}$, while β^{unequal}_s measures the effect of moving from $s^{\text{inc}}_{\text{below}} = s^{\text{inc}}_{\text{above}} \rightarrow s^{\text{inc}}_{\text{below}} > s^{\text{inc}}_{\text{above}}$. Therefore, both parameters measure the effect of increasing the low-income speeding rate while simultaneously decreasing the high-income speeding rate; the equal-compliance rationale predicts $\beta^{\text{equal}}_s > 0$ and $\beta^{\text{unequal}}_s < 0$. Furthermore, if respondents generally value decreases in inequality, then we would also expect $\beta^{\text{behavior}}_s > 0$: since larger behavior changes Δs imply more inequality under the fixed-fine policy (as discussed in section

7.1), inequality-averse respondents should place a higher relative value on income-based polices that induce larger behavior changes.

Turning to fine-schedule effects, β_τ^{prop} measures the effect of moving from $b < b_{\text{prop}} \rightarrow b = b_{\text{prop}}$, while $\beta_\tau^{\text{above-prop}}$ measures the effect of moving from $b = b_{\text{prop}} \rightarrow b > b_{\text{prop}}$. Therefore, both parameters measure the effect of decreasing the low-income fine while simultaneously increasing the high-income fine; the proportional-punishment rationale predicts $\beta_\tau^{\text{prop}} > 0$ and $\beta_\tau^{\text{above-prop}} < 0$.

We note that, in a slight abuse of notation, the β_s and β_τ parameters in regression equation 7.1 are re-scaled versions of the parameters that appear in our expression for the optimal fine-gradient (equation 3.2); in section 7.5.3, we convert our baseline estimates into implied optimal fine-gradients.

7.5 Results

7.5.1 Baseline Estimates

We begin by discussing respondents' average extensive-margin preferences for income-based fine policies. Table 7.2 shows the share of our main sample that strictly prefers the income-dependent fine policy (row 1), strictly prefers the income-dependent transfer (row 2), and strictly prefers income-based fines after accounting for redistributive preferences (row 3); we say a respondent has a strict preference for a policy when they indicate a positive WTP for it. More than 3/4 of our sample strictly prefers the income-dependent fine policy; for the income-dependent transfer, that share is slightly lower, but still more than 2/3. On net, 56% of our sample strictly prefers the income-dependent fine policy after accounting for redistributive preferences expressed in the transfer comparison.

Turning to our WTP measure, table 7.3 shows that respondents are on-average willing to forgo €75 per-capita of government revenue in order to implement an income-based fine

Table 7.2: Extensive-Margin Preferences for Hypothetical Survey Policies

	Share WTP > 0
$\text{WTP}^{\text{inc-fine}}$	0.768
$\text{WTP}^{\text{inc-transfer}}$	0.678
WTP^{net}	0.558
$\Delta\left(\text{WTP}^{\text{inc-fine}}, \text{WTP}^{\text{net}}\right)$	0.209^{***} (0.01)
N	2727

Share of the survey sample (see section 7.3) with positive willingness-to-pay for income-based fines (row 1) and income-based transfers (row 2). Row 3 displays the share of respondents with positive net-of-redistribution willingness-to-pay for income-based fines; see section 7.4 for details on this willingness-to-pay measure. Standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

policy. Given an average WTP for income-dependent transfers of €43 per-capita, average net-of-redistribution WTP for income-based fines is €32 per-capita, which corresponds to a total WTP of €144 million. We can also see from table 7.3 that median net WTP is about 25% as large as the average, though we can reject that the median is equal to 0 and the implied total WTP (€36 million) is still substantial. Together with our extensive-margin findings in table 7.2, these results indicate substantial fairness-driven support for income-based fines on-average; furthermore, they provide reassurance that our elicitation is not purely noise.[68]

Table 7.4 explores how WTP responds to randomized fine-policy characteristics using equation 7.1. The first row shows per-capita WTP for the omitted proportional-and-equal fine policy; rows 2-4 show the β_s coefficients, *i.e.* the effects of randomized behavior characteristics on WTP; rows 5-6 show the β_τ coefficients, *i.e* the effects of randomized fine-schedule steepness on WTP. To aid interpretation of our results, we include the size

[68] In particular, respondents can rank policies using our WTP measure.

Table 7.3: Mean & Median WTP for Hypothetical Survey Policies

	Per-Capita	Aggregate
Mean		
Income-Based Fines	€75.25	€339 million
Income-Based Transfers	€43.27	€195 million
Net WTP	€31.98	€144 million
Median		
Net WTP	€8	€36 million
N	2727	

Mean willingness-to-pay for income-based fines (row 1) and income-based transfers (row 2) using the survey sample described in section 7.3; row 3 displays the average difference between these measures and row 4 displays the median difference.

of the equal-and-opposite changes (relative to the omitted policy) that identify each coefficient.[69] We begin with column 1, which presents OLS estimates using $\text{WTP}^{\text{inc-fine}}$—*i.e.* WTP for income-based fines prior to accounting for redistribution—as the outcome. We see a very high per-capita WTP (€88) for the proportional-and-equal omitted policy in row 1; however, responses to randomized behavior and fine-schedule characteristics are insignificant and quantitatively small relative to WTP for the omitted policy. Column 2 presents OLS estimates using our primary outcome measure, net-of-redistribution WTP; while net WTP for the omitted policy is lower than gross WTP in column 1, the responses to randomized fine-policy characteristics are largely unchanged. In order to examine whether the top-coding of slider values drives our results, column 3 presents OLS estimates using an alternative net WTP measure that assigns larger values to respondents who max out the slider.[70] However, the results using this alternative measure

[69]Given the randomized values in table 7.1 and our choice of omitted policy, all β_s coefficients in equation 7.1 measure the effect of an equal-and-opposite speeding-rate change $ds = 0.01$ (10 crimes per 1,000). Furthermore, the implied equal-and-opposite fine change $d\tau$ between adjacent fine-schedule categories is €72.5, *i.e.* the average of 90 and 55. However, when we condition on a given fixed-fine level, this implies a different scaling for the fine changes.

[70]In our baseline net WTP measure, we assign the maximum value, €275, as the WTP for those who max out the slider. In our alternative measure, we assign a value of €325.

are nearly unchanged relative to column 2. Column 4 presents estimates from a quantile (median) regression with net WTP as the outcome and finds qualitatively similar results. Furthermore, in Appendix table 10.15, we replicate the same pattern of findings using extensive-margin outcomes.

Table 7.4: Effects of Fine-Policy Characteristics on Hypothetical WTP

	Per-Capita WTP for Income-Based Fines			
	(1) Gross	(2) Net	(3) Net (alt.)	(4) Net
Proportional Fines, Equal Speeding, $\Delta s = 0.02$	87.56***	45.74***	47.09***	16*
	(9.557)	(11.36)	(11.99)	(6.337)
Behavior (ds = 0.01)				
Equality ($s^{\text{inc}}_{\text{below}} = s^{\text{inc}}_{\text{above}}$)	-7.483	-3.473	-2.906	-2
	(6.104)	(7.256)	(7.659)	(4.048)
Inequality ($s^{\text{inc}}_{\text{below}} > s^{\text{inc}}_{\text{above}}$)	-7.899	-7.090	-6.567	-4
	(6.109)	(7.262)	(7.666)	(4.051)
Behavior Change Δs	-4.823	-8.496	-8.618	-6
	(4.951)	(5.886)	(6.213)	(3.283)
Fine Schedule (dτ =€72.5)				
Proportional ($b = b_{\text{prop}}$)	6.191	-5.022	-5.961	-4
	(6.044)	(7.185)	(7.584)	(4.008)
Above-Proportional ($b > b_{\text{prop}}$)	-8.429	-4.641	-5.081	2
	(6.137)	(7.295)	(7.700)	(4.069)
N	2727	2727	2727	2727
Model	OLS	OLS	OLS	Quantile

Estimates of equation 7.1 using the survey sample described in section 7.3. Standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In summary, respondents value income-based fines but are relatively insensitive to three randomized features of such policies: (1) the induced *distribution* of behavior, (2) the induced level of behavior *change*, and (3) the steepness of the fine schedule. We now probe the heterogeneity of these baseline results.

7.5.2 Heterogeneity

Table 7.5 presents OLS estimates of equation 7.1 separately for different randomized characteristics: the fixed-fine level, the level of induced behavior change, and the question order. Throughout, we focus our attention on the effects of randomized behavior and fine-schedule characteristics on net WTP, *i.e.* the β_s and β_τ coefficients from equation 7.1. Beginning with the fixed-fine level, columns 1 and 2 show results conditional on $\tau^{\text{fixed}} =$€200 and $\tau^{\text{fixed}} =$€330 respectively. While the results are quite noisy when $\tau^{\text{fixed}} =$€330, we see relatively precise responses to randomized behavior characteristics when $\tau^{\text{fixed}} =$€200. However, the signs for these effects are generally the *opposite* of those predicted by the equal-compliance rationale: respondents prefer income-based policies with $s^{\text{inc}}_{\text{below}} < s^{\text{inc}}_{\text{above}}$ to policies with equal speeding rates (row 2, column 1) and are willing-to-pay to *increase* behavioral inequality under the fixed-fine policy (row 4, column 1). We also see in column 1 that respondents dislike steeper-than-proportional fines when $\tau^{\text{fixed}} =$€200, though this result may be driven by the small value for the low-income fine (€35) under such a policy.[71] Despite the effects that emerge conditional on $\tau^{\text{fixed}} =$€200, only one of the coefficient differences between columns 1 and 2 is significant at the 5% level. In columns 3 and 4, we examine results conditional on different levels of behavior change and find neither significant effects nor discernible differences across the two columns. In columns 5 and 6, we examine the role of question order (whether respondents answer the transfer or fine comparison first) to test how salience and/or learning effects might influence our estimates. Within this comparison, we again find no significant differences.

Overall, across the 28 behavior and fine-schedule coefficients in table 7.5, only 3 are significant at the 5% level; furthermore, only one of the 17 across-group differences is significant at the 5% level.[72] In Appendix table 10.16, we explore heterogeneity with respect to both

[71]Note that we do not see any response to above-proportional fine schedules with $\tau^{\text{fixed}} =$€330, when the low-income fine is €60.

[72]This includes across-group differences in net WTP for the omitted policy.

Table 7.5: Heterogeneous Effects of Fine-Policy Characteristics on Hypothetical WTP (Randomized Features)

	Per-Capita Net WTP for Income-Based Fines					
	Fixed Fine		Behavior Change Δs		Question Order	
	(1)	(2)	(3)	(4)	(5)	(6)
	€200	€330	20	30	Transfer Q 1st	Fine Q 1st
Prop. Fines, Equal Speeding	59.92***	29.57	43.05***	23.46*	44.54**	46.70**
	(15.81)	(16.26)	(9.951)	(9.220)	(15.07)	(16.83)
Behavior (ds = 0.01)						
Equality ($s^{\text{inc}}_{\text{below}} = s^{\text{inc}}_{\text{above}}$)	-16.12	8.938	-1.682	-4.707	5.135	-10.59
	(10.04)	(10.43)	(10.41)	(10.15)	(9.520)	(10.87)
Inequality ($s^{\text{inc}}_{\text{below}} > s^{\text{inc}}_{\text{above}}$)	-4.643	-8.067	-16.02	1.623	-17.13	1.208
	(9.916)	(10.59)	(10.49)	(10.09)	(9.649)	(10.75)
Behavior Change Δs	-18.96*	3.472			-5.181	-11.21
	(8.055)	(8.566)			(7.758)	(8.770)
Fine Schedule						
Proportional ($b = b_{\text{prop}}$)	1.513	-10.20	-1.168	-9.167	-19.85*	7.014
	(9.863)	(10.41)	(10.26)	(10.07)	(9.517)	(10.68)
Above-Prop. ($b > b_{\text{prop}}$)	-22.85*	14.31	-6.743	-2.981	2.912	-11.72
	(9.937)	(10.64)	(10.47)	(10.19)	(9.494)	(11.00)
N	1358	1369	1356	1371	1313	1414

Estimates of equation 7.1 using the survey sample described in section 7.3. Standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

background characteristics—income, age, and education—and beliefs collected from respondents at the end of the survey. For the latter, we ask respondents (1) how often they think other drivers are aware of the speeding fine they would pay and (2) whether they believe speeding behavior is responsive to fine size.[73] Overall, we find remarkably little heterogeneity across sub-groups: net WTP for the omitted proportional-and-equal policy is consistently large and there are no significant differences within any given sub-group comparison.

Given this set of results, we conclude that our baseline findings in table 7.5 are ro-

[73]See Appendix section 10.2.6 for additional detail on these elicitations.

bust: respondents place high value on income-based fine policies but are insensitive to the behavioral and fine-schedule features of such policies over the range of variation we consider.

7.5.3 Implied Fine-Gradient

We now translate our estimates for the behavior and fine-schedule effects into implied fairness-driven fine gradients; we use these implied gradients to assess whether our estimates are sufficiently small and precise to reject fairness-driven rationales for income-based fines in Finland. Our procedure involves two steps. First, we use our survey estimates to recover the β_s and β_τ fairness parameters that appear in our formula for the optimal fine gradient (equation 3.2). Second, while β_s directly enters the expression for the optimal gradient, β_τ must be transformed into the appropriate units.[74] To accomplish the latter, we leverage the empirical moments from section 6.4 that we used to derive the optimal redistributive fine gradient.[75]

To recover an estimate of β_s, we note that, in our main results from column 2 of table 7.4, the behavior effects are approximately homogeneous: the effect of increasing the low-income speeding rate by a given amount (while simultaneously decreasing the high-income speeding rate by the same amount) is roughly constant. Therefore, we combine the three sources of behavioral variation to recover our estimate of β_s;[76] to properly scale this parameter, we use the identification results in Appendix table 10.13. Because our main fine-schedule effects in column 2 of table 7.4 are also homogeneous—*i.e.* equal-and-opposite fine changes have a constant effect on respondents' WTP—we take an analogous

[74]In particular, welfare effects arising from behavioral fairness are proportional to behavioral responses, whereas welfare effects arising from fine fairness are proportional to fine changes. We therefore divide the latter by the average fine semi-elasticity; see equation 3.2.

[75]In particular, we use the general expression for the optimal fine gradient (Appendix equation 10.3), the average fine elasticity from Kaila (2024), and the distributional statistics from panel c of Appendix table 10.12.

[76]In particular, our independent variable becomes $\underbrace{s^{\text{inc}}_{\text{below}} - s^{\text{fixed}}_{\text{below}}}_{\text{Behavior change}} + \underbrace{s^{\text{inc}}_{\text{below}} - 0.06}_{\text{Behavior distribution}}$, which captures both dimensions (behavior change and induced distribution) that we use to increase the speeding rate of the low-income group relative to the high-income group.

approach to identify β_τ.

Table 7.6 shows results from regressing respondents' net WTP on the unit-adjusted behavior and fine-schedule variation in our survey; the resulting coefficients yield estimates for the implied fine gradients. Column 1 presents OLS results, while column 2 uses a quantile (median) regression; to ensure that policy comparisons occur only within each fixed-fine level, we also include an indicator for the fixed-fine level in our regressions.[77] We note four main findings. First, consistent with our baseline findings in table 7.4 and our extensive-margin results from table 10.15, the implied slopes are negative; this arises because respondents express a slight distaste for fine policies with steeper schedules and higher speeding rates among the low-income group. Second, the implied slopes are extremely large; this follows because respondents are willing to sacrifice several euros *per-capita* for policies that ultimately affect a small share of the population. As is well-known, hypothetical WTP elicitations may be subject to inflation bias; while our differencing approach deals with some of this, some residual likely remains.[78] Third, after combining the various behavioral effects, we have enough power to reject behavior-related motives for income-based fines, though we caution that the combination of these coefficients was not pre-registered. Fourth, even after assuming a constant fine-schedule effect, we cannot reject a schedule-motivated fairness rationale. In particular, while the estimated fine-fairness coefficient is negative, the size of the standard errors are such that we cannot rule out the degree of income-dependence found in the day-fine region of the Finnish speeding schedule (€88, see figure 10.3). This uncertainty is emphasized in column 2, which estimates a quantile regression and finds a (median) fine fairness effect equal to 0, with large standard errors.

To assess both rationales together, we calculate a total implied gradient (last row of table

[77]If we did not do this, we would confound changes in the fine schedule with comparisons across different fixed-fine levels.

[78]E.g. if a respondent likes the income-dependent transfers half as much as the income-dependent fine policy, then inflation can persist if the base elicitation (how much they like the fine policy) is inflated.

Table 7.6: Implied Fairness-Driven Fine Gradients

	Optimal Fine Gradient b^* (per €10,000 of income) (1)	(2)
Behavioral Fairness	-331.9*	-222.2*
	(168.0)	(98.11)
Fine Fairness	-268.8	0
	(307.5)	(179.6)
N	2727	2727
Fixed-Fine Control	X	X
Estimation	OLS	Quantile
Total Gradient	-600.7	-222.2
	(350.2)	(204.5)

Estimates of the optimal fairness-driven fine gradients (equation 10.3) implied by our survey results (table 7.4); see section 7.5.3 for details on the construction of these estimates. Standard errors are in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

7.6) as the sum of the behavioral- and fine-fairness slopes; in our OLS results, the estimated day-fine gradient of €88 lies just outside our 95% confidence interval: [-1287.1, 85.7]. Therefore, considering the impact of both rationales on the optimal fine gradient, we can reject the degree of income-dependence in the day-fine region of the Finnish speeding-fine schedule.

In summary, when considered in aggregate, responses to increases in the low-income speeding rate (relative to the high-income rate) are precise enough to reject the behavioral rationale for income-based fines. Our estimates of the response to increases in the steepness of the fine schedule are more imprecise; we therefore cannot reject schedule-based rationales in isolation, though we do find that respondents place slightly lower value on fine policies with steeper fine schedules. Considering both rationales jointly (table 7.6), our results are precise enough to reject the degree of income-dependence in the day-fine region of the Finnish speeding-fine schedule. Therefore, our read is that the survey evidence overall is inconsistent with either behavior- or schedule-based fairness motives.

We now discuss some additional analyses to probe the robustness of our results.

7.6 Robustness

Safety Concerns While we emphasize to respondents that total speeding is unresponsive to which fine policy is implemented, respondents might therefore believe that the incidence of accident costs is different between the two policies (or that particular groups' speeding behavior is more dangerous). Therefore, to address this potential confound, we asked respondents, after making their slider decision, to provide free-text reasoning for why they made their particular fine-policy choice.[79] To examine the prevalence of safety-related concerns, we tag all responses that include one of the following strings: {safe, danger, accident}; we do not address the tone in which these strings are mentioned.[80] In total, only 113 responses (4%) include one of these strings; furthermore, average net WTP among this group is €28, not meaningfully different from all other responses (€32).[81] This evidence suggests that safety-related concerns do not drive our results.[82]

8 Discussion

We now conduct two simple counterfactual exercises in order to compare magnitudes of the effects we've identified. In particular, we estimate the welfare consequences of setting different fine gradients b using equation 10.3. First we consider decreasing the total fine

[79] *E.g.* someone who chooses the income-based fine policy sees the following prompt: *"Please use the text box below to describe the factors that influenced your decision to choose Income-based fines over Fixed fines."*

[80] *E.g.*, someone might write "Accidents are the same in both policies."

[81] In unreported results, we re-estimate equation 7.1 after dropping this group and find little difference in our coefficients.

[82] For an expanded notion of safety (which may also include fairness-motivated rationales), we also include responses with the string "deter;" this expands the number of safety-related responses to 230, nearly 10% of our sample. Again, this group has similar net WTP to the rest of the sample, and re-estimating equation 7.1 with or without these responses makes little difference.

gradient (measured by combining all speeding crime) from its observed level (€12.4 per €10,000 of earnings) to the optimum implied by redistribution- and externality-based rationales (-€19.6 under our preferred estimates from sections 5.4 and 6.4). This exercise clarifies the consequences of eliminating income-dependence altogether and using standard economic rationales to generate policy prescriptions. In our second counterfactual, we show the consequences of proportionality by increasing the fine gradient from €12.4 to €88, the latter being the gradient we estimate in the day-fine region of the Finnish fine schedule. Throughout, we assume that the fixed component of the fine schedule $a^*(b)$ and the income tax schedule $T(\cdot)$ optimally adjust; together, these imply that our counterfactuals amount to rotations of the fine schedule around response-weighted average earnings $\overline{z}^{rw}$.[83]

Throughout, we assume that our estimated behavioral responses (and moments of the income distribution) are constant over the range of changes we consider.[84] As a result, these exercises are merely suggestive.[85] We show the effects in figure 8.1; we separate the effects into revenue, externality (*i.e.* reallocation of optimal deterrence), and redistribution components. The first bar shows the effects of optimally setting the gradient to optimally redistribute and correct externalities; overall the effects are quite small. Most of the positive gains are on the redistribution side, with the revenue losses being quite modest. As a result, this reform implies a marginal value of public funds (MVPF) around 5.5. The proportionality reform, on the other hand, is quite distortionary, generating over €5.5 million in losses, which is around 8% of total spending on speeding fines per year in Finland.

How do these effects compare to respondents' WTP for income-based fine policies? On average, respondents are willing to forgo €144 million in government revenue in order

[83] In particular, equation 10.2 for the optimal fixed component implies that, at any level of b, the fine for those earning $\overline{z}^{rw}$ is held constant.

[84] For details on the moments we use, which are common with those used in our section 6.4, see table 10.12.

[85] A proper account would allow fine semi-elasticities to change as the gradient changes, for example.

Figure 8.1: Welfare Effects of Counterfactual Fine Gradients

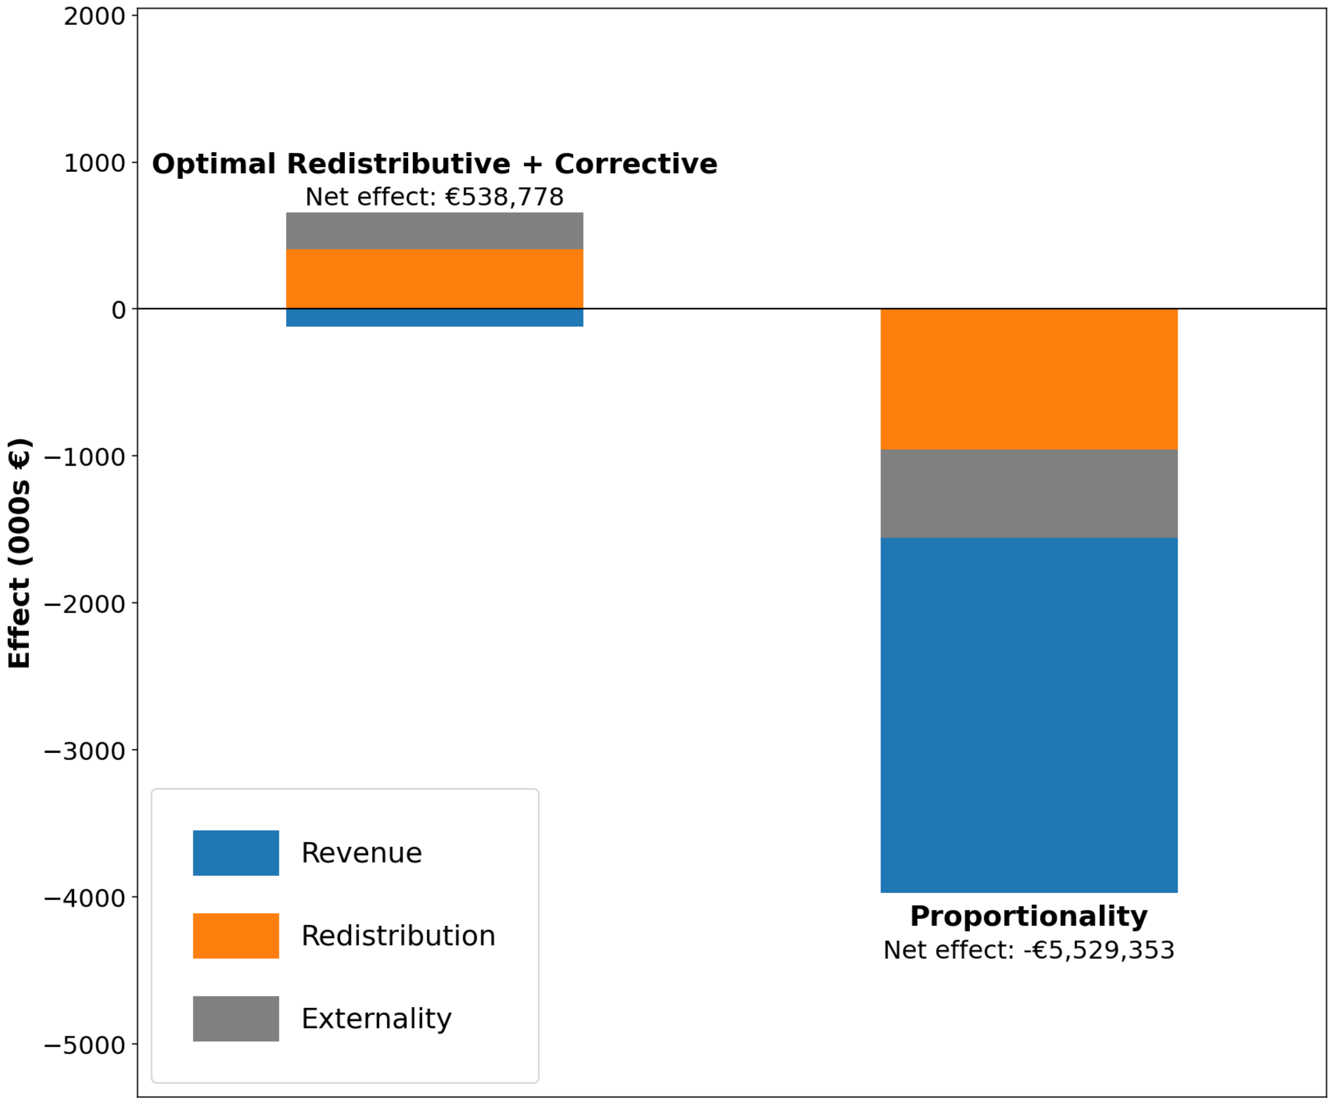

to implement income-based fine policies, and our results indicate that this WTP is relatively insensitive to fine-schedule steepness or induced behavior. Therefore, we assume this value would be forgone by moving to the optimal redistributive + corrective gradient and would not augment by moving to proportionality. As a result, the fairness losses of setting optimal redistributive + corrective fines would be orders of magnitude larger than losses on other dimensions; this is true even if we use the median WTP in our survey (€36 million). On the other hand, our survey results show no compelling fairness-driven reason to increase the steepness of the fine schedule; as a result, the proportionality counterfactual generates sizable large losses and no fairness gains.

Overall, if we take the WTP of respondents in our survey at face-value, the optimal fine gradient should include *some* modest level of income-dependence, but need not include the extremes found in the income-based region of the Finnish speeding-fine schedule, given the sizable welfare losses associated with such a system.[86]

9 Conclusion

In this paper, we evaluate rationales for a controversial policy: income-based fines. Using the Finnish speeding fine system as our laboratory, we first consider the Beckerian motive: does the marginal social cost of speeding scale with the income of the offender? Drawing on linked data from accident reports, income tax returns, and crime histories, we exploit cross-sectional variation in speeding histories conditional on income, together with a job-loss design, to recover the income-gradient of marginal social speeding costs. We find that this gradient is *negative*, leading us to reject this rationale.

We then evaluate a standard motivation in the public finance literature: in a Mirrleesian redistributive framework, fines can be exploited to lower the distortionary costs of re-

[86] In our survey, the most minimal income-based policy we pose to respondents includes a gradient of around €30, somewhat higher than the overall gradient of €12.4 but well below the proportional gradient of €88.

distribution if speeding is an informative signal for latent earnings ability. As clarified by recent work in optimal commodity taxation, this signal can be identified by differencing across-income variation in speeding rates with respect to causal earnings effects on speeding. To identify these causal earnings effects, we leverage two complementary strategies: within-individual earnings variation and differences in the timing and size of inheritance shocks. We find that the estimated earnings ability signal is *negatively* correlated with income, implying that Mirrleesian redistributive rationales can only justify a *lower* speeding fine for the rich.

Finally, we consider two fairness rationales alluded to in theoretical and statutory motivations for the policy: equal compliance across incomes and proportional punishment. To evaluate these motives, we use data from an original survey in which Finnish residents make trade-offs between fixed- and income-based fine policies, where we randomly vary the latter with respect to both schedule "steepness" and the induced across-income distribution of speeding rates. We find that income-based fines are popular: net of redistribution, respondents are willing to forgo €144 million in government revenue in order to implement an income-based speeding fine policy. However, these valuations are insensitive to both the distribution of behavior induced by fines and the steepness of the fine schedule itself. Therefore, our evidence is inconsistent with the equal-compliance and proportional-punishment rationales.

Overall, our evidence suggests that oft-cited rationales for income-based fines are insufficient to justify their popularity. Furthermore, we explicitly reject two natural economic objectives for such policies: externality mitigation and redistribution. Therefore, accounting for the popularity of income-based fines requires an appeal to either nonstandard economic motives (e.g. income-correlated mis-perceptions) or fairness preferences that are not tied to salient observable features of such policies.

References

Hunt Allcott, Benjamin B Lockwood, and Dmitry Taubinsky. Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *The Quarterly Journal of Economics*, 134(3):1557–1626, 05 2019. ISSN 0033-5533. doi: 10.1093/qje/qjz017. URL https://doi.org/10.1093/qje/qjz017.

Anthony B. Atkinson and Joseph E. Stiglitz. The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1):55–75, 1976. ISSN 0047-2727. doi: https://doi.org/10.1016/0047-2727(76)90041-4. URL https://www.sciencedirect.com/science/article/pii/0047272776900414.

Olivier Bargain, Mathias Dolls, Dirk Neumann, Andreas Peichl, and Sebastian Siegloch. Tax-benefit revealed social preferences in europe and the us. *Annals of Economics and Statistics*, (113/114):257–289, 2014. ISSN 21154430, 19683863. URL http://www.jstor.org/stable/10.15609/annaeconstat2009.113-114.257.

Gary S. Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217, 1968. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/1830482.

Gary E. Bolton and Axel Ockenfels. Erc: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1):166–193, March 2000. doi: 10.1257/aer.90.1.166. URL https://www.aeaweb.org/articles?id=10.1257/aer.90.1.166.

François Bourguignon and Amedeo Spadaro. Tax–benefit revealed social preferences. *The Journal of Economic Inequality*, 10(1):75–108, March 2012. ISSN 1573-8701. doi: 10.1007/s10888-010-9153-0. URL https://doi.org/10.1007/s10888-010-9153-0.

Aviv Caspi, Julia Gilman, and Charlie Rafkin. Preferences for rights. 2024.

Gary Charness and Matthew Rabin. Understanding social preferences with simple tests.

The Quarterly Journal of Economics, 117(3):817–869, 2002. ISSN 00335533, 15314650. URL http://www.jstor.org/stable/4132490.

Vidar Christiansen and Eilev S. Jansen. Implicit social preferences in the norwegian system of indirect taxation. *Journal of Public Economics*, 10(2):217–245, 1978. ISSN 0047-2727. doi: https://doi.org/10.1016/0047-2727(78)90036-1. URL https://www.sciencedirect.com/science/article/pii/0047272778900361.

John DiNardo, Nicole M. Fortin, and Thomas Lemieux. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5): 1001–1044, 1996. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2171954.

Dirk Engelmann and Martin Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, 94(4):857–869, September 2004. doi: 10.1257/0002828042002741. URL https://www.aeaweb.org/articles?id=10.1257/0002828042002741.

Ernst Fehr and Klaus M. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, 1999. ISSN 00335533, 15314650. URL http://www.jstor.org/stable/2586885.

Antoine Ferey, Benjamin Lockwood, and Dmitry Taubinsky. Sufficient statistics for nonlinear tax systems with general across-income heterogeneity. *American Economic Review*, Forthcoming, 2024. URL https://www.aeaweb.org/articles?id=10.1257/aer.20221053&from=f.

Nathaniel Hendren. Measuring economic efficiency using inverse-optimum weights. *Journal of Public Economics*, 187:104198, 2020. ISSN 0047-2727. doi: https://doi.org/10.1016/j.jpubeco.2020.104198. URL https://www.sciencedirect.com/science/article/pii/S0047272720300621.

Martti Kaila. How do people react to income-based fines? evidence from speeding tickets discontinuities. Working Paper 11064, CESifo, 2024. URL https://ssrn.com/abstract=4803851. Available at SSRN.

Elena Kantorowicz-Reznichenko. Day fines: Should the rich pay more? *Review of Law and Economics*, 11(2):481–501, 2013. doi: 10.2139/ssrn.2387442. URL http://dx.doi.org/10.2139/ssrn.2387442.

Louis Kaplow. On the undesirability of commodity taxation even when income taxation is not optimal. *Journal of Public Economics*, 90(6):1235–1250, 2006. ISSN 0047-2727. doi: https://doi.org/10.1016/j.jpubeco.2005.07.001. URL https://www.sciencedirect.com/science/article/pii/S0047272705001131.

Raimo Lahti. *Day Fines in Finland*, page 24–43. Cambridge University Press, 2021.

Benjamin B. Lockwood and Matthew Weinzierl. Positive and normative judgments implicit in u.s. tax policy, and the costs of unequal growth and recessions. *Journal of Monetary Economics*, 77:30–47, 2016. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2015.10.006. URL https://www.sciencedirect.com/science/article/pii/S0304393215001257. "Inequality, Institutions, and Redistribution"held at the Stern School of Business,New York University,April 24-25, 2015.

J. A. Mirrlees. An exploration in the theory of optimum income taxation12. *The Review of Economic Studies*, 38(2):175–208, 04 1971. ISSN 0034-6527. doi: 10.2307/2296779. URL https://doi.org/10.2307/2296779.

Arash Nekoei and David Seim. How Do Inheritances Shape Wealth Inequality? Theory and Evidence from Sweden. *The Review of Economic Studies*, 90(1):463–498, 03 2022. ISSN 0034-6527. doi: 10.1093/restud/rdac016. URL https://doi.org/10.1093/restud/rdac016.

Charlie Rafkin and Evan Soltas. Social preferences and bargaining failure in eviction. Working paper, 2025.

Alex Rees-Jones and Dmitry Taubinsky. Measuring "schmeduling". *The Review of Economic Studies*, 87(5):2399–2438, 09 2019. ISSN 0034-6527. doi: 10.1093/restud/rdz045. URL `https://doi.org/10.1093/restud/rdz045`.

Emmanuel Saez. The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes. *Journal of Public Economics*, 83(2):217–230, 2002a. ISSN 0047-2727. doi: https://doi.org/10.1016/S0047-2727(00)00159-6. URL `https://www.sciencedirect.com/science/article/pii/S0047272700001596`.

Emmanuel Saez. Optimal income transfer programs: Intensive versus extensive labor supply responses. *The Quarterly Journal of Economics*, 117(3):1039–1073, 2002b. ISSN 00335533, 15314650. URL `http://www.jstor.org/stable/4132495`.

James Tobin. On limiting the domain of inequality. *The Journal of Law Economics*, 13 (2):263–277, 1970. ISSN 00222186, 15375285. URL `http://www.jstor.org/stable/725025`.

Arthur van Benthem. What is the optimal speed limit on freeways? *Journal of Public Economics*, 124:44–62, 2015. ISSN 0047-2727. doi: https://doi.org/10.1016/j.jpubeco.2015.02.001. URL `https://www.sciencedirect.com/science/article/pii/S0047272715000146`.

Till von Wachter and Stefan Bender. In the right place at the wrong time: The role of firms and luck in young workers' careers. *American Economic Review*, 96(5): 1679–1705, December 2006. doi: 10.1257/aer.96.5.1679. URL `https://www.aeaweb.org/articles?id=10.1257/aer.96.5.1679`.

10 Appendix

10.1 Expressions for Fines and Taxes

Given $\tau(z) = a + b * z$, the optimal marginal income tax rate satisfies the following at each earnings level z^0:

$$T'(z^0) + b * s(z^0) + [\tau(z^0) - \overline{e}'(z^0)]s'_{inc}(z^0) = \frac{\int_{z \geq z^0} (\hat{g}(z) - 1)dH(z)}{h(z^0) z'_{T_z}(z^0)} \tag{10.1}$$

where $z'_{T_z}(z^0)$ is the (semi-)elasticity of labor supply with respect to the marginal income tax rate and $h(\cdot)$ is the earnings density function. Under an optimal income-tax schedule, if assumption 3.1 holds, then the fixed component of fines satisfies the following for any given b:

$$a^*(b) = \frac{\mathbb{E}[s'_{\tau_s}(z) * \overline{e}'(z)] + \mathrm{Cov}[s'(z) - s'_{inc}(z), \hat{g}(z)] - b * \mathbb{E}[s'_{\tau_s}(z) * z]}{\overline{s}'_{\tau_s}} \tag{10.2}$$

We now write the general expression for the optimal fine gradient that under a constant fine elasticity ζ_τ (to replace the semi-elasticities in our previous equations). To reflect our empirical approches, we make the following assumption:

Assumption 10.1. *Externality $\overline{e}(z)$, behavioral fairness $\overline{e}_s(z)$, and fine fairness $\overline{e}_\tau(z)$ are all composed of intercept and slope terms:*

$$\overline{e}(z) = \alpha_e + \beta_e * z$$
$$\overline{e}_s(z) = \alpha_s + \beta_s * z$$
$$\overline{e}_\tau(z) = \alpha_\tau - \beta_\tau * z$$

Furthermore, either $\alpha_\tau \leq 0$ or there exists some earnings level z_0 such that $\overline{e}_\tau(z_0) = 0$.

The last part of assumption 10.1 guarantees that β_τ determines the sign that fine-fairness contributes to the optimal fine gradient by ensuring that the intercept α_τ is not too large.

Then, if the income tax schedule $T(z)$ satisfies equation 10.1, the fixed-fine component $a^*(b)$ satisfies equation 10.2, and the fine elasticity ζ_τ is constant, then assumption 10.1 implies that the optimal fine gradient satisfies:

$$b^* = \beta_e + \beta_s + \frac{-\beta_\tau \int_Z z * (z - \overline{z}^{rw}) dH(z) + \alpha_\tau(\overline{z} - \overline{z}^{rw}) + \text{Cov}\Big[\hat{g}(z), \int_0^z (y - \overline{z}^{rw}) s'_{het}(y) dy\Big]}{\zeta_\tau \int_Z z * (z - \overline{z}^{rw}) \frac{s(z)}{\tau(z)} dH(z)} \tag{10.3}$$

where $\overline{z}^{rw} := \frac{\mathbb{E}[\frac{s(z)}{\tau(z)} * z]}{\overline{s}/\overline{\tau}}$ is response-weighted average income.

10.2 Empirical Appendix

10.2.1 Imputing Driving Distance Using Odometer Data

We start with a quarterly panel of vehicles from the vehicle registration data and keep only quarters in which we observe an odometer reading for a vehicle. We then group our data into vehicle-by-primary-user cells. Each of these cells constitutes a panel where we observe odometer readings over time for each vehicle-user. Between any two successive odometer readings for a vehicle-user cell, we calculate distance driven between inspections. We then uniformly assign this distance across quarters between the two inspection dates, assuming that inspections occur at the beginning of each quarter.[87] Aggregating across vehicles and quarters, we obtain a user-by-year panel of driving distance. Driving distance is non-missing for a given user-year (i, t) in our panel if and only if at least one quarter in year t falls between two inspections for a given vehicle registered to user i. Note that, if a user switches from an old vehicle h to a new vehicle h', we will not be able to calculate distance driven between the last inspection date for the old vehicle h and the first inspection date for the new vehicle h'.

[87] Assuming the inspections occur at the end of the quarter makes no difference for our results.

10.2.2 Procedure for Calculating Money-Metric Cost of Injuries and Deaths from Accidents

We proceed as follows

1. For each driver j involved in an accident a_t in year t, we identify the set of externality-relevant accident victims $\{v_1(j, a_t), v_2(j, a_t), ..., v_n(j, a_t\}$.

2. For each victim $v_m(j, a_t)$ we randomly select two "control" individuals, $c_1(v_m)$ and $c_2(v_m)$, from the FOLK population register who have the same birth year, sex, and education level as v_m in year $t-1$

3. We compute disposable income differences for each victim-year

$$\{\Delta\hat{y}_{v_m,t-1}, \Delta\hat{y}_{v_m,t}, ..., \Delta\hat{y}_{v_m,2023}\}$$

   until 2023, where

$$\Delta\hat{y}_{v_m,t+r} \coloneqq y_{v_m,t+r} - [(y_{c_1(v_m),t+r} + y_{c_2(v_m),t+r})/2]$$

   for all $r \in \{-1, 0, 1, ..., 2023-t\}$

4. Similarly, we compute differences in net transfers received for each victim-year

$$\{\Delta\hat{T}_{v_m,t-1}, \Delta\hat{T}_{v_m,t}, ..., \Delta\hat{T}_{v_m,2023}\}$$

   until 2023, where $\Delta\hat{T}_{v_m,t+r} \coloneqq T_{v_m,t+r} - [(T_{c_1(v_m),t+r} + T_{c_2(v_m),t+r})/2]$ for all $r \in \{-1, 0, 1, ..., 202$
   $t\}$

5. We calculate victim-specific DID net-income losses in each event year $r \geq 0$ as $\Delta\hat{y}^{DID}_{v_m,t+r} \coloneqq \Delta\hat{y}_{v_m,t+r} - \Delta\hat{y}_{v_m,t-1}$. We analogously calculate victim-specific DID for net transfer received, $\Delta\hat{T}^{DID}_{v_m,t+r}$

6. We calculate a welfare-weighted, NPV cost for each victim v_m as

$$M_{v_m} := \sum_{k=0}^{2023-t} \delta^k [\underbrace{\Delta \hat{T}^{DID}_{v_m,t+k}}_{\text{Transfers Paid by Govt.}} - \underbrace{g(y_{v_m,t-1}) * \Delta \hat{y}^{DID}_{v_m,t+k}}_{\text{Welfare-weighted income gain}}]$$

under some exponential discount rate δ and $g(\cdot)$ denotes a set of social marginal welfare weights

7. The total welfare-weighted external income loss imposed by driver j in accident a_t is $M^{ext}_j(a_t) := \sum_{m=1}^{n} M_{v_m}$

10.2.3 Bias in Income Effects Estimated Using Inheritance Shocks

Let α be the death effect (assumed homogeneous), η_q be the income effect for quintile q, and ΔW_q be the size of the total inheritance shock received by quintile q. Then, for a given q, the total effect of parental death is $\alpha + \eta_q * \Delta W_q$. Differencing effects across two quintiles q and q' yields $\eta_q * \Delta W_q - \eta_{q'} * \Delta W_{q'} = (\eta_q - \eta_{q'}) * \Delta W_q + \eta_{q'}(\Delta W_q - \Delta W_{q'})$. Dividing by $\Delta W_q - \Delta W_{q'}$ (*i.e.* the first-stage) yields $\eta_{q'} + (\eta_q - \eta_{q'}) * \frac{1}{1 - \Delta W_{q'}/\Delta W_q}$, which equals η_q if $\Delta W_{q'} = 0$, $\eta_{q'} = \eta_q$, or $\eta_{q'} = 0$. Furthermore, as $\Delta W_{q'}/\Delta W_q \to 0$, this expression converges to η_q.

10.2.4 Construction of Survey Transfer Policies

Let $\{\bar{s}^{inc}_{below}, \bar{s}^{inc}_{above}\}$ be speeding rates under the income-dependent fine policy and $d\tau = \tau^{inc}_{above} - \tau^{fixed}_{above}$ be the fine difference for the high-income group. We ask individuals to compare a fixed transfer policy that gives 25 € to all Finnish adults with an income-dependent transfer that allocates $25 + d\tau * s^{inc}_{below}$ to the low-income group and $25 - d\tau * s^{inc}_{above}$ to the high-income group. We use the speeding rates under the income-dependent policy, rather than those under the fixed-fine policy, since this choice guarantees that revenue effects are the same for income-dependent transfer and fine policies. To see this, note that the revenue effect of the income-dependent fine policy is given by:

$$\Delta R_{fine} = 0.5[s^{inc}_{below} * \tau^{inc}_{below} + s^{inc}_{above} * \tau^{inc}_{above}] - 0.5 * \tau^{fixed} * [s^{fixed}_{below} + s^{fixed}_{above}]$$

Since total speeding is the same under both policies, we can replace $s^{fixed}_{below} + s^{fixed}_{above}$ with $s^{inc}_{below} + s^{inc}_{above}$ and it immediately follows that:

$$\Delta R_{fine} = 0.5 * d\tau * [s^{inc}_{above} - s^{inc}_{below}]$$

Thus, increasing the low-income transfer by $d\tau * s^{inc}_{below}$ (*i.e.* spending an additional $d\tau * s^{inc}_{below}$ per low-income person) and lowering the transfer by $d\tau * s^{inc}_{above}$ for the high-income group yields an equivalent revenue effect.

In general, our method for constructing the income-dependent transfer will tend to inflate the role of redistribution when (1) "altruism" weights α on each group are positive and (2) $\alpha_{below} > \alpha_{above}$. In particular, for baseline speeding rates $\{s_{below}, s_{above}\}$, the welfare effect of transfers $\{d\tau * s_{below}, -d\tau * s_{above}$ is given by:

$$dW = d\tau * \alpha_{below} * (s_{below} - \frac{\alpha_{above}}{\alpha_{below}} * s_{above})$$

With $\alpha_{above} > 0$ and $\alpha_{above} < \alpha_{below}$, this value will be maximized when s_{below} is large and s_{above} is small. Since the income-dependent policies we consider increase the speeding rate for the low-income and decrease it for the high-income group, dW will be larger if we use the speeding rates under the income-dependent policy to construct the income-dependent transfer, making our estimates for the value of income-dependent fines more conservative. This result also holds if $\alpha_{above} < 0$ with $\alpha_{below} > |\alpha_{above}| > 0$.

10.2.5 Choice of Fine Policy Parameters

Average Incomes: We set the above- and below-median incomes (€15,000 and €51,000 respectively) using the 2023 distribution of earned income; we lower the income of the low-income group relative to these calculations in order to make proportional fines more apparent in the income-based policy.

$\tau^{\textbf{fixed}} = \mathbf{330}$: We calculate the higher fixed-fine of €330 as the average of the proportional fines €150 and €510; these proportional fines, which are 1% of average income for each group, are reflective of the degree of income-dependence we estimate on-average within the day-fine region of the Finnish fine schedule (see figure 10.3.

$\tau^{\textbf{fixed}} = \mathbf{200}$: The fixed-fine of €200 is close to the average speeding fine received in Finland (around 215 €). We calculate proportional income-based fines relative to this average fine level using a slightly lower proportionality rate of 0.6%, which guarantees that the fine difference between the fixed and income-based policies is the same in relative terms for both levels of the average fine. In particular, (200-90)/200 $\approx$ (330-150)/330.

$\mathbf{d}\tau$: We choose the size of the fine perturbations $d\tau$ (€90 for $\tau^{fixed} = 330$ and €55 for $\tau^{fixed} = 200$) around the proportional policy to be roughly equivalent relative to the level of the fixed fine.

Speeding Rate: We choose the average speeding rate to reflect the average in Finland in 2019, our last complete year of data.

10.2.6 Survey Belief Elicitations

Fine Awareness We ask respondents to indicate how often (from never to always) they believe "drivers are aware of the fine they would receive if they were caught speeding." In Appendix table 10.16, we separate respondents into those who answer "rarely" or "never" (column 3) and those who answer at least as often as "sometimes" (column 4); see Appendix figure 10.24 for the distribution of responses.

Fine Responsiveness We ask respondents to indicate how much they agree with the following statement:

"Drivers will speed less if speeding fines are increased."

In Appendix table 10.16, we group respondents into {"strongly disagree," "disagree," "neither a (column 5) and {"agree," "strongly agree"} (column 6); the distribution of responses is shown in Appendix figure 10.25.

10.2.7 Counterfactual Analysis

When the fixed component of the fine schedule $a^*(b)$ is optimal and the income tax $T(\cdot)$ is optimal, then a small increase in the fine slope b yields a welfare effect (ignoring fairness considerations) given by:

$$\begin{aligned} dW =&(b - \beta^M) * db * \int_Z s'_{\tau_s}(z) * z * (z - \overline{z}^{rw}) dH(z) \\ &- db * \mathrm{Cov}\big[\hat{g}(z), \int_{z_{min}}^{z} (y - \overline{z}^{rw})\big(s'(y) - s'_{inc}(y)\big)\big] \end{aligned}$$

where β^M is the income-gradient of the marginal social cost of speeding (assumed constant). To calculate the revenue, efficiency, and redistributive consequences of a change in the slope from b to b^{alt}, we integrate the expression above.

10.3 Supplemental Tables

Table 10.1: Income-Gradient of Excess Speed

	Excess Speed				
	(1)	(2)	(3)	(4)	(5)
Earned Income (€10,000s)	-0.342***	-0.0175***	-0.0436***	-0.00708**	-0.00959**
	(0.00358)	(0.00371)	(0.00366)	(0.00224)	(0.00348)
Constant	15.66***	14.46***	14.56***	11.23***	14.43***
	(0.0164)	(0.0166)	(0.0163)	(0.0103)	(0.0155)
N	753881	753877	753876	603762	753876
Age FE		X			
Age-Year FE			X	X	X
Camera Only				X	
Municipality FE					X

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.2: Summary Statistics for Collapsed Accident Panel

	Mean (SD)
Earned Income (10,000s €)	2.576 (1.788)
Car Ownership Rate	0.508 (0.446)
Speeding Rate	0.0573 (0.104)
Ever Receive Speeding Ticket (%)	35.12 (47.73)
Other Traffic Crime Rate	0.0238 (0.0657)
Ever Receive Other Traffic Ticket (%)	18.17 (38.56)
Ever in Accident (%)	7.965 (27.08)
Ever in Injurious Accident (%)	1.765 (13.17)
Ever Injured in Accident (%)	0.798 (8.898)
N	5033598

Table 10.3: Marginal Speeding Gradient (Other-Person Injuries)

	(1) Other-Person Injuries (per 1,000 drivers)	(2)	(3)	(4)	(5)
Speeding Rate (MSC Intercept)	10.70***	9.202***	10.58***	2.911***	1.498***
	(0.198)	(0.330)	(0.199)	(0.207)	(0.254)
Speeding X Earnings (MSC Gradient)	-1.103***	-0.837***	-1.105***	-0.463***	-0.253***
	(0.0445)	(0.0695)	(0.0445)	(0.0463)	(0.0593)
Earned Income (10,000s €)	0.145***	0.107***	0.157***	0.148***	0.176***
	(0.00479)	(0.0131)	(0.00491)	(0.00698)	(0.0289)
Constant	0.338***	0.680***	0.316***	-0.171***	-0.124
	(0.0128)	(0.0546)	(0.0129)	(0.0163)	(0.100)
N	5033597	1588308	5033597	5033597	2223285
Cohort FE	X	X	X	X	X
Reverse-Causality Controls	X	X	X	X	X
Ever Speeding		X			
Municipality FE			X		
Car Ownership & Traffic Crime Controls				X	X
Annual Driving Distance					X

Heteroskedasticity-robust standard errors in parentheses. Car ownership rates, non-speeding crime rates, and driving distance are also interacted with earnings when included. Each individual is weighted by the number of years they appear in the population register. In column 5, this weight is multiplied by the number of years for which driving distance information is available from 2013-2018. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.4: Conditional Marginal Speeding Gradient

	(1) Other-Party Injuries (per 1,000 accidents)	(2)
Speeding Rate (MSC Intercept)	18.64	2.581
	(9.811)	(8.853)
Speeding X Earnings (MSC Gradient)	-10.85***	-3.872
	(2.626)	(2.370)
Pre-Period Earnings (10,000s €)	5.235**	-0.471
	(1.965)	(1.773)
Other-Party Participants		153.7***
		(0.574)
Constant	129.9***	-7.789
	(6.855)	(6.207)
N	315026	315026
Age FE	X	X
Year FE	X	X
Reverse Causality	X	X
Car Ownership & Other Traffic Crime	X	X
Vehicle Mass	X	X

Sample consists of the first accident in which individuals from our cross-sectional sample (described in section 5.1) are involved as the driver of a motor vehicle (car, truck, bus, moped, motorbike, motorcycle, or van) from 2006-2018. Earnings are measured in the year prior to the accident. Vehicle mass is measured as the average mass of vehicles registered to a given individual from 2013-2018. "Reverse Causality Controls" refers to a control for the total number of accidents in which an individual is involved from 2006-2018. For descriptions of the remaining variables, see section 5.1. Car ownership rates, non-speeding crime rates, driving distance, and vehicle mass controls include both a baseline effect and an interaction of these variables with earnings. Each individual is weighted by the number of years they appear in the population register from 2006-2018, multiplied by the number of years for which vehicle mass information is available from 2013-2018. Standard errors are in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.5: Cumulative Social Costs of Accidents

	(1) Non-Fatal Injuries	(2) Fatal Injuries
# of Post-periods	12.98	12.91
Age at accident	44.43	50.92
Cumulative Fiscal Cost	€13,250	-€41,957
Cumulative Net Inc. Loss	€8,791	€244,961
$\delta = 0.03$		
Discounted Fiscal Cost	€10,564	-€35,261
Discounted Net Inc. Loss	€6,473	€199,547
$\delta = 0.05$		
Discounted Fiscal Cost	€9,226	-€31,817
Discounted Net Inc. Loss	€5,358	€176,700
$\delta = 0.07$		
Discounted Fiscal Cost	€8,149	-€28,973
Discounted Net Inc. Loss	€4,486	€158,157
N	78378	3489

Table 10.6: Income Loss Gradients, Separately by ν

	Cumulative Net Income Losses			
	(1)	(2)	(3)	(4)
	$\nu = 0$	$\nu = 0.25$	$\nu = 1$	$\nu = 4$
Pre-Period Net Income	3.707***	2.497***	1.826**	-1.176
	(0.591)	(0.230)	(0.601)	(1.363)
Constant	-46385.8***	-30074.9***	-39834.9*	45111.5
	(9882.7)	(3743.1)	(16552.4)	(49712.7)
N	100761	99803	99803	99803
R^2	0.027	0.028	0.007	0.004
Age-Year FE	X	X	X	X

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.7: Marginal Speeding Cost Gradients by Discount Rate and Inequality Aversion: Other-Person Victims

	Discount Rate		
	3%	5%	7%
Welfare Weight			
$\nu = 0$	-21.45*	-18.45*	-15.99*
	(8.425)	(7.302)	(6.398)
$\nu = 0.25$	-18.06*	-15.55*	-13.5*
	(8.146)	(7.079)	(6.215)
$\nu = 1$	-9.027	-7.74	-6.703
	(4.809)	(4.193)	(3.696)
$\nu = 4$	-9.064	-7.761	-6.713
	(4.783)	(4.169)	(3.674)
N	5033597	5033597	5033597
Cohort FE	X	X	X
Reverse Causality	X	X	X
Car Ownership	X	X	X
Other Traffic Crime	X	X	X
Annual Driving Distance	X	X	X

Heteroskedasticity-robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.8: DDD-IV Income Effect Estimates, Excluding Transport-Related Causes-of-Death

	(1) Reduced-Form	(2) First-Stage	(3) IV
DDD Coefficient	0.311***	1.424***	
	(0.0818)	(0.0118)	
Inherit. Shock (10,000s €)			0.219***
			(0.0575)
Constant	7.361***	0.00648***	7.359***
	(0.0348)	(0.000552)	(0.0349)
N	10491513	10491513	10491513

Standard errors (clustered at the individual level) in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.9: DDD-IV Income Effect Estimates with Controls

	(1) Number of Heirs	(2) Number of Revisions
Inherit. Shock (10,000s €)	0.279^{***}	0.258^{**}
	(0.0831)	(0.0876)
Constant	6.944^{***}	7.313^{***}
	(0.0944)	(0.0541)
N	4912424	4912424
DDD-IV	X	X

Table 10.10: DDD-IV Income Effect Estimates: 5-year Effect Horizon

	(1) Reduced-Form	(2) First-Stage	(3) IV
DDD Coefficient	0.364^{***}	1.219^{***}	
	(0.108)	(0.0272)	
Inherit. Shock (10,000s €)			0.299^{***}
			(0.0892)
Constant	7.358^{***}	0.00471^{***}	7.356^{***}
	(0.0401)	(0.000674)	(0.0402)
N	10179538	10179538	10179538

Table 10.11: Labor Supply Effect: Top Inheritance Group

	(1)	(2)
	Speeding Rate (pp)	
Taxable Earned Income (10,000s €)	1.165***	0.474***
	(0.0133)	(0.0267)
Re-scaled Labor Supply (10,000s €)		0.281***
		(0.0227)
Constant	3.285***	4.426***
	(0.0389)	(0.0751)
N	9593075	9593075
R^2	0.014	0.206
Cohort-Year FE	X	X
Individual FE		X

Standard errors (clustered at the individual level) in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.12: Redistributive Value: Sufficient Statistics

	Earnings Group							
	1	2	3	4	5	6	7	8
Panel (a): Earnings Effects & Cross-Sectional Variation								
Inherit. Income Effect	0.211	0.211	0.211	0.211	0.287	0.287	0.287	0.287
	(0.0866)	(0.0866)	(0.0866)	(0.0866)	(0.105)	(0.105)	(0.105)	(0.105)
FE Labor Supply Effect	0.599	0.281	0.261	0.283	0.345	0.277	0.478	0.167
	(0.0917)	(0.0619)	(0.0778)	(0.0852)	(0.0936)	(0.0840)	(0.0845)	(0.0576)
$s'_{inc}(z)$	0.810	0.492	0.472	0.494	0.632	0.564	0.765	0.454
$s'(z)$	0.333	0.513	1.328	2.753	1.675	1.784	0.890	0.376
Panel (b): Earnings Ability Tag								
$s'_{het}(z) := s'(z) - s'_{inc}(z)$	-0.477	0.021	0.856	2.259	1.043	1.220	0.125	-0.078
Panel (c): Additional Moments (2012)								
Avg. Earnings (€)	2350	13720	20405	25231	29943	35992	44927	72940
Avg. Disposable Income (€)	8143	14803	18600	21561	24707	28423	33764	50834
Speeding Rate (%)	4.368	4.747	5.090	5.731	7.028	8.041	9.635	12.129
Observed Fine Level (€)	173	188	196	202	208	215	226	261
$\overline{z}^{rw}$	€35,124							

Table 10.13: Fairness Effects of Equal-and-Opposite Fine Reforms

	$z < z_{med}$	$z > z_{med}$	Total
Fine Changes	$-d\tau$	$d\tau$	0
Fine-Fairness Effects	$0.5(\alpha_\tau - \beta_\tau \overline{z}_{below})d\tau$	$-0.5(\alpha_\tau - \beta_\tau \overline{z}_{above})d\tau$	$\beta_\tau * d\tau * 0.5 * \Delta\overline{z}$
Behavioral Responses	ds	$-ds$	0
Behavioral-Fairness Effects	$-0.5(\alpha_s + \beta_s \overline{z}_{below})ds$	$0.5(\alpha_s + \beta_s \overline{z}_{above})ds$	$\beta_s * ds * 0.5 * \Delta\overline{z}$
Residual	$\beta_0 - \beta_0 * I(b > 0)$		

Table 10.14: Survey Respondent Summary Statistics

	(1) All Responses	(2) Main Sample
Income Group (%)		
0 € - 14,999 € (1st quintile)	20.3	19.6
15,000 € - 24,999 € (2nd quintile)	16.3	16.4
25,000 € - 34,999 € (3rd quintile)	20.9	21.0
35,000 € - 49,999 € (4th quintile)	22.3	22.7
50,000 € or above (5th quintile)	20.2	20.2
Education completed (%)		
Primary Education or lower	7.0	6.2
Secondary Education	48.1	47.7
Higher Education	44.9	46.1
Age group (%)		
18-34 (1st quartile)	27.1	26.3
35-50 (2nd quartile)	26.0	25.9
51-65 (3rd quartile)	21.3	21.3
Over 65 (4th quartile)	25.6	26.4
Gender (%)		
Man	49.9	49.3
Woman	49.3	49.9
Other/decline to answer	0.8	0.8
Driving frequency (%)		
Never	13.0	12.3
Less than once per month	10.6	10.4
Monthly	6.8	6.8
Weekly	69.5	70.5
I(Speeding ticket in last 2 years) (%)	7.27	7.00
HH Car Ownership Rate (%)	79.3	79.9
Median completion time (mins.)	12.9	13.05
N	3038	2727

Table 10.15: Extensive Margin Regressions

	Share WTP > 0 (pp)	
	(1) $\mathrm{WTP}^{\text{inc-fine}}$	(2) $\mathrm{WTP}^{\text{net}}$
Proportional Fines, Equal Speeding, $\Delta s = 0.02$	79.69***	58.26***
	(3.125)	(3.674)
Behavior (ds = 0.01)		
Equality ($s^{\text{inc}}_{\text{below}} = s^{\text{inc}}_{\text{above}}$)	-2.922	-2.625
	(1.996)	(2.346)
Inequality ($s^{\text{inc}}_{\text{below}} > s^{\text{inc}}_{\text{above}}$)	-1.888	-2.130
	(1.998)	(2.348)
Behavior Change Δs	-2.018	-2.831
	(1.619)	(1.903)
Fine Schedule (dτ =€72.5)		
Proportional ($b = b_{\text{prop}}$)	-0.537	-3.791
	(1.976)	(2.323)
Above-Proportional ($b > b_{\text{prop}}$)	-1.395	0.929
	(2.007)	(2.359)
N	2727	

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 10.16: Heterogeneity by Background Characteristics

	Outcome: Net WTP									
	Income		Fine Awareness Beliefs		Fine Responsiveness Beliefs		Age		Education	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	<35,000 €	≥35,000 €	Unaware	Aware	Unresponsive	Responsive	≤ 50	> 50	Secondary Educ.	Higher Educ.
Proportional Fines & Equal Speeding	52.19***	38.59*	45.76*	44.71**	34.32	53.96***	32.46*	61.11***	38.65*	52.54**
	(14.90)	(17.51)	(18.08)	(14.60)	(17.81)	(14.68)	(14.83)	(17.50)	(16.63)	(17.05)
Behavior										
$\Delta s_{below}^{inc} = 0.01$ ($s_{below}^{inc} < s_{above}^{inc}$)	1.240	-10.72	6.412	-10.27	-2.134	-5.123	-14.33	9.088	-7.467	1.205
	(9.572)	(11.11)	(11.61)	(9.313)	(11.37)	(9.376)	(9.405)	(11.14)	(10.43)	(10.82)
$\Delta s_{below}^{inc} = 0.01$ ($s_{below}^{inc} = s_{above}^{inc}$)	-11.07	-2.505	-10.46	-4.758	-16.16	0.907	-1.558	-13.72	-6.580	-8.769
	(9.434)	(11.35)	(11.79)	(9.216)	(11.26)	(9.463)	(10.03)	(10.82)	(10.48)	(11.02)
$\Delta s_{below}^{fixed} = 0.01$	-5.707	-12.60	-9.041	-7.361	-14.56	-3.927	-3.668	-13.91	-0.723	-15.84
	(7.712)	(9.092)	(9.411)	(7.553)	(9.147)	(7.650)	(7.817)	(8.924)	(8.370)	(8.912)
Fine Schedule										
$\Delta\tau_{below}$=€72.5 (Below-Proportional)	-6.203	-1.836	-2.291	-7.086	-15.94	4.450	-9.195	0.299	-12.59	-1.197
	(9.487)	(11.00)	(11.46)	(9.223)	(11.07)	(9.409)	(9.350)	(10.97)	(10.60)	(10.67)
$\Delta\tau_{below}$=€72.5 (Proportional)	-2.974	-8.436	9.145	-13.78	21.97	-24.52**	-7.855	-0.887	4.913	-19.78
	(9.617)	(11.20)	(11.77)	(9.293)	(11.33)	(9.494)	(9.559)	(10.85)	(10.37)	(10.91)
N	1555	1172	1110	1617	1117	1610	1425	1302	1300	1257

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

10.4 Supplemental Figures

Figure 10.1: Hypothetical Fine Schedule (100 km/h limit)

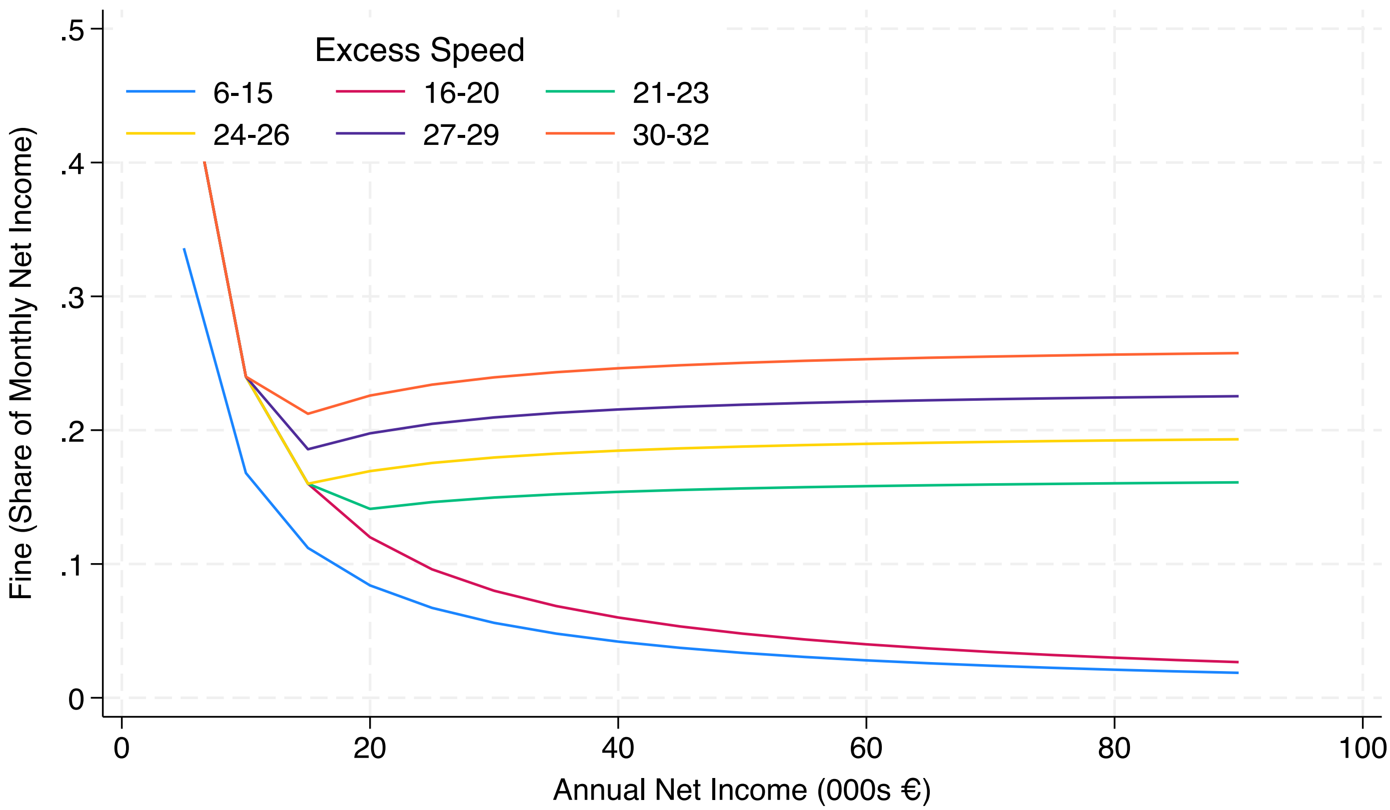

Net-of-tax income consists of earned and capital income, as well as any transfers received, after the payment of taxes. We assume no dependents for the formula in equation 2.1.

Figure 10.2: Binscatter Relationships Between Income and Speeding Crime Characteristics

(a) Speed Limit

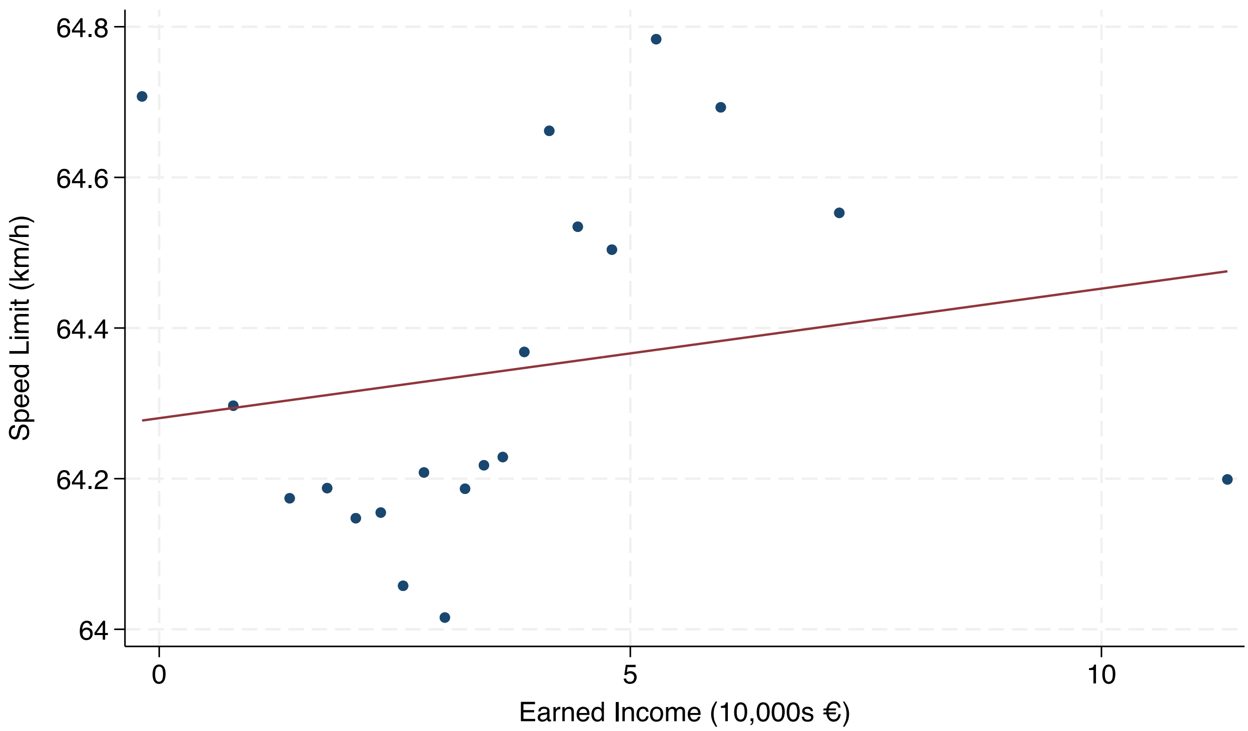

(b) Excess Speed

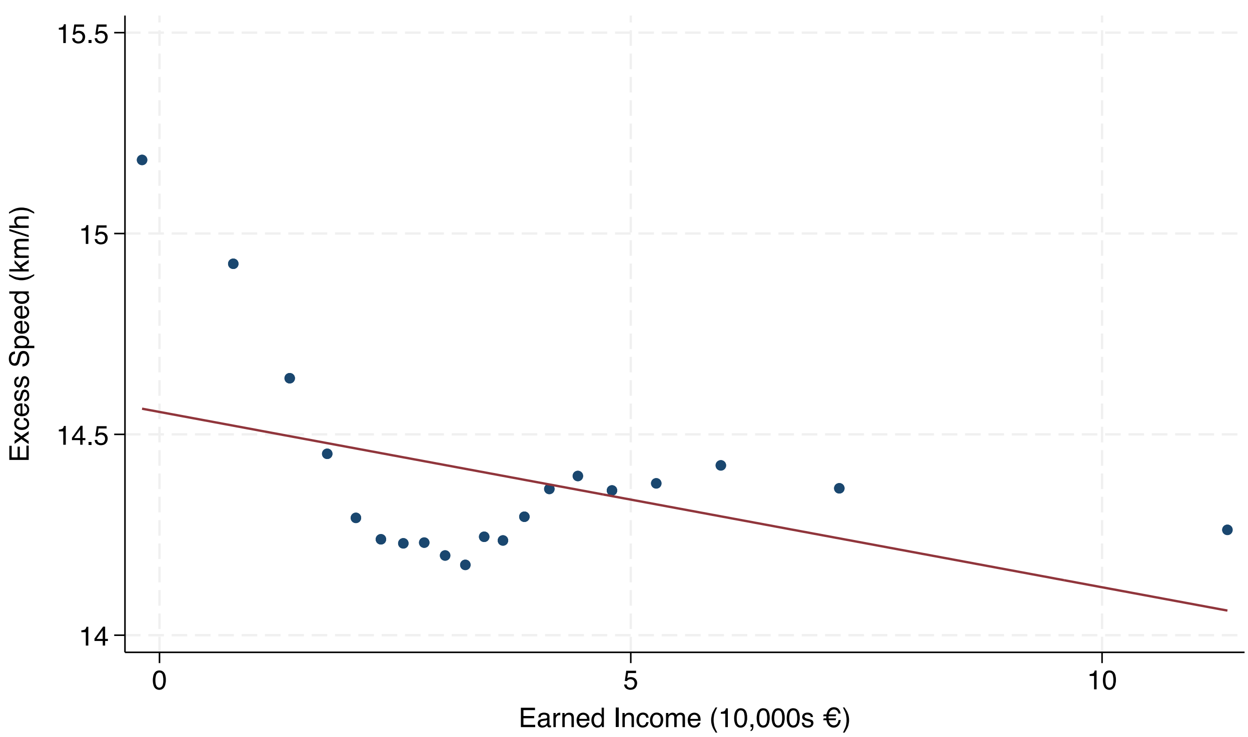

Figure 10.3: Day Fine Slope

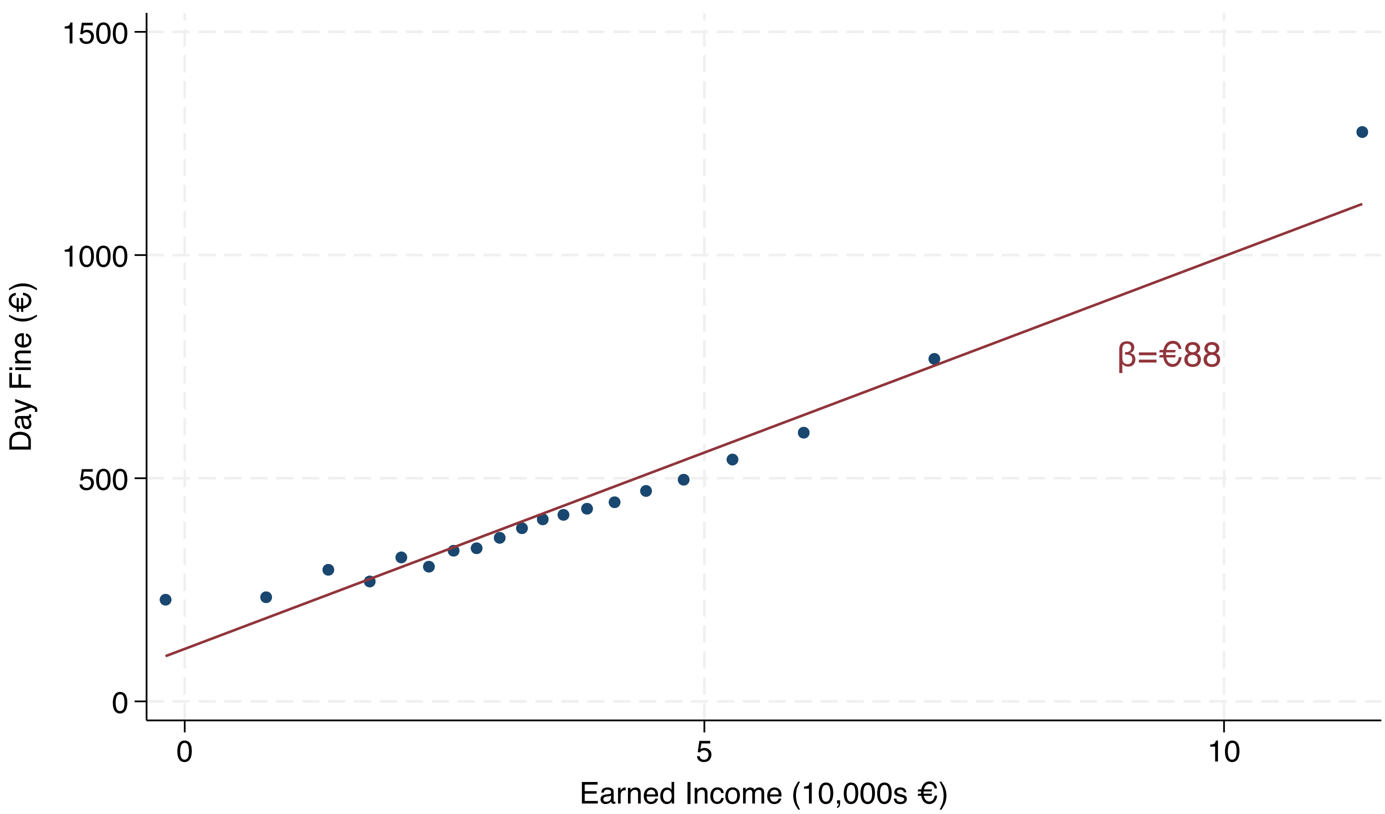

Binscatter based on police ticket data from 2018-2020. Income and the fine amount are residualized on age-by-year fixed-effects.

Figure 10.4: Overall Fine Slope

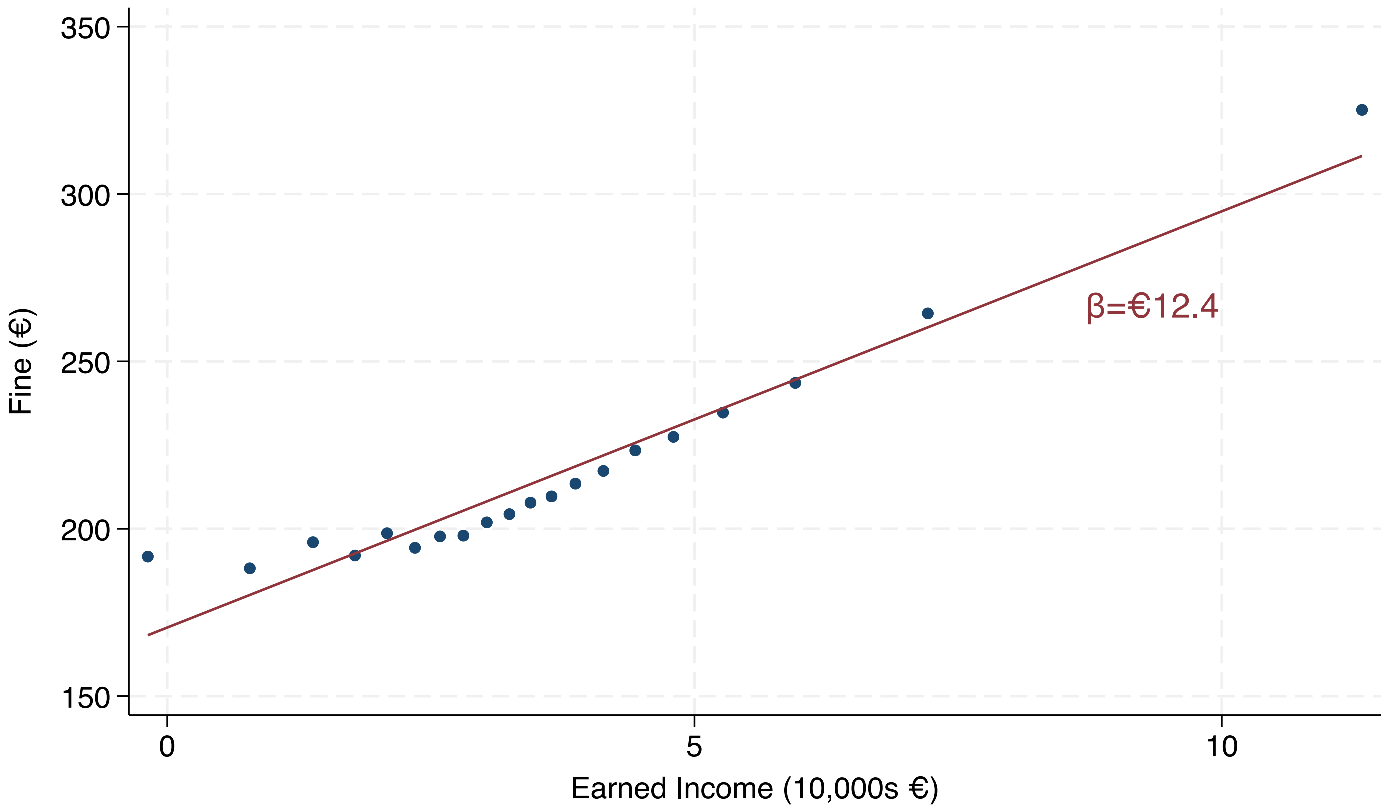

Binscatter based on police ticket data from 2018-2020. Income and the fine amount are residualized on age-by-year fixed-effects.

Figure 10.5: Binscatter Correlations Between Driver and Victim Income

(a) Other-Party Victims

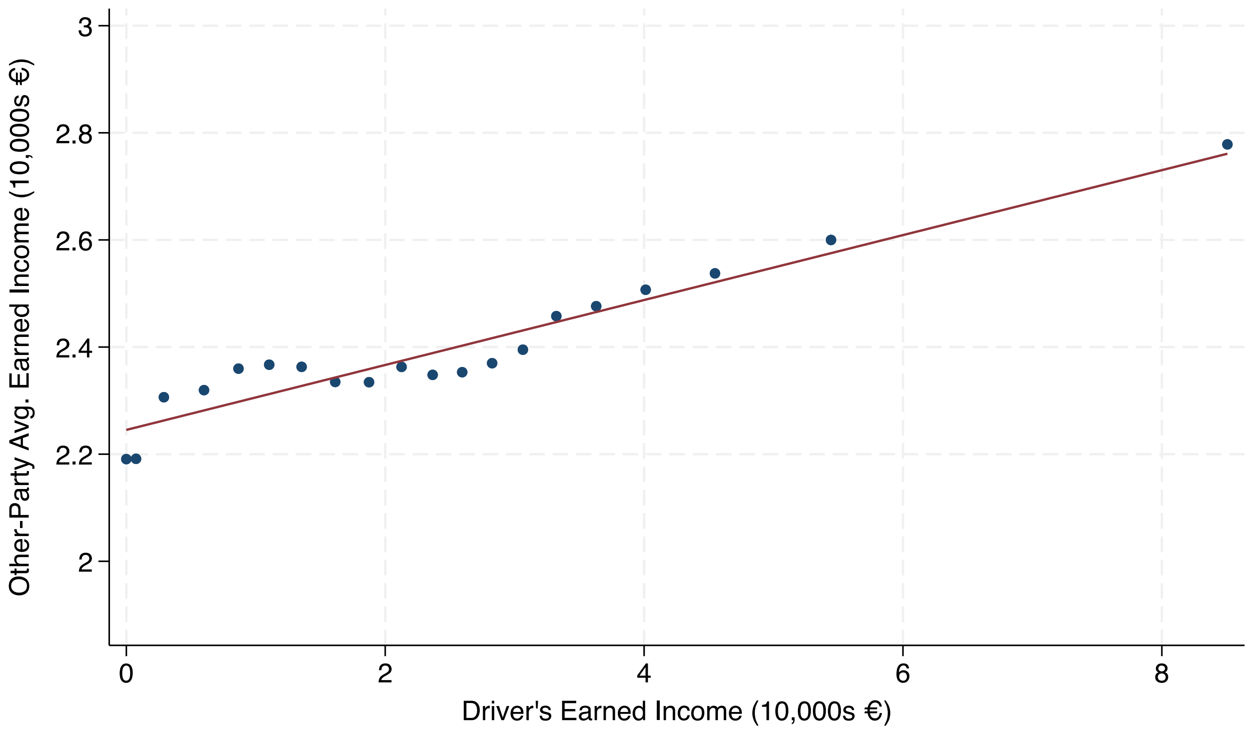

(b) Other-Person Victims

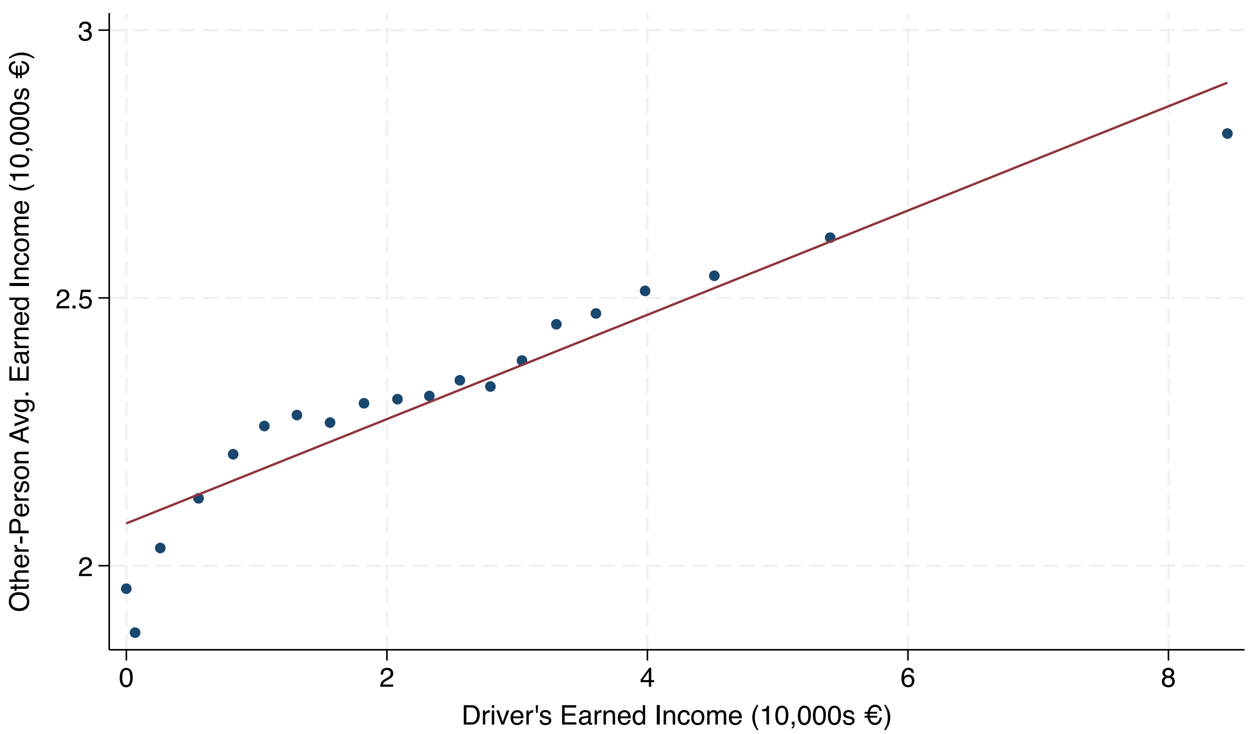

Binscatter relationships between driver income and victim income. We use the sample of individuals described in 5.1 and restrict to the first accident in which an individual is involved as a driver from 2006-2018. We further restrict to drivers of a car, truck, bus, moped, motor bike, motorcycle, or van and exclude accidents involving trains, trams, tractors, snow mobiles, trailers, or other unidentified motor vehicles. Earnings are measured in the year prior to the accident.

Figure 10.6: Event Study Effects of Injury/Death on Disposable Income

(a) Drivers

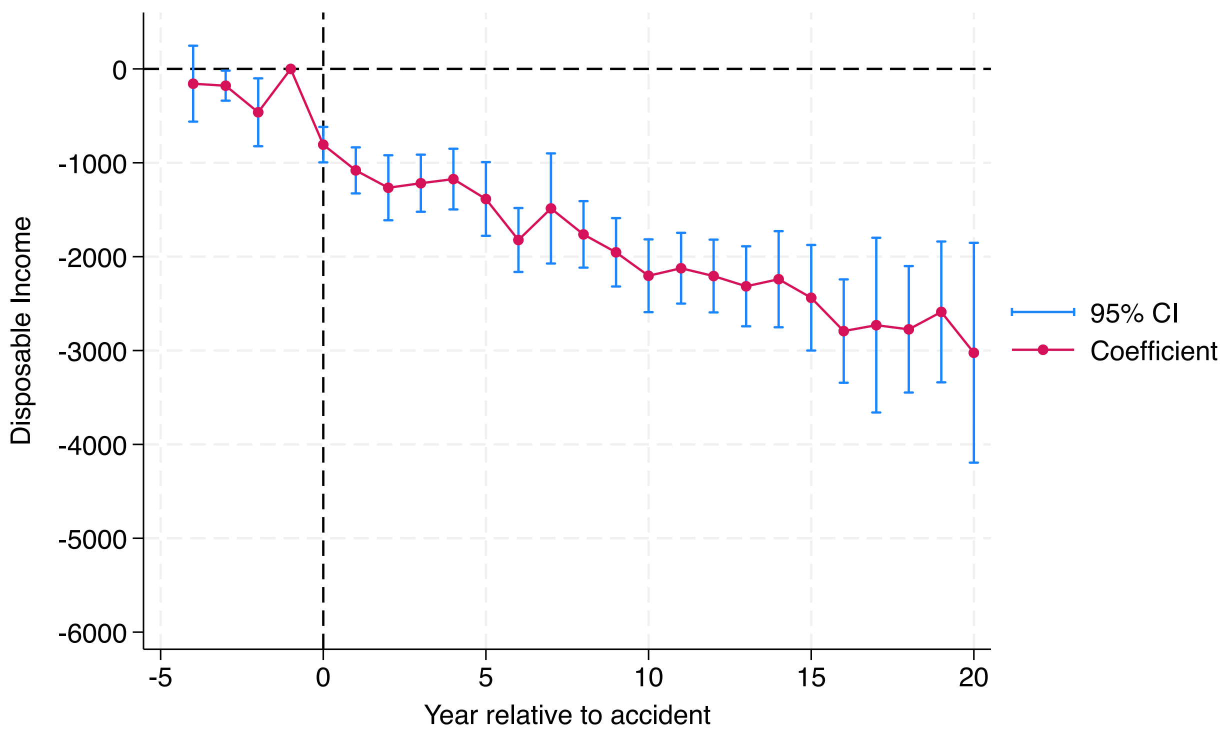

(b) Passengers

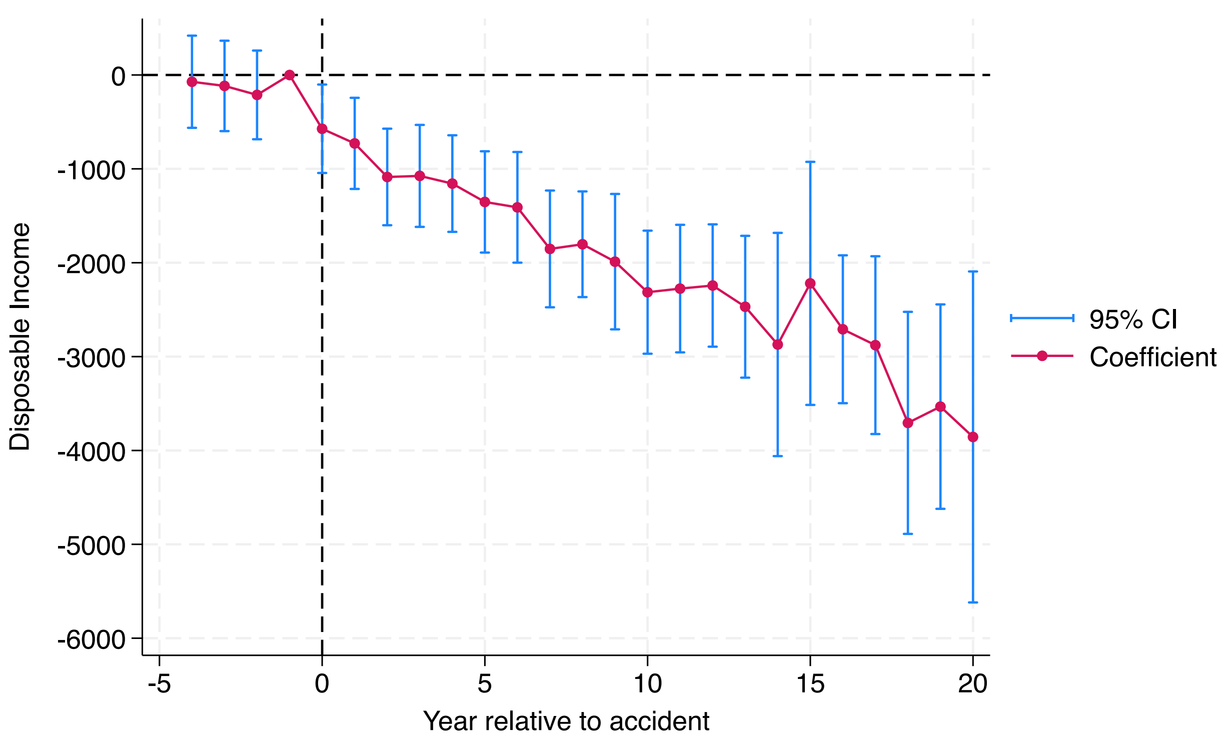

Figure 10.7: Event Study Effects of Non-Fatal Injuries on Probability of Working

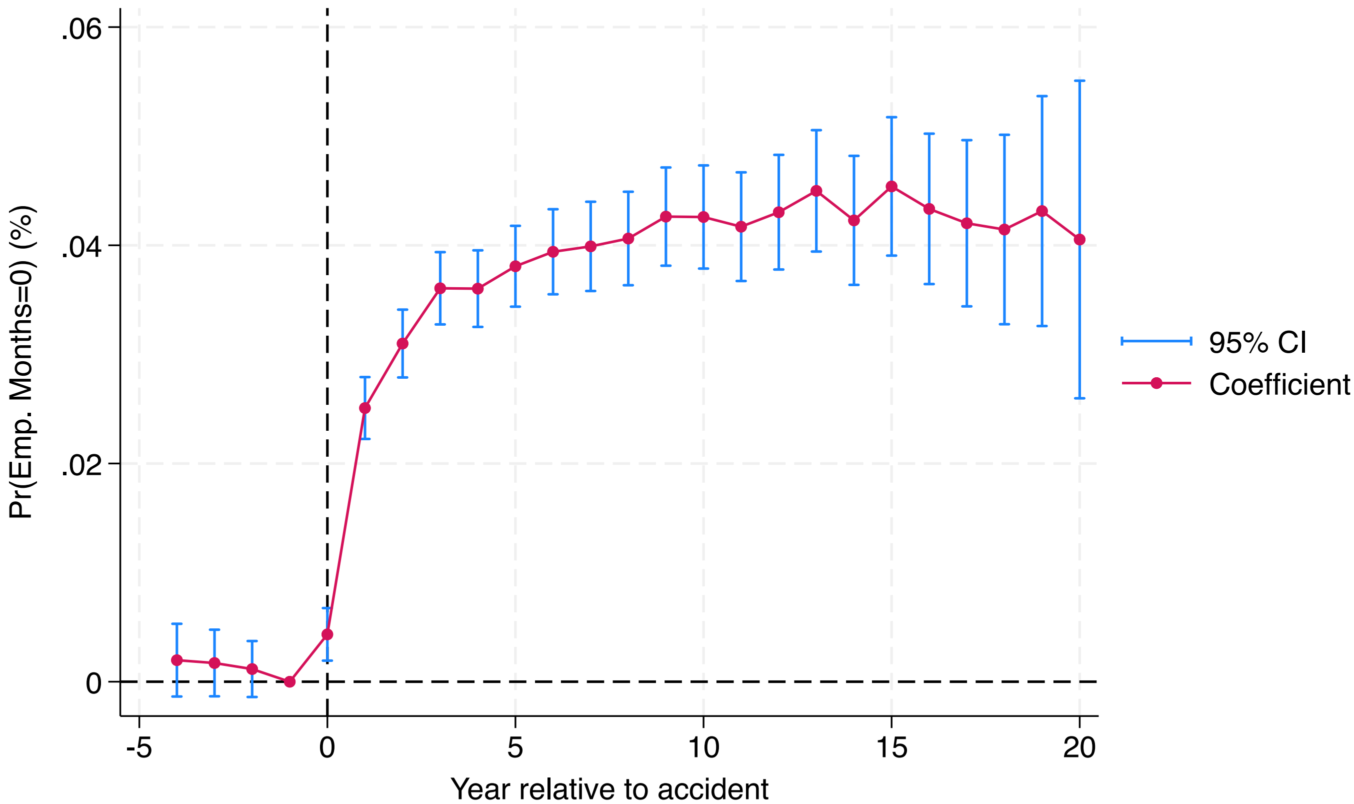

Estimates of β_k from equation 5.2. The sample includes all accident-related non-fatal injuries from 2003-2022, restricted to accidents that include a motor vehicle (cars, trucks, buses, vans, mopeds, motor bikes, and motorcycles). Injuries from accidents involving a snow mobile, trailer, tram, train, tractor, or other unidentified motor vehicles are excluded.

Figure 10.8: Event Study Effects of Injury/Death on Transfers/Taxes

(a) Transfers Received

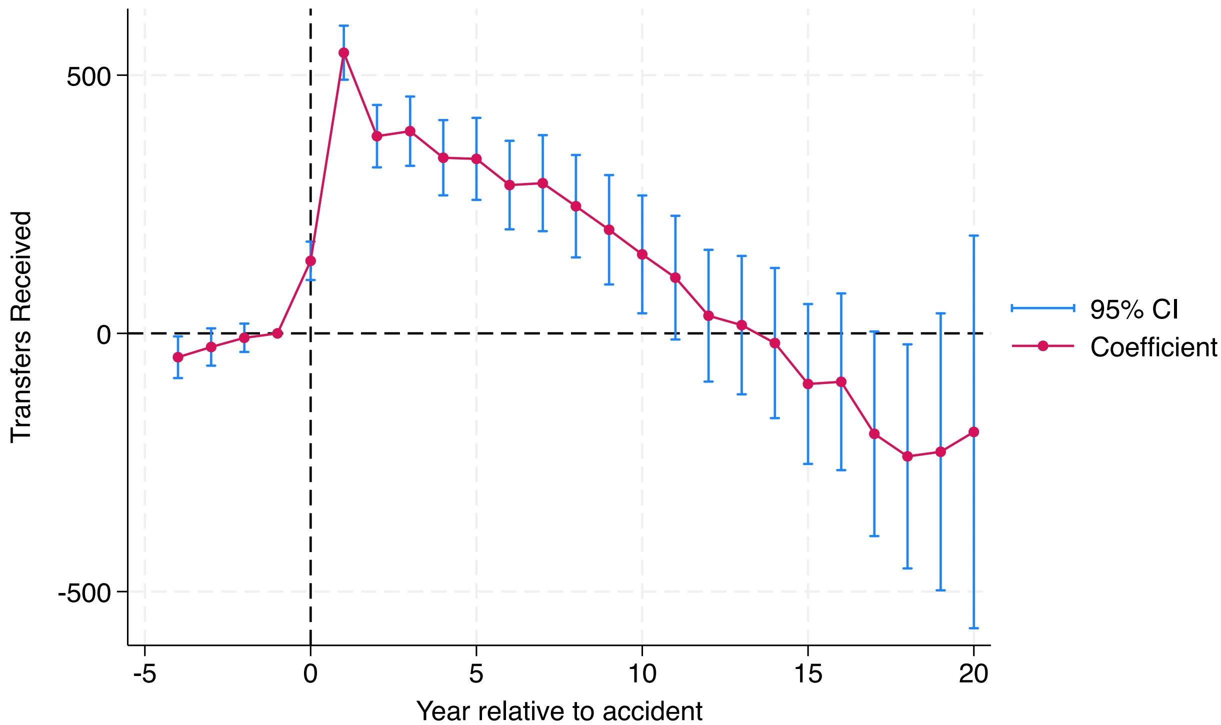

(b) Taxes Paid

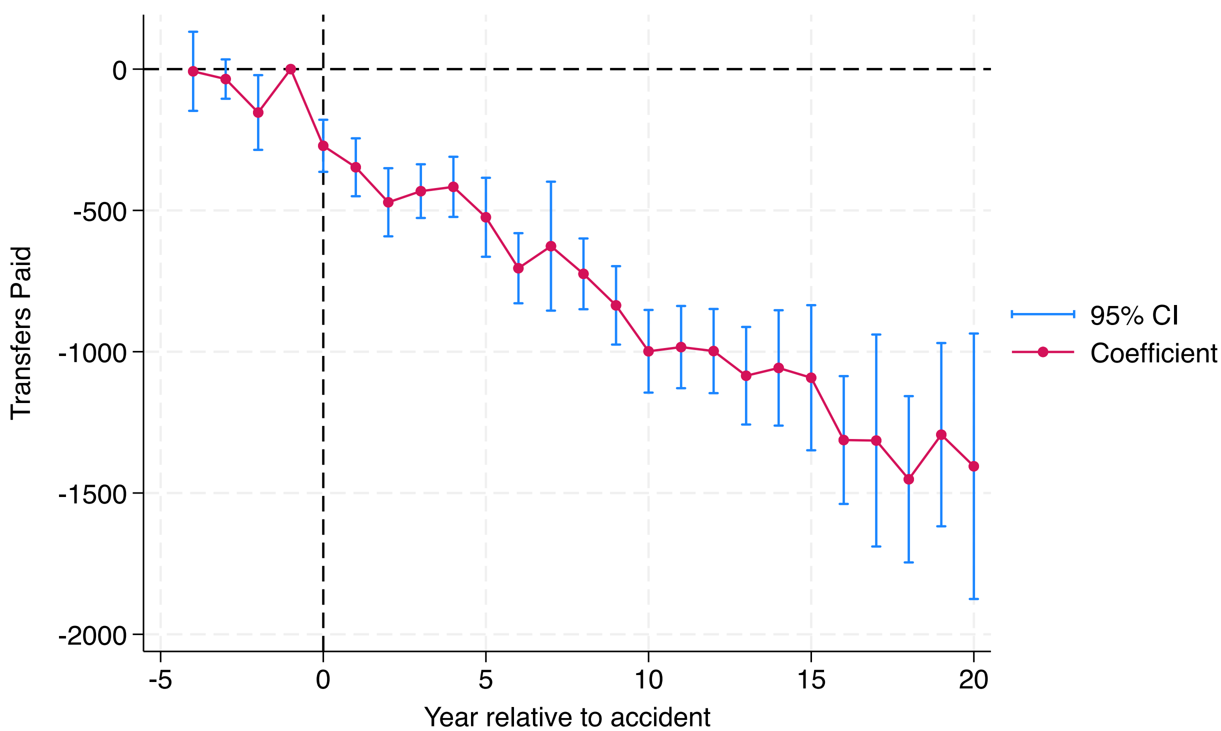

Estimates of β_k from equation 5.2. The sample includes all accident-related injuries from 2003-2022, restricted to accidents that include a motor vehicle (cars, trucks, buses, vans, mopeds, motor bikes, and motorcycles). Injuries from accidents involving a snow mobile, trailer, tram, train, tractor, or other unidentified motor vehicles are excluded.

Figure 10.9: Event Study Effects of Non-Fatal Injury on Disposable Income

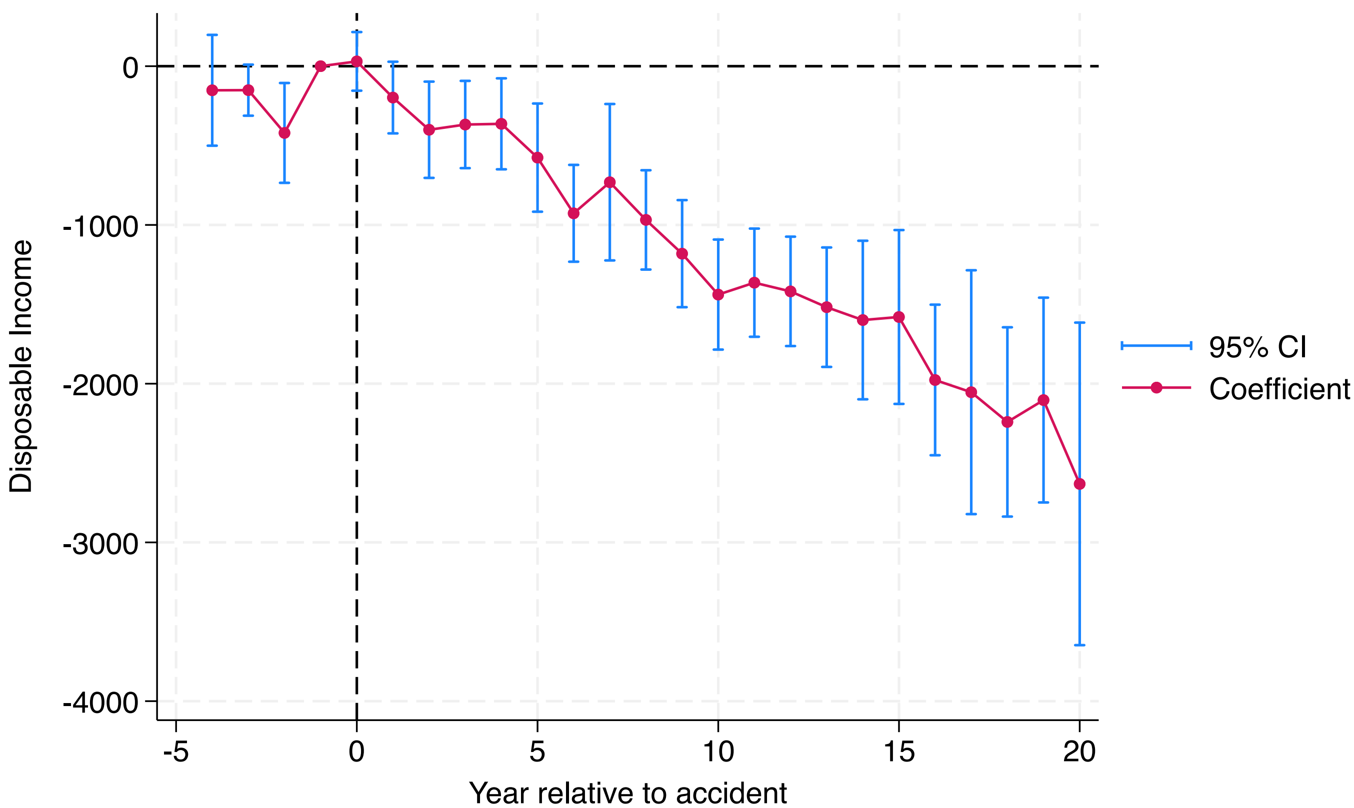

Estimates of β_k from equation 5.2. The sample includes all non-fatal accident-related injuries from 2003-2022, restricted to accidents that include a motor vehicle (cars, trucks, buses, vans, mopeds, motor bikes, and motorcycles). Injuries from accidents involving a snow mobile, trailer, tram, train, tractor, or other unidentified motor vehicles are excluded.

Figure 10.10: Event Study Effects of Non-Fatal Injuries on Social Costs

Figure 10.11: Cumulative Social Costs of Injury/Death, by Pre-Period Income

(a) Fiscal Externality

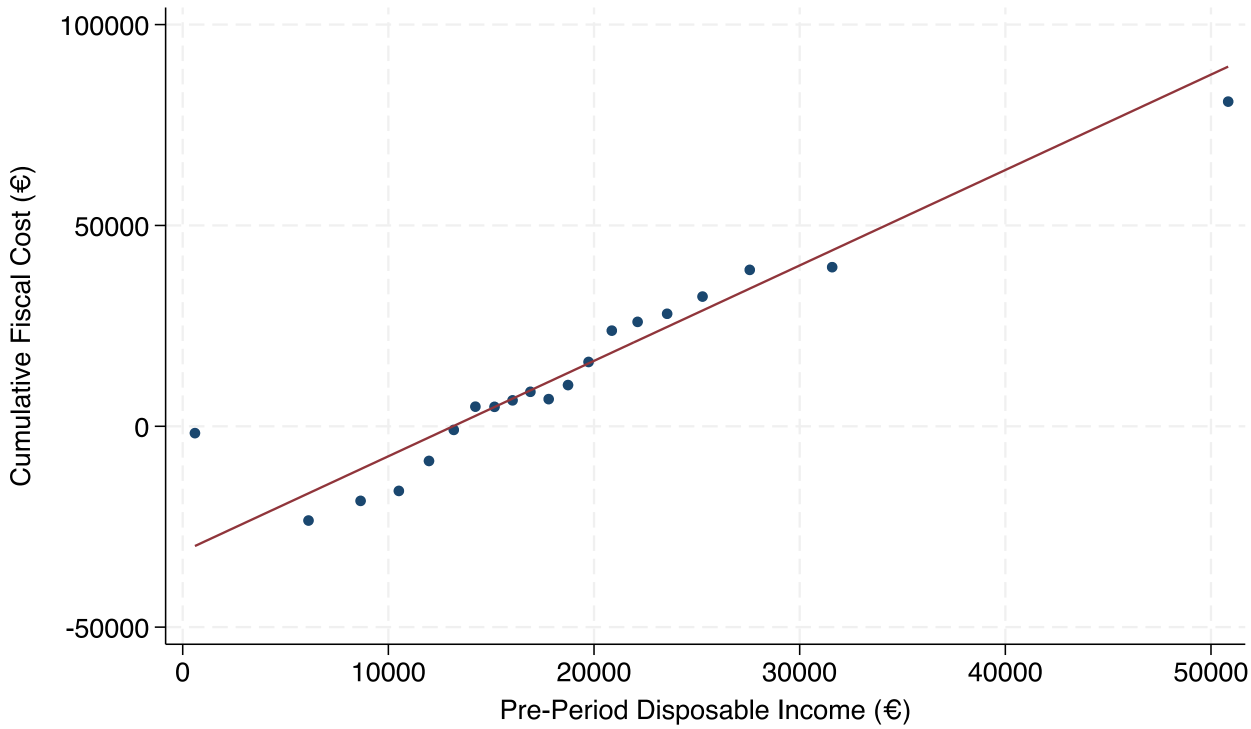

(b) Disposable Income Loss

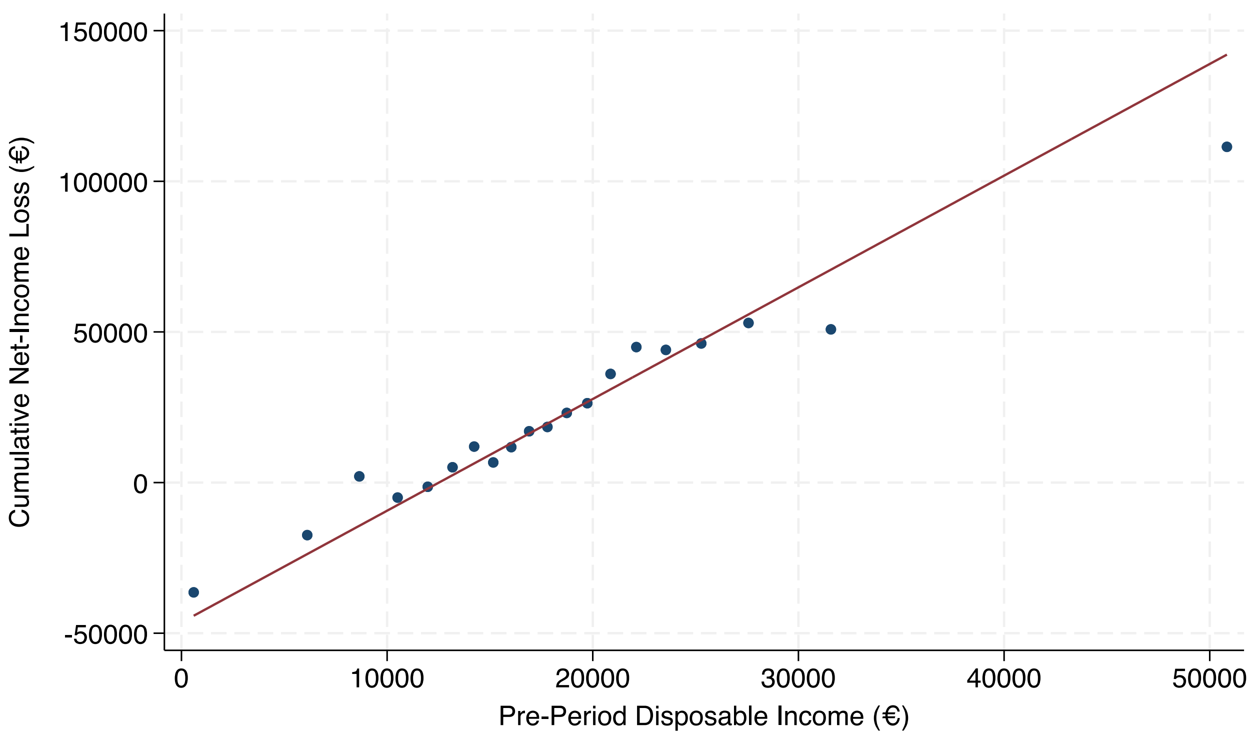

Figure 10.12: Event Study of Taxable Earned Income around Parental Death

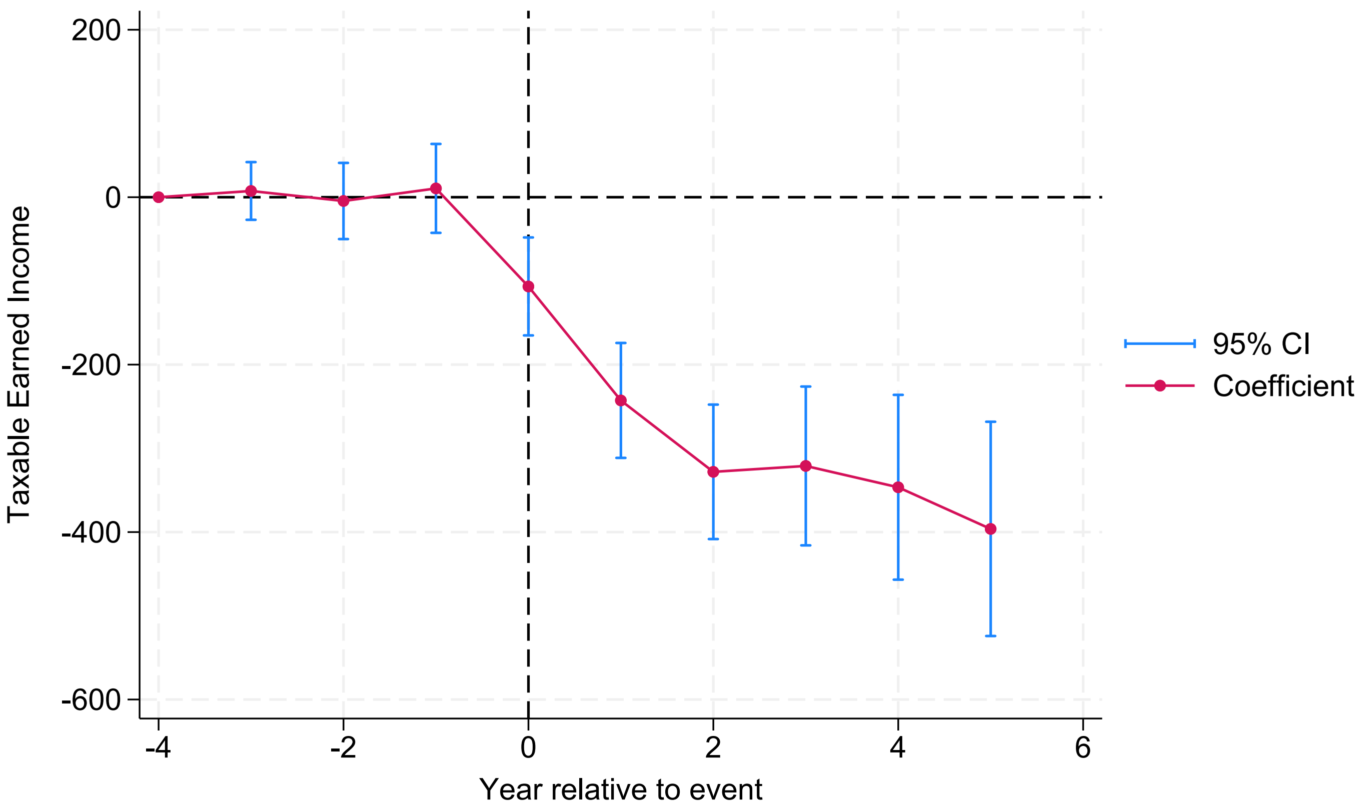

Figure 10.13: Event Study of Car Ownership around Parental Death

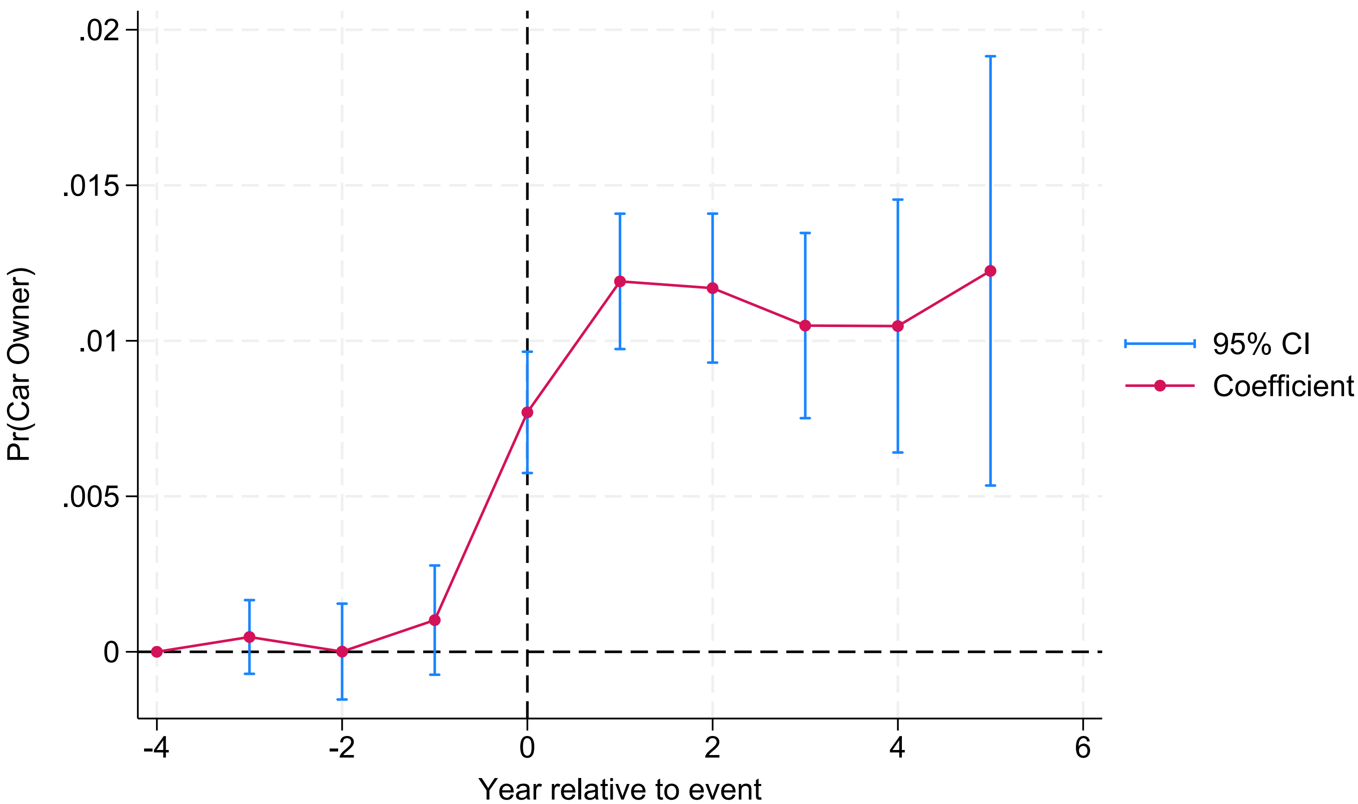

Figure 10.14: Triple-Difference Effects on Labor Supply

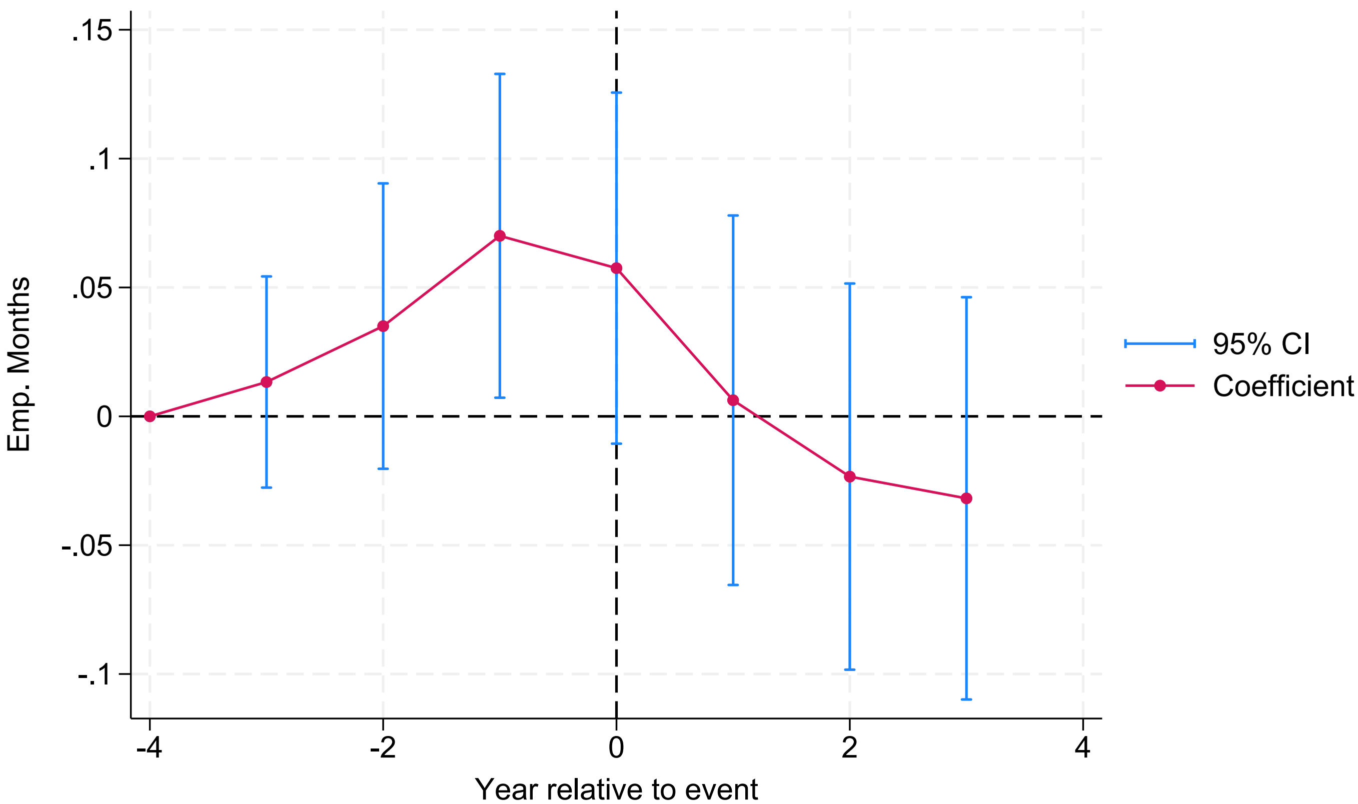

Figure 10.15: Triple-Difference Effects on Earnings

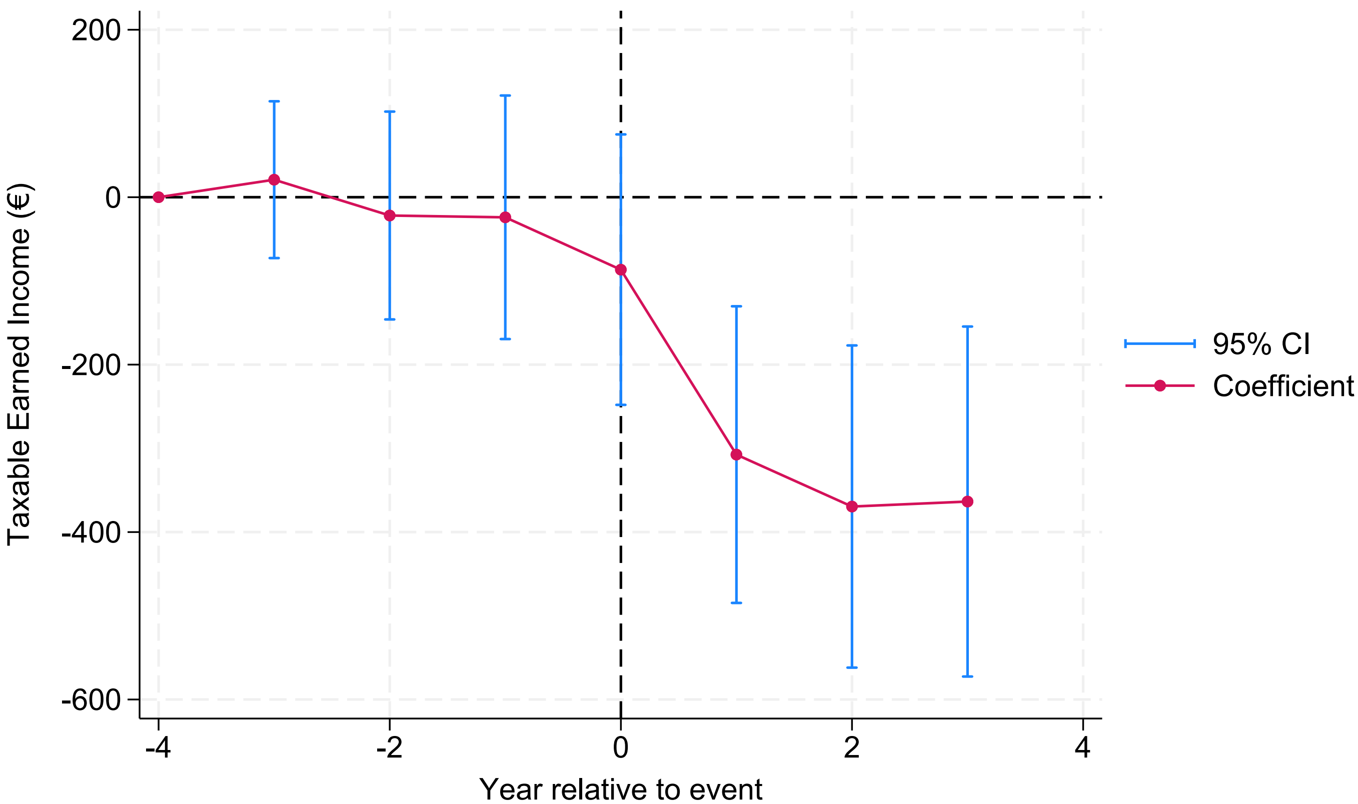

Figure 10.16: Potential Child Heirs and Tax Revisions by Inheritance Quintile

(a) Child Heirs

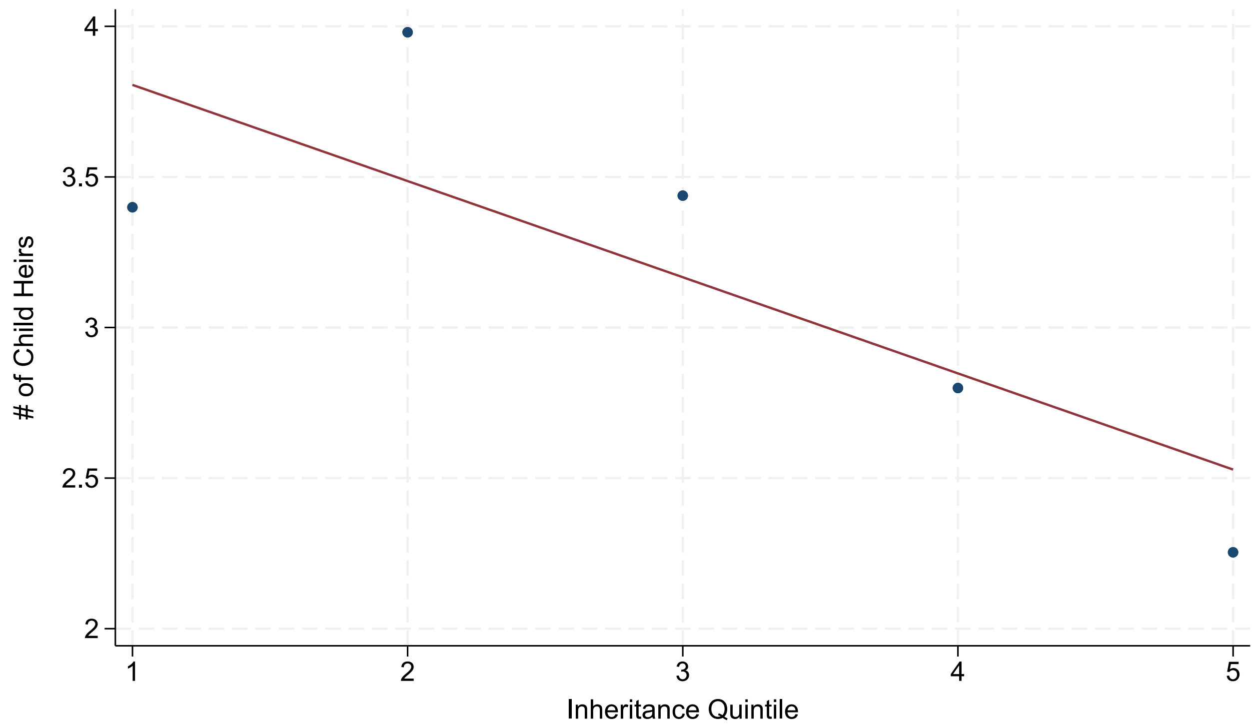

(b) Tax Revisions

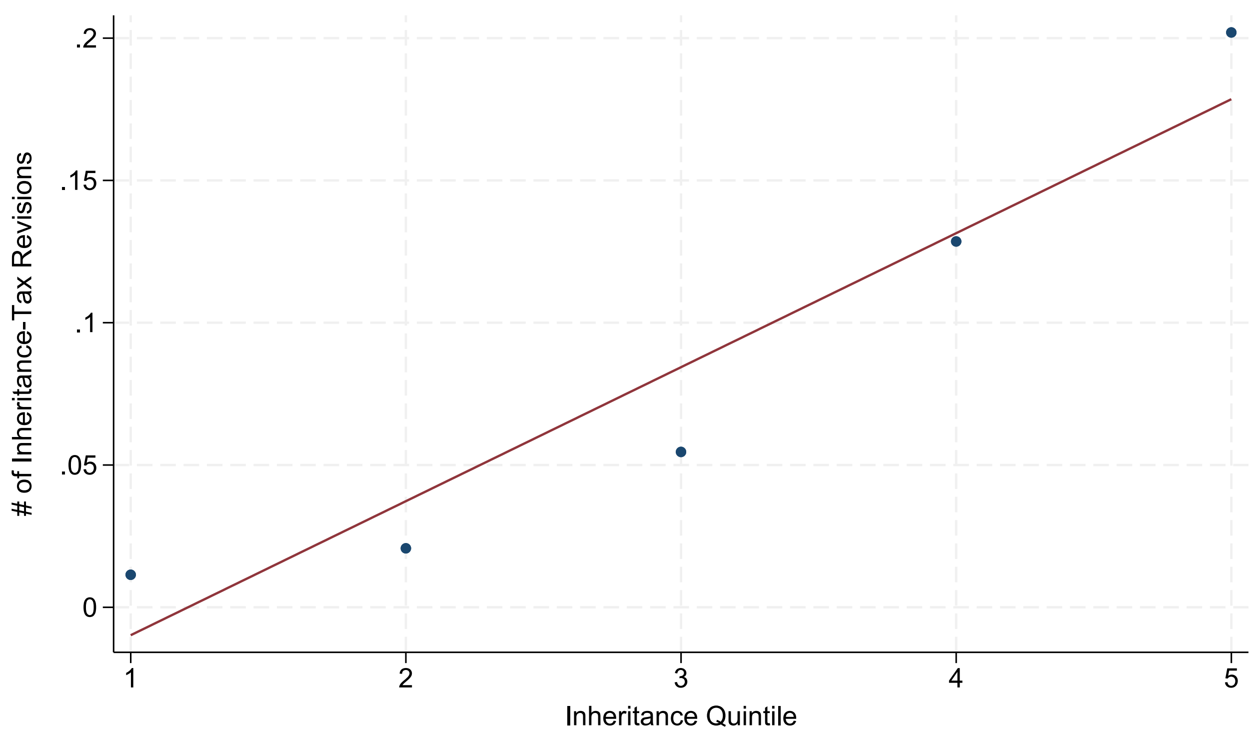

Figure 10.17: DDD Event Study of Speeding Rates around Parental Death: 5-year Event Horizon

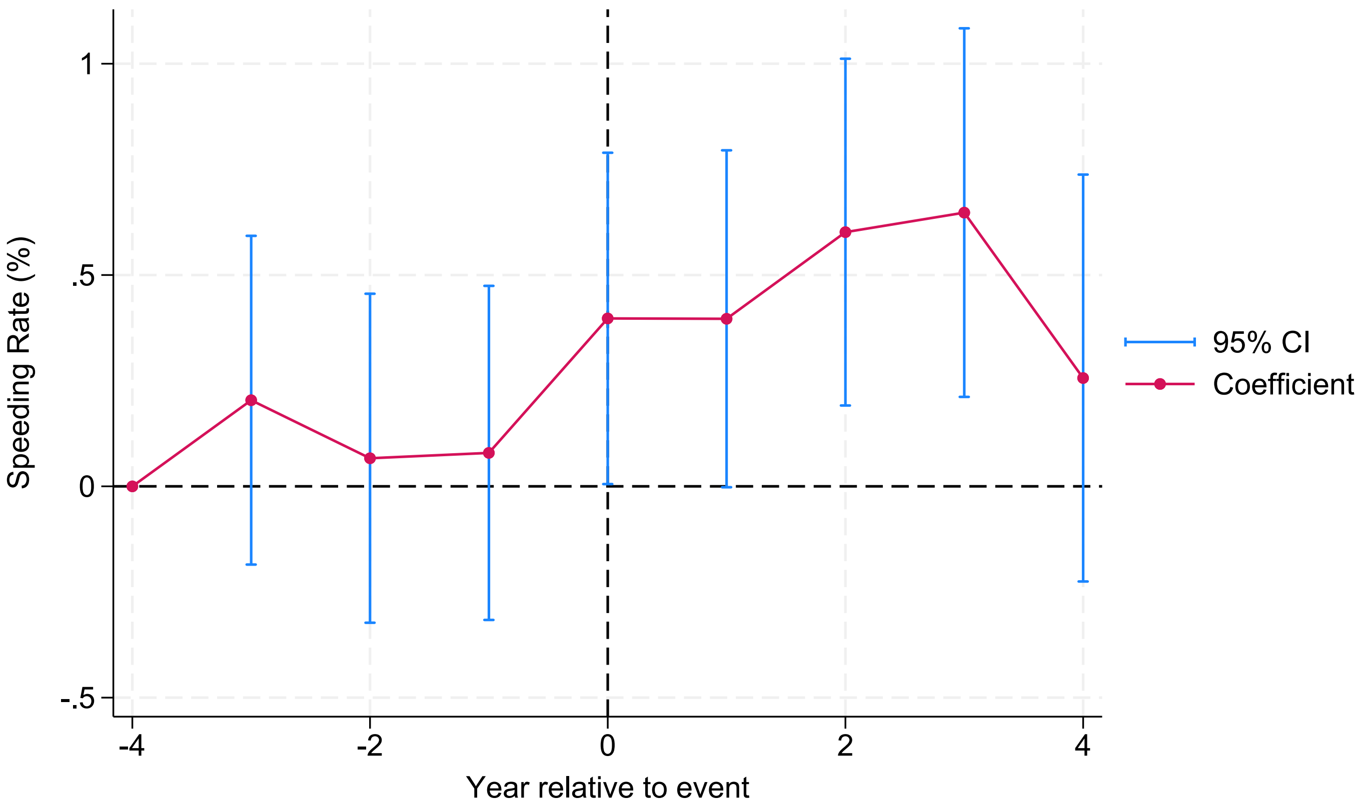

Figure 10.18: Relationship between Earnings and Speeding Rates (2012)

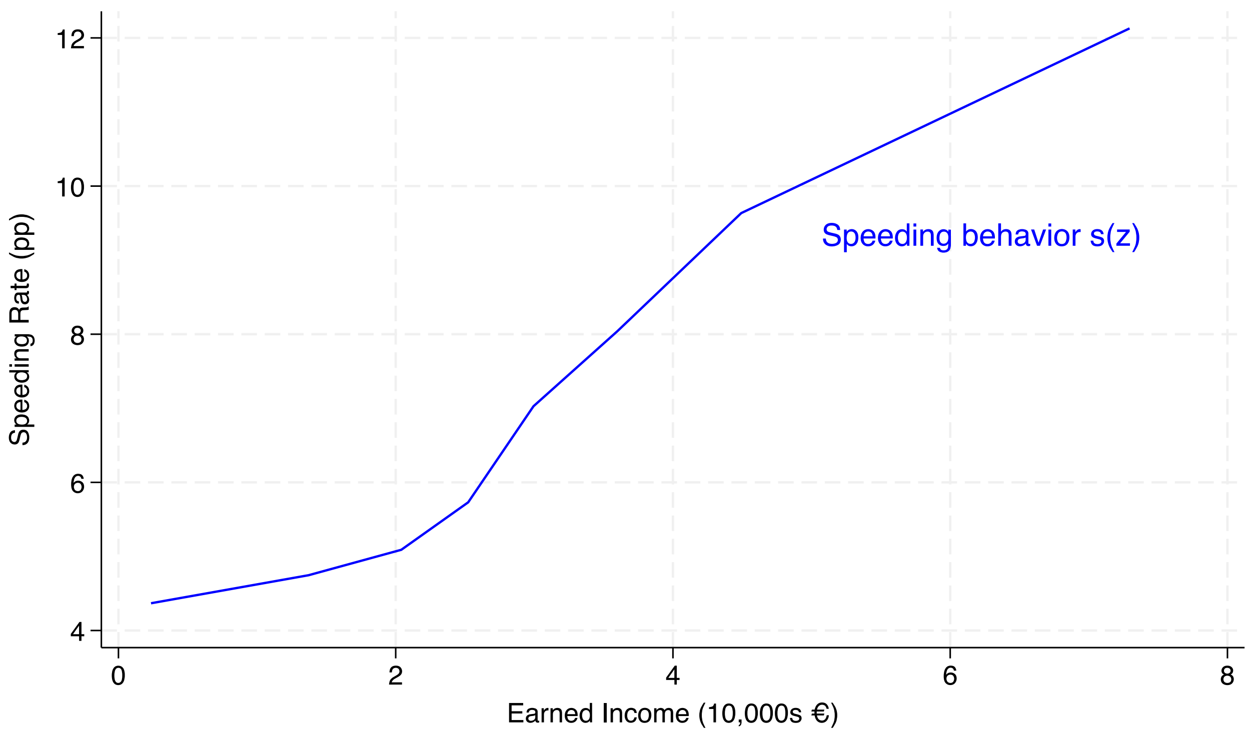

Figure 10.19: Heatmap of Optimal Redistributive Gradients

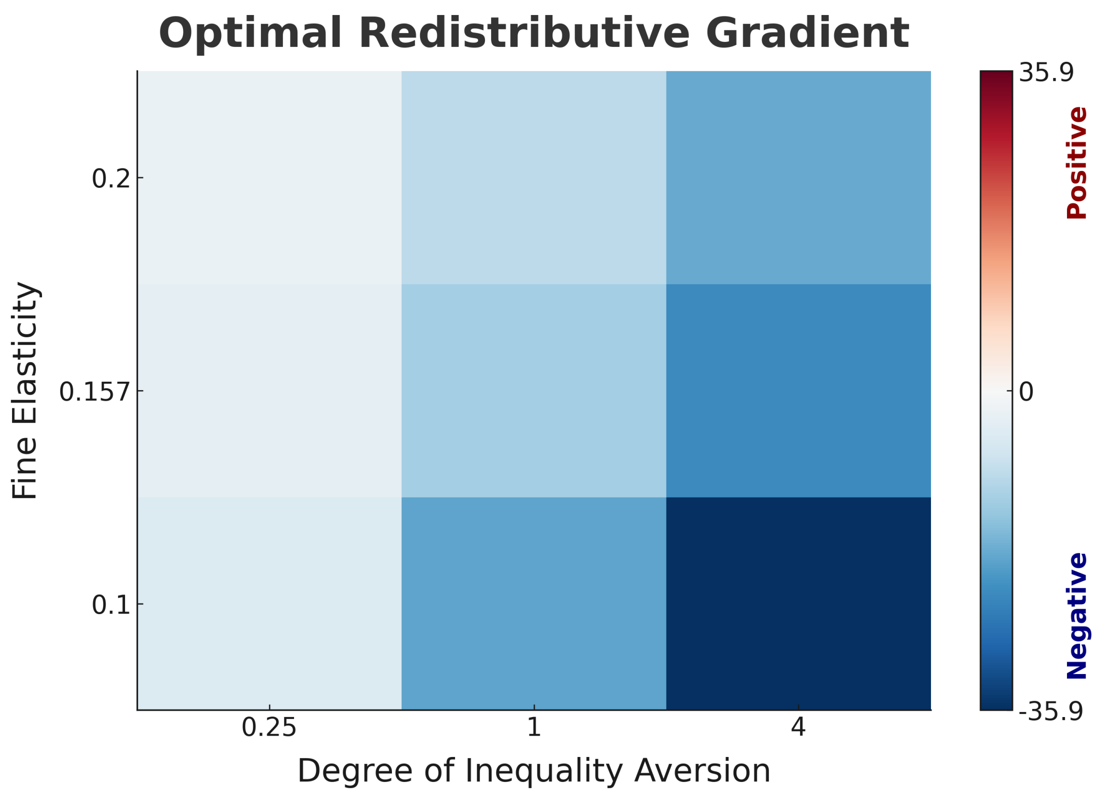

Figure 10.20: Speeding Rate Comprehension

The next questions concern hypothetical speeding fines that different income groups might have to pay. We describe how common speeding is in each income group by reporting how many speeding fines the group receives per 1,000 people per year.

For example, if an income group has 2,000 people and a total of 20 of them receive a speeding fine in a year, then proportionally this group receives 10 fines per 1,000 people per year.

If group A receives more speeding fines per 1,000 people than group B, then group A receives proportionally more fines.

The next question ensures that you understand what we mean by proportional (per-capita) numbers.

Suppose that a hypothetical group receives a total of 500 speeding fines per year and its size is 5,000 people. Estimate how many speeding fines the group receives per 1,000 people per year.

Figure 10.21: Fine-Policy Reasoning

We would appreciate if you could describe the basis on which you initially chose the Income-based fines policy. The tables below present the policies again.

Fixed fines		
Income Group	Speeding Fine	Fines per 1,000 people
Low (*Average Income*: *15 000 €*)	**330 €**	**40**
High (*Average Income*: *51 000 €*)	**330 €**	**100**
Total speeding tickets per year: 315 000		

Income-based fines		
Income Group	Speeding Fine	Fines per 1,000 people
Low (*Average Income*: *15 000 €*)	**150 €**	**70**
High (*Average Income*: *51 000 €*)	**510 €**	**70**
Total speeding tickets per year: 315 000		

Please use the text box below to describe the factors that influenced your decision to choose Income-based fines rather than Fixed fines.

Figure 10.22: Survey: Transition from Comparison to Slider

Let's imagine that the Finnish government could save on administrative costs by adopting the system that you did *not* choose in the previous question (Fixed fines).

Next, we ask you to assess how large the *minimum* annual savings would need to be for you to be willing to switch to the system you did not choose in the previous question.

Click the "Next" button to move forward.

Figure 10.23: Transfer Comparison Question

Imagine that the Finnish government receives a **one-time** unexpected revenue increase. The government plans to distribute some of this extra revenue to the Finnish adult population in the form of cash transfers. There are two alternative systems for distributing the funds:

1. **Fixed cash transfers -** every adult living in Finland receives the same amount.
2. **Income-based cash transfers -** the amount distributed to a person depends on whether they belong to a low- or high-income group.

We ask you to choose **which system would be better for Finland overall.** Suppose that your choice between the two ways of distributing the funds **wouldn't affect you: YOU would not receive a cash transfer in either case.**

The table below shows the two ways of distributing the funds described above. **Each income group covers half of the adult Finnish population. The low-income group includes the half of the population with the lowest incomes; the high-income group includes the half with the highest incomes.**

	Cash transfer to each individual in a group	
Income Group	**Fixed cash transfers**	**Income-based cash transfers**
Low (*Average Income*: 15 000 €)	**25 €**	**32 €**
High (*Average Income*: 51 000 €)	**25 €**	**20 €**

Which system do you prefer?

Please note that the systems differ in how much cash they distribute to different income groups. In addition, the systems do not necessarily distribute the same amount of transfers in total. The systems are otherwise completely similar.

Figure 10.24: Frequency of Fine Awareness Beliefs

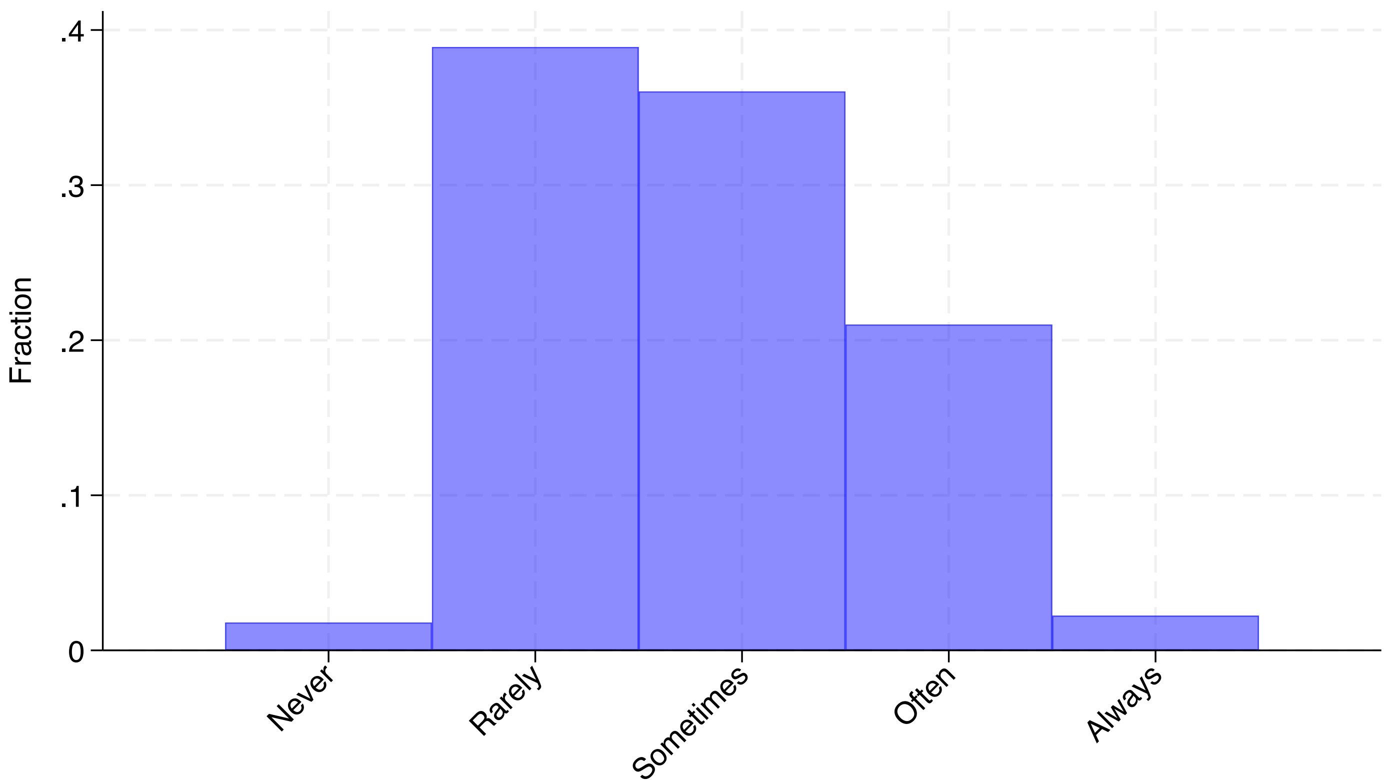

Figure 10.25: Frequency of Fine Responsiveness Beliefs