Optimal income-based traffic fines with labor supply responses

Max Ghenis

Monday 23rd February, 2026

1 Literature review

This paper sits at the intersection of several literatures: optimal income taxation, the economics of crime and deterrence, the speed-safety relationship, Pigouvian taxation with pre-existing distortions, the Finnish day-fine system, and equity in monetary sanctions. We review each in turn, emphasizing the gap our model fills.

1.1 Optimal taxation and behavioral responses

The modern theory of optimal taxation, initiated by Mirrlees [1971] and extended by Diamond and Mirrlees [1971], derives tax schedules that balance redistribution against efficiency costs arising from behavioral responses. Saez [2001] showed that optimal marginal tax rates can be expressed in terms of sufficient statistics—principally the elasticity of taxable income (ETI)—without specifying the full structure of preferences. Feldstein [1999] emphasized that the ETI captures all behavioral margins, including labor supply, tax avoidance, and evasion.

The empirical literature on the ETI, surveyed by Saez et al. [2012], places central estimates around 0.25 for broad income, with a range of 0.12 to 0.40. Chetty [2012] reconciles micro and macro labor supply estimates by accounting for optimization frictions, arriving at a Hicksian elasticity of approximately 0.25. Keane [2011] provides a comprehensive survey, noting that intensive-margin elasticities for prime-age men are typically 0.1–0.3, while estimates for women and secondary earners are larger.

Our contribution extends this literature by identifying a novel source of effective marginal taxation: income-linked regulatory penalties. Maag et al. [2012] document how benefit phase-outs create implicit marginal tax rates exceeding 80% for low-income families. We show that income-based fines create analogous distortions that vary with violation behavior, a channel not previously analyzed.

1.2 Crime, deterrence, and fine design

Becker [1968] established the framework for analyzing crime as a rational choice, where individuals weigh expected penalties against the benefits of offending. Polinsky and Shavell [1979] extended this to optimal fine design, showing that maximal fines with low detection probabilities minimize enforcement costs under risk neutrality. When offenders vary in wealth, Polinsky and Shavell [1991] showed that optimal fines should be wealth-dependent: flat fines over-deter the poor and under-deter the wealthy. Garoupa [1997] and Polinsky and Shavell [2007] provide comprehensive surveys.

Empirically, Chalfin and McCrary [2017] review the deterrence literature and find that both police presence and sanction severity reduce crime, with elasticities of crime to police around -0.3 to -0.5. Hansen [2015] exploits sharp punishment thresholds for drunk driving in Washington State, finding that crossing a blood alcohol content cutoff that triggers harsher sanctions reduces recidivism.

DeAngelo and Hansen [2014] show that reduced police enforcement during budget crises increases traffic fatalities in Oregon. These studies confirm that traffic penalties have real deterrent effects.

The behavioral dimension complicates standard deterrence theory. Gneezy and Rustichini [2000] demonstrated in a daycare experiment that introducing small fines for late pickup *increased* lateness, suggesting fines can crowd out intrinsic motivation when set too low. Chetty et al. [2009] showed that less salient taxes generate smaller behavioral responses, and Taubinsky and Rees-Jones [2018] provided experimental evidence on how attention variation affects welfare. These findings raise the question of whether the deterrent effect of income-based fines depends partly on their psychological salience. Makowsky and Stratmann [2009] demonstrate that enforcement intensity—and thus the effective fine burden—varies with local political economy, adding another dimension of heterogeneity beyond income.

Our model takes the deterrence motive seriously: income-based fines achieve more uniform expected disutility across the income distribution. The question is whether this benefit survives once labor supply distortions are accounted for.

1.3 Speed-safety relationship

The relationship between vehicle speed and crash severity is central to our calibration. Nilsson [2004] proposed the *power model*, in which fatal crashes are proportional to $(v_1/v_0)^n$ where v_1 and v_0 are the after and before speeds and $n \approx 4$ for fatalities, 3 for serious injuries, and 2 for all injury crashes. This nonlinear relationship implies that speeding carries sharply increasing marginal risk: a 10% speed increase roughly doubles fatality risk.

Elvik [2019] re-estimated the power model across a broader set of studies, finding somewhat lower exponents (around 3.5 for fatalities) with variation by road type. We adopt the Nilsson parameterization with $n \sim \mathcal{N}(4.0, 0.5)$ as our prior, encompassing both estimates.

The power model is attractive for our purposes because it provides a microfounded relationship between individual speeding intensity and mortality risk, avoiding the need to model crash mechanics directly. We express death probability as $p(s) = p_{\text{base}} \cdot (1 + s)^n$, where $s \in [0, 1]$ is fractional speed above the limit and p_{base} is the baseline annual traffic fatality probability.

1.4 Pigouvian taxation and labor market interaction

Our analysis closely parallels the literature on environmental taxes in the presence of pre-existing distortions. Building on Sandmo [1975], who first analyzed optimal corrective taxation with pre-existing distortions, Bovenberg and de Mooij [1994] showed that the optimal pollution tax falls *below* the Pigouvian level (marginal external damage) when it interacts with a distortionary labor tax. The intuition is that the environmental tax raises the cost of consumption, reducing the real wage and amplifying labor supply distortions—the *tax-interaction effect*. Bovenberg and Goulder [1996] showed that recycling environmental tax revenue through labor tax cuts (the *revenue-recycling effect*) partially but not fully offsets this interaction.

Jacobs and de Mooij [2015] proved a striking irrelevance result: when the income tax is set optimally, the optimal Pigouvian tax equals marginal external damage regardless of pre-existing distortions. The tax-interaction and revenue-recycling effects exactly cancel at the second-best optimum.

Allcott et al. [2019] extended this line of research to regressive sin taxes, analyzing the tension between corrective taxation and distributional concerns with heterogeneous agents and marginal social welfare weights—a tension directly relevant to income-based fines.

Our setting differs from the standard Pigouvian framework in a crucial respect. Environmental taxes apply uniformly to all consumers of a polluting good, whereas income-based fines apply only to offenders. This creates *heterogeneous* effective tax rates: agents who speed more face higher marginal taxation on labor income. The resulting distortion cannot be offset by adjusting the uniform income tax schedule, because the planner cannot condition taxes on speeding behavior (which is the rationale for fines in the first place). This heterogeneity is why the Jacobs–de Mooij irrelevance result does not carry over to our setting.

1.5 Finnish day-fines

Finland introduced the day-fine (*p"aiv"asakko*) system in 1921, making it the longest-running income-based penalty system. The formula sets the daily fine amount to roughly (monthly net income − 255)/60, multiplied by a number of day-fine units determined by offense severity Kantorowicz-Reznichenko and Faure [2021]. For speeding violations exceeding 20 km/h above the limit, fines become income-based; below this threshold, fixed petty fines (*rikesakko*) of around \euro200 apply.

Kaila [2024] provides the most rigorous empirical analysis, exploiting discontinuities in the Finnish system. Key findings include: (i) income-based fines reduce reoffending, though effects attenuate after six months; (ii) there is limited evidence of labor supply responses, potentially because day-fines are calculated from *previous-year* tax returns rather than current income; and (iii) there is no bunching in the income distribution near fine thresholds. The backward-looking income assessment is an important institutional feature: it severs the contemporaneous link between current work effort and fine liability that drives the labor distortion in our model. We discuss this distinction in Section 6.

Several other countries use related systems. Switzerland and several German states employ day-fine variants for criminal offenses, though not for traffic violations. Kantorowicz-Reznichenko and Faure [2021] provide a comparative analysis, finding broad support for the deterrence equity rationale but limited formal welfare analysis.

1.6 Equity and monetary sanctions

A growing literature documents the regressive impact of flat monetary sanctions. Harris [2016] shows that fines, fees, and court costs disproportionately burden low-income individuals, creating cycles of debt, license suspension, and incarceration. Lerman and Weaver [2014] demonstrate that contact with the criminal justice system—including monetary sanctions—erodes civic participation and trust in government, with downstream consequences for democratic engagement.

These findings motivate the equity case for income-based fines: if flat penalties trap the poor in punitive cycles while failing to deter the wealthy, calibrating fines to ability to pay seems both fair and efficient. Saez and Stantcheva [2016] provide a framework for incorporating such concerns through generalized social marginal welfare weights that can reflect a society's equity preferences.

However, the public economics literature counsels caution about using individual policy instruments for redistribution. Kaplow and Shavell [2002] argue on both theoretical and practical grounds that distributional concerns are better addressed through the tax-and-transfer system than through the design of specific regulatory policies. Our analysis provides a quantitative test of this principle: we compare the welfare effects of addressing fine regressivity through income-linking (which creates labor distortions) versus maintaining flat fines and relying on the existing fiscal system for redistribution.

1.7 Synthesis

No previous work has formally modeled the interaction between income-based fines and labor supply decisions in a setting with heterogeneous agents, endogenous speeding, and mean-field equilibrium. The closest contributions are the environmental tax literature—which analyzes Pigouvian taxes with labor interactions but not heterogeneous exposure—and the law and economics literature on wealth-dependent sanctions—which considers optimal deterrence but not labor supply responses. Our model fills this gap by combining elements from both traditions in a calibrated simulation framework.

2 Model

We develop a model of heterogeneous agents who jointly choose labor supply and speeding intensity under alternative fine structures. The environment features a proportional income tax, a fine system (flat or income-based), and a universal transfer funded by tax and fine revenue. We solve for mean-field equilibrium where agents optimize given aggregate outcomes and aggregate outcomes are consistent with individual choices.

2.1 Preferences and technology

Consider an economy with N agents indexed by productivity (wage) w_i, drawn from a lognormal distribution. Each agent chooses annual work hours $h_i \in [1, H]$ and speeding intensity $s_i \in [0, 1]$, where s represents the fractional speed above the posted limit and H is the maximum feasible work year.

Assumption 1. *Agents maximize:*

$$U_i = \underbrace{\log(1 + c_i)}_{\text{consumption}} + \underbrace{\alpha \log(1 + s_i)}_{\text{speeding benefit}} - \underbrace{\frac{\beta}{2}\left(\frac{h_i}{H}\right)^2}_{\text{labor cost}} - \underbrace{\frac{p(s_i) \cdot V}{1 + c_i}}_{\text{death cost}} \tag{1}$$

where c_i is consumption, $\alpha > 0$ is the weight on speeding utility, $\beta > 0$ governs labor disutility, V is the value of statistical life (VSL), and $p(s_i)$ is the annual death probability.

Several features merit discussion. Log consumption utility ensures declining marginal utility and positive consumption. The speeding term $\alpha \log(1 + s)$ captures time savings and private benefits from higher speeds, with diminishing returns. The quadratic labor cost in normalized hours h/H generates interior solutions and well-behaved comparative statics. The death cost term $p(s) \cdot V/(1 + c)$ converts the monetary VSL into utility units by multiplying by the marginal utility of consumption $u'(c) = 1/(1 + c)$, ensuring that risk valuation is consistent with the consumption utility function.

Assumption 2 (Power model). *The annual death probability follows the Nilsson [2004] power model:*

$$p(s) = p_{\text{base}} \cdot (1 + s)^n, \quad n \approx 4 \tag{2}$$

where p_{base} is the baseline death probability and n is the speed-fatality exponent Nilsson [2004].

This specification implies that fatality risk is highly convex in speeding intensity. A 10% increase in speed ($s = 0.1$) raises death probability by a factor of $(1.1)^4 \approx 1.46$, while a 50% increase ($s = 0.5$) raises it by $(1.5)^4 \approx 5.06$.

2.2 Budget constraint and fine structures

Agent i earns gross income $y_i = w_i h_i$, pays income tax, pays a fine that depends on the fine system, and receives a uniform transfer T. In the budget constraints below, τ denotes the agent-specific tax rate MTR_i. We use each agent's marginal tax rate from the CPS/PolicyEngine data as a linearization of the progressive tax schedule around the observed income level. This is a standard approximation in the public finance simulation literature Saez [2001]: the marginal rate governs the agent's labor supply response at the margin, which is the relevant object for the welfare comparison between fine systems. Because both fine systems use the same linearized tax treatment, any level bias from this approximation is symmetric and cancels in the welfare *difference* $\Delta W = W_{\mathrm{IB}} - W_{\mathrm{flat}}$, which is our primary object of interest. We discuss the limitations of this approximation in Section 6. We consider two fine structures.

Flat fine. Under a flat fine F, the penalty is $F \cdot s_i$—proportional to speeding intensity but independent of income:

$$c_i = w_i h_i (1 - \tau) - F s_i + T \tag{3}$$

Income-based fine. Under an income-based fine with rate ϕ, the penalty is $\phi \cdot y_i \cdot s_i$—proportional to both income and speeding:

$$c_i = w_i h_i (1 - \tau - \phi s_i) + T \tag{4}$$

The crucial difference is visible in (4): the income-based fine enters the budget constraint multiplicatively with income, creating an effective marginal tax rate of $\tau + \phi s_i$ that varies with speeding behavior.

2.3 First-order conditions

2.3.1 Flat fine system

Taking derivatives of (1) subject to (3):

$$\frac{\partial U_i}{\partial h_i} = 0: \quad \frac{w_i(1-\tau)}{1+c_i} + \frac{p(s_i) V w_i (1-\tau)}{(1+c_i)^2} = \frac{\beta h_i}{H^2} \tag{5}$$

$$\frac{\partial U_i}{\partial s_i} = 0: \quad \frac{\alpha}{1+s_i} = \frac{F}{1+c_i} + \frac{F \cdot p(s_i) V}{(1+c_i)^2} + \frac{p'(s_i) V}{1+c_i} \tag{6}$$

Under flat fines, the labor supply FOC (5) depends on speeding only through $p(s_i)$ and c_i, and the speeding FOC (6) depends on labor only through c_i. Labor and speeding decisions interact through consumption but the fine itself does not create a direct coupling.

2.3.2 Income-based fine system

Under the income-based system with budget constraint (4):

$$\frac{\partial U_i}{\partial h_i} = 0: \quad \frac{w_i(1-\tau-\phi s_i)}{1+c_i} + \frac{p(s_i) V w_i (1-\tau-\phi s_i)}{(1+c_i)^2} = \frac{\beta h_i}{H^2} \tag{7}$$

$$\frac{\partial U_i}{\partial s_i} = 0: \quad \frac{\alpha}{1+s_i} = \frac{\phi w_i h_i}{1+c_i} + \frac{\phi w_i h_i \cdot p(s_i) V}{(1+c_i)^2} + \frac{p'(s_i) V}{1+c_i} \tag{8}$$

The key difference is in (7): the marginal return to labor is reduced from $w_i(1-\tau)$ to $w_i(1-\tau-\phi s_i)$. Agents who speed more face a higher effective tax rate and supply less labor, all else equal. Similarly, (8) shows that the marginal cost of speeding now depends on income $w_i h_i$, creating the income-based deterrence.

Proposition 1. *Under income-based fines, the effective marginal tax rate on labor income is $\tau + \phi s_i$, which is increasing in both the fine rate ϕ and the agent's speeding intensity s_i.*

Proposition 2. *An increase in the fine rate ϕ reduces both speeding and labor supply:*

$$\frac{\partial s_i}{\partial \phi} < 0, \quad \frac{\partial h_i}{\partial \phi} < 0 \tag{9}$$

The first effect is the intended deterrence; the second is the unintended labor distortion.

Proof sketch. From (8), the marginal cost of speeding includes $\phi w_i h_i/(1+c_i)$, which is increasing in ϕ, so the equilibrium s_i falls. From (7), the marginal return to labor is $w_i(1-\tau-\phi s_i)/(1+c_i)$, which is decreasing in ϕ (holding s_i fixed). While the reduction in s_i partially offsets this by raising the net-of-fine return, the direct effect dominates when ϕ is small relative to $1-\tau$, which holds at empirically relevant fine rates. Numerical verification confirms this for all parameter draws in our Monte Carlo analysis.

Proposition 3. *The labor supply reduction from income-based fines is increasing in productivity w_i, provided that hours h_i are increasing in w_i.*

Proof sketch. The labor distortion arises from the additional effective tax ϕs_i on labor income $w_i h_i$. From (7), the wedge between the marginal return to labor and the marginal disutility is proportional to $\phi s_i w_i/(1+c_i)$. When w_i is higher, the absolute reduction in labor income from any given percentage reduction in hours is larger, and the deadweight loss—which is proportional to w_i^2 in a linear approximation—is therefore increasing in w_i. This holds when hours are interior and increasing in wages, which is the case in our simulations for all agents.

2.4 Mean-field equilibrium

Fine and tax revenue funds a uniform transfer. The government budget constraint is:

$$T = \frac{1}{N}\left[\sum_{i=1}^{N} \tau w_i h_i + \sum_{i=1}^{N} f_i(w_i h_i, s_i)\right] \tag{10}$$

where $f_i(\cdot)$ is the fine paid by agent i.

Definition 1 (Mean-field equilibrium). *A mean-field equilibrium is a collection of choices $\{(h_i^*, s_i^*)\}_{i=1}^N$ and a transfer T^* such that:*

1. *Each agent optimizes: $(h_i^*, s_i^*) = \arg\max_{h,s} U_i(h, s; T^*, w_i)$ for all i.*

2. *The government budget balances: T^* satisfies (10) given $\{(h_i^*, s_i^*)\}$.*

We solve for equilibrium using damped fixed-point iteration:

1. Initialize $T^{(0)} = 0$.
2. Given $T^{(k)}$, each agent solves their optimization problem via L-BFGS-B.
3. Compute the implied transfer $\hat{T}^{(k+1)}$ from (10).
4. Update: $T^{(k+1)} = \lambda \hat{T}^{(k+1)} + (1-\lambda)T^{(k)}$, with damping $\lambda \in (0, 1]$.
5. Iterate until $|T^{(k+1)} - T^{(k)}| / \max(|T^{(k)}|, 1) < \varepsilon$.

The damping parameter λ (set to 0.5 in our baseline) prevents oscillations that arise because agents' labor supply and speeding responses to the transfer create feedback loops.

2.5 Social welfare

We evaluate outcomes under three social welfare functions:

$$W_{\text{util}} = \sum_{i=1}^{N} U_i \quad \text{(utilitarian)} \tag{11}$$

$$W_{\text{rawls}} = \min_i U_i \quad \text{(Rawlsian)} \tag{12}$$

$$W_{\text{atk}}(\varepsilon) = N \left[\frac{1}{N} \sum_{i=1}^{N} U_i^{1-\varepsilon} \right]^{1/(1-\varepsilon)} \qquad \text{(Atkinson, } \varepsilon \geq 0\text{)} \tag{13}$$

The Atkinson family nests the utilitarian ($\varepsilon = 0$) and Rawlsian ($\varepsilon \to \infty$) criteria as special cases, allowing us to trace how social preferences over inequality affect the welfare ranking of fine systems.

2.6 Welfare decomposition

To understand the sources of welfare differences, we decompose $\Delta W = W_{\text{IB}} - W_{\text{flat}}$ into three channels:

$$\Delta W = \underbrace{\Delta W_{\text{deterrence}}}_{\text{safety gain}} + \underbrace{\Delta W_{\text{labor}}}_{\text{distortion loss}} + \underbrace{\Delta W_{\text{revenue}}}_{\text{fiscal effect}} \tag{14}$$

The deterrence gain captures welfare improvements from more effectively targeted penalties across the income distribution. The labor distortion loss captures the efficiency cost of the implicit tax on earnings. The revenue effect captures differences in equilibrium transfers arising from different revenue levels under the two systems.

3 Calibration

We calibrate the model to the United States, using real microdata from the Enhanced Current Population Survey (CPS) via PolicyEngine PolicyEngine [2024]. Rather than fixing parameters at point estimates, we specify informative priors for behavioral parameters and propagate uncertainty through forward Monte Carlo simulation. The key innovation is using empirically estimated per-agent marginal tax rates that capture the full complexity of the US tax-benefit system.

3.1 Income distribution and marginal tax rates

3.1.1 CPS microdata

Agent wages and tax rates are drawn from the Enhanced CPS microdata, processed through PolicyEngine's US microsimulation model PolicyEngine [2024], US Census Bureau [2024]. We restrict the sample to working-age adults (18–64) with positive employment income, yielding approximately 19,400 observations representing the US working population.

For each observation, PolicyEngine computes the marginal tax rate accounting for federal income tax (including bracket structure and standard deduction), state income taxes (varying across all 50 states plus DC), Federal Insurance Contributions Act (FICA) payroll taxes (Social Security 6.2% plus Medicare 1.45%, with the Social Security wage cap), Earned Income Tax Credit (EITC) phase-in and phase-out, and benefit phase-outs including the Supplemental Nutrition Assistance Program (SNAP), Medicaid, and housing assistance.

The resulting MTR distribution is dramatically heterogeneous. Workers in the EITC phase-out region face effective marginal rates near 40% Maag et al. [2012], while some middle-income workers above the EITC range but below higher tax brackets face rates around 22–25%. High earners face combined federal-state rates of 40–50%. This heterogeneity is a key empirical contribution: it means income-based fines interact very differently with the tax-benefit system at different income levels, amplifying the double distortion for workers already facing high marginal rates.

We clip marginal tax rates to the range $[0, 0.95]$ to handle occasional values exceeding 100% that arise from benefit cliffs (discrete jumps in program eligibility).

3.1.2 Sampling procedure

Each Monte Carlo draw samples $N = 50$ agents with replacement from the CPS microdata, using calibrated household weights as sampling probabilities. This weighted bootstrap preserves the representative structure of the CPS while allowing independent draws across simulations. Each sampled agent carries their empirical wage (employment income / 2,080 hours) and marginal tax rate.

Hourly wages are derived as $w_i = y_i/H$ where y_i is annual employment income and $H = 2{,}080$ hours.

3.2 Preference parameters

3.2.1 Speeding utility weight (α)

The parameter α governs the private benefit agents derive from speeding. We set $\alpha \sim \mathcal{N}(0.5, 0.15)$. Higher values of α generate more speeding in equilibrium and make deterrence more valuable, tilting the comparison toward income-based fines. In our baseline calibration, equilibrium mean speeding intensity is approximately 0.37, higher than typical observed behavior (most speeders exceed limits by 5–15 mph, corresponding to $s \approx 0.08$–0.23). This elevated speeding arises because our model does not include detection probability or non-monetary penalties (points, license suspension), which in practice constrain speeding. The continuous fine formulation $F \cdot s$ or $\phi \cdot y \cdot s$ can be interpreted as the expected fine conditional on the full enforcement technology Becker [1968], and the implied detection probability would then scale down the equilibrium speeding intensity. We discuss the implications of this calibration in Section 6.

3.2.2 Labor disutility (β)

The parameter β scales the labor disutility function $\beta(h/H)^2/2$. We set $\beta \sim \mathcal{N}(1.0, 0.3)$, which governs equilibrium hours and the level of labor disutility.

The quadratic specification implies a Frisch elasticity of labor supply that is identically 1.0, regardless of β. To see this, note that the Frisch elasticity is $\varepsilon^F = v'(h)/[h \cdot v''(h)]$. For $v(h) = \frac{\beta}{2}(h/H)^2$, we have $v'(h) = \beta h/H^2$ and $v''(h) = \beta/H^2$, so $\varepsilon^F = (\beta h/H^2)/(h \cdot \beta/H^2) = 1$. The parameter β affects equilibrium hours but not the curvature ratio that determines the elasticity.

This implied Frisch elasticity of 1.0 is higher than the meta-analytic consensus of approximately 0.25 from Chetty [2012] and Keane [2011]. Targeting $\varepsilon^F = 0.25$ would require a higher-power specification such as $v(h) \propto (h/H)^{1+1/\varepsilon} = (h/H)^5$. We retain the quadratic form for analytical tractability and because the higher elasticity makes our results *more conservative*: since the labor distortion from income-based fines is increasing in the elasticity, the welfare advantage of income-based fines would be even larger at the empirically estimated elasticity of 0.25 than at our model's value of 1.0. The finding that income-based fines dominate in 95% of draws despite an elasticity four times the consensus estimate strengthens the qualitative conclusion.

3.2.3 Maximum hours (H)

We fix $H = 2{,}080$ hours per year (40 hours/week $\times$ 52 weeks) with no uncertainty. This serves as a normalization; the effective labor supply margin operates through the interior choice of h.

3.3 Safety parameters

3.3.1 Value of statistical life (V)

We adopt $V \sim \mathcal{N}(11{,}600{,}000,\ 2{,}900{,}000)$ USD, following the US EPA's central estimate for regulatory impact analyses US Environmental Protection Agency [2024]. The 25% coefficient of variation reflects the range in meta-analytic estimates. The VSL enters the model as a scaling factor in the death cost term $p(s) \cdot V/(1 + c)$; higher values increase the private cost of speeding and reduce equilibrium speeding intensity under both fine systems.

3.3.2 Baseline death probability (p_{base})

The baseline annual traffic fatality probability is $p_{\text{base}} \sim \mathcal{N}(0.00012, 0.00006)$. This is derived from NHTSA FARS data National Highway Traffic Safety Administration [2024]: 1.37 deaths per 100 million vehicle miles traveled, combined with average annual driving of approximately 13,400 miles per licensed driver, yields an annual fatality probability of approximately 1.2×10^{-4}.

3.3.3 Speed-fatality exponent (n)

The power model exponent is $n \sim \mathcal{N}(4.0, 0.5)$, following Nilsson [2004], who estimated $n \approx 4$ for fatalities across international data. Elvik [2019] found somewhat lower values ($n \approx 3.5$) in updated estimates. Our prior encompasses both, with a range from approximately 3 to 5. This parameter is critical: higher exponents make speeding more dangerous and increase the value of deterrence, favoring income-based fines.

3.4 Fine system parameters

3.4.1 Flat fine

The flat fine baseline is $F = \$130$, approximating the US national average for speeding tickets based on data from the National Center for State Courts.

3.4.2 Income-based fine

The income-based fine rate is $\phi = 0.02$ per unit speeding intensity. There is no established US income-based fine system to calibrate against directly; we adopt a rate consistent with European day-fine systems Kantorowicz-Reznichenko and Faure [2021] and the San Francisco pilot San Francisco Municipal Transportation Agency [2025]. At median US income (~\$56,000), a speeding intensity of $s = 0.1$ yields a fine of $0.02 \times 56{,}000 \times 0.1 = \$112$—comparable to the flat fine at median income, providing a revenue-neutral comparison while creating a measurable implicit tax on labor income.

3.5 Monte Carlo procedure

For each Monte Carlo draw, the procedure is as follows. Sample N agents with replacement from CPS microdata using calibrated household weights as sampling probabilities, obtaining per-agent wages w_i and marginal tax rates MTR_i. Draw $(\alpha, \beta, V, p_{\text{base}}, n)$ independently from their Normal priors, clipping to valid ranges ($\alpha, \beta > 0.01$; $p_{\text{base}} \in [10^{-8}, 0.1]$; $n \in [0.5, 10]$). Find the welfare-maximizing flat fine F^* and income-based rate ϕ^* by grid search. Solve mean-field equilibrium under each optimal fine system. Record utilitarian welfare, mean speeding, Gini coefficient, and equilibrium transfers for each system. Compute the welfare difference $\Delta W = W_{\text{IB}} - W_{\text{flat}}$ and its decomposition.

The baseline specification uses 100 Monte Carlo draws with 50 agents per draw. The moderate sample sizes reflect the computational cost of solving mean-field equilibrium with per-agent L-BFGS-B optimization inside a damped fixed-point iteration loop; each draw requires solving two equilibria (flat and income-based), each involving repeated optimization of all agents until convergence. The standard error of the estimated probability $\Pr(\Delta W > 0)$ at $\hat{p} = 0.95$ is approximately $\sqrt{0.95 \times 0.05/100} \approx 0.022$, adequate to establish the direction of the welfare comparison.

Tax rates are **not drawn from priors**—they are empirical, fixed per agent, drawn from the CPS microdata. This is a key methodological improvement over using a single scalar tax rate parameter.

The following table summarizes the prior specifications.

Parameter	Symbol	Mean	SD	Source
Speeding utility weight	α	0.50	0.15	Calibrated to US speeding rates
Labor disutility	β	1.00	0.30	Keane [2011], Chetty [2012]
Maximum hours	H	2,080	0	Standard full-time year
Value of statistical life	V	11,600,000	2,900,000	US Environmental Protection Agency [2024]
Baseline death probability	p_{base}	0.00012	0.00006	National Highway Traffic Safety Administration [2024]
Speed-fatality exponent	n	4.0	0.5	Nilsson [2004], Elvik [2019]
Frisch elasticity (implied by quadratic v)	ε^F	1.0	—	Implied by $v(h) = \frac{\beta}{2}(h/H)^2$; see text

4 Results

This section presents the main findings from the Monte Carlo analysis. We report welfare comparisons between flat and income-based fines, distributional outcomes, welfare decompositions, and sensitivity analysis. Results are computed over 100 parameter draws from the priors specified in Section 4, with each draw sampling 50 agents from US CPS microdata with empirically estimated marginal tax rates. The moderate sample sizes reflect the computational cost of solving mean-field equilibrium with per-agent optimization; we verify robustness to sample size in the convergence diagnostics.

4.1 Baseline welfare comparison

The central question is whether flat or income-based fines generate higher social welfare. Under utilitarian welfare (sum of individual utilities), the Monte Carlo analysis yields a clear result.

Finding 1. *Under the baseline calibration using US CPS data, income-based fines generate higher utilitarian welfare than flat fines in 95% of Monte Carlo draws.*

The mean welfare difference is $\Delta W = W_{\text{IB}} - W_{\text{flat}} = 0.83$ (95% CI: $[-0.02, 3.22]$), indicating that income-based fines dominate. This result reflects the fact that the distributional gain from income-proportional penalties—lower fines for low-income agents and the resulting welfare improvement through concave utility—exceeds the labor distortion cost from the implicit tax on earnings. While the double distortion mechanism is present, its magnitude is quantitatively small: the additional effective marginal tax rate is 0.2–2 percentage points at the baseline fine rate, modest relative to the pre-existing MTR heterogeneity (which ranges from near-zero to over 50%).

The result is robust across parameter draws but not unanimous: flat fines dominate in approximately 5% of draws, typically when the labor supply elasticity is drawn from the upper tail of its prior distribution and the speeding utility weight is low.

4.2 Distributional analysis

The welfare advantage of income-based fines is reinforced by distributional improvements.

Finding 2. *Income-based fines reduce consumption inequality. The mean Gini coefficient under income-based fines is 0.325, compared with 0.343 under flat fines.*

This finding reflects two channels. First, income-based fines redistribute from high-income speeders to all agents via the universal transfer: high-income agents pay more in fines, increasing total revenue and the equilibrium transfer. Second, flat fines are regressive in the sense that they represent a larger share of income for low-wage agents, exacerbating pre-existing inequality.

Finding 3. *The welfare ranking is even stronger under inequality-averse social welfare functions. Under Rawlsian preferences, income-based fines dominate flat fines in virtually all draws.*

The Rawlsian criterion places all weight on the worst-off agent—typically the lowest-wage individual. Since income-based fines already dominate under utilitarian preferences, the Atkinson crossover from flat to income-based dominance occurs at $\varepsilon = 0$: no inequality aversion is needed to prefer income-based fines. This contrasts with the theoretical prediction that the double distortion would create a meaningful equity-efficiency trade-off; with realistic US calibration, the efficiency cost is too small to outweigh the distributional gains.

4.3 Welfare decomposition

We decompose the welfare difference into three components following (14).

Finding 4. *The deterrence gain from income-based fines exceeds the labor distortion loss, explaining the overall welfare advantage.*

The decomposition reveals that income-based fines achieve more uniform deterrence across the income distribution: high-income agents face larger penalties proportional to their ability to pay, reducing the under-deterrence problem inherent in flat fines. The labor distortion channel is present but small relative to the 0–50 percentage point range of pre-existing MTRs: the implicit tax ϕs adds only 0.2–2 percentage points to effective marginal rates at the baseline fine rate $\phi = 0.02$, generating small deadweight loss relative to the deterrence gain. The revenue effect—arising from differences in equilibrium transfers between the two systems—is the smallest component.

4.4 Aggregate speeding

An important finding is that aggregate speeding is slightly *higher* under income-based fines (mean $s = 0.370$) than under flat fines (mean $s = 0.358$). This may appear to contradict the "deterrence equity" narrative, but it reflects the composition of optimal fine levels. The welfare-maximizing income-based rate ($\phi^* \approx 0.076$) generates lower per-unit deterrence for low-income agents than the optimal flat fine (\$3,262), because $\phi^* \times y$ falls below \$3,262 for agents with income below approximately \$43,000—roughly the bottom half of the income distribution. The welfare gain from income-based fines therefore comes not from reducing aggregate speeding but from the distributional improvement: low-income agents pay less and high-income agents pay more, reducing consumption inequality and improving utilitarian welfare through the concavity of log utility. This is consistent with the welfare decomposition (Finding 4): the deterrence gain arises from redistributing fine burdens across the income distribution, not from reducing aggregate speeding.

4.5 Optimal fine levels

The welfare-maximizing fine levels provide additional insight.

Finding 5. *The optimal flat fine averages \$3,262 (95% CI: [\$1,238, \$5,000]), far exceeding the current US national average of approximately \$130. The optimal income-based fine rate averages* $\phi^* = 0.076$ *(95% CI: [0.025, 0.10]).*

Both confidence intervals exhibit grid boundary effects: the upper end of the flat fine CI coincides with the grid maximum of \$5,000, and the upper end of the income-based CI coincides with $\phi = 0.10$. Some draws may favor values above these bounds. Extending either grid would likely shift the respective optimal levels slightly upward without affecting the welfare ranking, since the welfare function is relatively flat near the optimum.

The high optimal flat fine reflects the model's emphasis on mortality risk: with $p_{\text{base}} = 0.00012$ and $n = 4$, even moderate speeding carries substantial fatality risk, and the model favors strong deterrence. However, these optimal levels should be interpreted through the lens of detection probability, which the model abstracts from: at a plausible detection rate of $\pi = 0.05$, the expected fine per speeding event would be $0.05 \times 3{,}262 \approx \$163$, close to current fine levels. The optimal income-based rate of $\phi^* = 0.076$ generates fines of approximately $\phi^* \times y \times s = 0.076 \times 56{,}000 \times 0.37 \approx \$1{,}575$ for a median-income agent with typical speeding, distributed proportionally to income.

4.6 Effective marginal tax rates

The implicit tax created by income-based fines varies across agents and represents the core mechanism of the double distortion. Because we use empirical per-agent MTRs from the CPS, the interaction between fines and existing taxes is heterogeneous across the income distribution.

Finding 6. *Under income-based fines, effective marginal tax rates for regular speeders exceed their CPS-based marginal tax rates by 0.2–2 percentage points, with larger increases for agents who speed more intensively.*

For an agent with speeding intensity s facing marginal tax rate MTR_i and fine rate ϕ, the effective marginal tax rate on labor income is:

$$\text{EMTR}_i = \text{MTR}_i + \phi s \tag{15}$$

At the baseline fine rate $\phi = 0.02$, an agent with moderate speeding ($s = 0.1$) faces an additional 0.2 percentage point effective tax; with high speeding ($s = 0.5$) the additional tax rises to 1 percentage point. At the welfare-maximizing rate $\phi^* \approx 0.076$, these figures rise to 0.76 and 3.8 percentage points respectively. While the welfare cost of these additional tax wedges is amplified by the pre-existing marginal tax rate—because deadweight loss is convex in the total tax rate Harberger [1964]—the magnitude is small enough that the distributional gain dominates.

Workers in the EITC phase-out region (earning roughly \$20,000–\$50,000) face the highest baseline MTRs (near 40%), so the additional distortion from income-based fines is disproportionate there. However, these workers also benefit most from the income-scaling of fines, which reduces their fine burden relative to a flat system, partially offsetting the labor distortion through higher after-fine consumption.

4.7 Sensitivity to key parameters

4.7.1 Labor supply elasticity

The welfare ranking is most sensitive to the labor supply elasticity.

Finding 7. *The probability that flat fines dominate increases with the labor supply elasticity, but income-based fines continue to dominate in the large majority of draws even at the upper end of empirically plausible elasticities.*

This finding has a simple intuition: the labor distortion channel matters more when labor supply is elastic. The quadratic disutility specification implies a Frisch elasticity of 1.0, four times the meta-analytic consensus of 0.25 Chetty [2012]. Because income-based fines dominate even at this elevated elasticity, the result would hold a fortiori at the empirically estimated value. Variation in β across Monte Carlo draws modulates equilibrium hours and the absolute magnitude of labor distortion, but the curvature ratio that determines the elasticity remains fixed at 1.0.

4.7.2 Value of statistical life

Finding 8. *Higher VSL values favor income-based fines by increasing the value of deterrence.*

The VSL enters the death cost term $p(s) \cdot V/(1 + c)$, scaling the private cost of speeding risk. When V is large, agents internalize more mortality risk, making effective deterrence across the income distribution more valuable. The interaction with fine structure favors income-based fines because they achieve deterrence more equitably.

4.7.3 Speed-fatality exponent

Finding 9. *Higher power model exponents (n) favor income-based fines. When speeding is more dangerous, the deterrence benefit of income-proportional penalties increases.*

The exponent n controls the convexity of the death probability function $p(s) = p_{\text{base}}(1 + s)^n$. Higher n means that even moderate speeding carries substantial mortality risk, making effective deterrence across the full income distribution more valuable.

4.8 Convergence diagnostics

The mean-field equilibrium is solved by damped fixed-point iteration with damping parameter $\lambda = 0.5$.

Finding 10. *The equilibrium solver converges within 200 iterations for more than 99% of Monte Carlo draws, with typical convergence in 20–50 iterations.*

Convergence is measured by the relative change in the universal transfer: $|T^{(k+1)} - T^{(k)}| / \max(|T^{(k)}|, 1) < 10^{-4}$. The damping parameter prevents oscillations that can arise when large changes in the transfer induce large changes in labor supply and speeding, which in turn change revenue and the implied transfer. Draws that fail to converge within 200 iterations are flagged and excluded from welfare comparisons.

5 Discussion

Our analysis reveals that, despite the theoretical double distortion mechanism, income-based fines generate higher welfare than flat fines under realistic US calibration. The deterrence equity gain from income-proportional penalties exceeds the labor distortion cost in the large majority of parameter configurations. This section discusses the sensitivity of our results to modeling choices, the role of institutional context, and implications for policy design.

5.1 Backward-looking income assessment

The most important institutional qualification concerns the timing of income measurement. Finland's day-fine system calculates penalties from *previous-year* tax returns, not current income Kaila [2024]. US systems vary: San Francisco's pilot San Francisco Municipal Transportation Agency [2025] may use more contemporaneous income data, as would be feasible with real-time income verification through IRS data sharing. The Staten Island day-fine experiment Hillsman [1990] used recent pay stubs to assess income at the time of sentencing.

When fines are backward-looking, the contemporaneous link between current labor effort and fine liability is severed, and the effective marginal tax rate on current labor is MTR_i regardless of speeding behavior. This design feature substantially weakens the double distortion mechanism. However, it does not eliminate it entirely. Forward-looking agents who anticipate future fines will internalize that higher *future* income implies higher *future* fines, generating a dynamic version of the labor supply distortion. The strength of this channel depends on the persistence of income shocks, the frequency of violations, and discount rates. Kaila [2024] finds limited evidence of labor supply responses in the Finnish system, consistent with the attenuation we would expect from backward-looking assessment.

A contemporaneous system—where fines are linked to current-year income—would exhibit the full force of the double distortion. As US jurisdictions design income-based fine programs, the choice between backward-looking and contemporaneous assessment is a critical design parameter that directly affects the magnitude of the labor distortion we identify.

5.2 Sensitivity to the labor supply elasticity

Our results are most sensitive to the labor supply elasticity, which governs the responsiveness of work effort to the effective marginal tax rate. The quadratic labor disutility specification in our model implies a Frisch elasticity of 1.0 (see Section 4), substantially higher than the meta-analytic consensus of approximately 0.25 from Chetty [2012] and Saez et al. [2012]. Despite operating at this elevated elasticity—which amplifies the labor distortion channel by a factor of roughly four relative to the empirical estimate—income-based fines still dominate in 95% of draws. At the consensus elasticity of 0.25, the labor distortion would be correspondingly smaller, and the welfare advantage of income-based fines would be even more pronounced.

At the intensive margin, elasticities for prime-age men are typically 0.1–0.3 Keane [2011]. At the extensive margin (participation), elasticities can be substantially larger, particularly for secondary earners and low-income workers. If income-based fines push some agents below their participation threshold—inducing them to exit the labor force entirely—the efficiency cost could be larger than our continuous model suggests, potentially narrowing or reversing the welfare advantage of income-based fines.

The elasticity also varies across institutional settings and income levels within the US. Workers in the EITC phase-out range may exhibit different behavioral responses than high-income earners, because the EITC phase-out already creates strong implicit taxes that interact with the fine-induced distortion. The US labor market, with weaker employment protections and more flexible hours than many European systems, may also permit larger intensive-margin responses, which would strengthen the double distortion and narrow the welfare gap.

5.3 The EITC phase-out and the double distortion

A distinctive feature of the US calibration is the interaction between income-based fines and the Earned Income Tax Credit phase-out. Workers earning roughly $20,000--50$,000 face effective marginal tax rates near 40% due to the combined effect of federal income tax, FICA payroll taxes, and the EITC phase-out Maag et al. [2012]. Adding an income-based fine on top of these already-high rates creates particularly large deadweight loss because of the quadratic relationship between tax rates and efficiency costs Harberger [1964].

This interaction is especially concerning from an equity perspective. The EITC phase-out region contains many working-poor families—precisely the population that income-based fines are intended to help. While income-based fines reduce the *level* of the fine for these workers (relative to a flat fine), they increase the *marginal tax rate* on labor, potentially discouraging additional work effort. The net welfare effect depends on the relative magnitudes of the fine reduction and the labor distortion, which our Monte Carlo analysis quantifies.

5.4 Inequality and social preferences

The welfare ranking of fine systems depends on how society values equality, but with realistic US calibration the ranking favors income-based fines even under utilitarian preferences. Income-based fines dominate in 95% of Monte Carlo draws without any inequality aversion, and the advantage strengthens under more egalitarian social welfare functions.

Saez and Stantcheva [2016] provide a framework for incorporating diverse social preferences through generalized social marginal welfare weights. In their framework, the optimal degree of income-basedness depends on the weight society places on equity versus efficiency—a normative choice that our model can inform but not resolve. Our contribution is to show that, at empirically calibrated parameters, income-based fines are welfare-improving even before accounting for inequality aversion. The efficiency cost of the implicit labor tax exists but is quantitatively small relative to the deterrence equity gain.

This finding should not be interpreted as meaning the double distortion is irrelevant. At higher labor supply elasticities, higher fine rates, or with contemporaneous income assessment, the labor distortion channel would be larger. The relevant policy question is whether the efficiency cost is large enough to overcome the equity and deterrence benefits—and our calibration suggests it is not, at least under the baseline parameterization.

5.5 The value of targeted deterrence

Income-based fines are motivated by the observation that flat fines under-deter the wealthy and over-deter the poor. Our model captures this asymmetry: under flat fines, high-income agents speed more because the fine represents a smaller fraction of their consumption. Income-based fines equalize the deterrence margin across the income distribution, and our results confirm that this equalization generates a welfare gain of $\Delta W = 0.83$ (positive in 95% of draws).

The welfare value of this equalization depends on the curvature of the death probability function. Under the power model with $n = 4$, speeding carries sharply increasing risk, and under-deterrence of high-income agents imposes significant costs (through higher aggregate speeding and death probability). Our model focuses on private mortality risk; incorporating external harm to other road users would further strengthen the case for income-based fines by increasing the social value of deterrence.

5.6 Alternative policy instruments

Although our results favor income-based fines, alternative policy instruments may achieve similar objectives with different trade-offs.

One approach is payment flexibility: flat fines combined with income-contingent payment plans can address the regressivity concern without creating labor supply distortions. The fine amount remains flat, preserving the labor supply incentives, while the payment schedule accommodates liquidity constraints. Several US jurisdictions have adopted such systems, and the San Francisco pilot includes payment plan provisions alongside its income-scaling mechanism San Francisco Municipal Transportation Agency [2025].

Many US courts already offer community service as an alternative to monetary fines for defendants who cannot pay. This avoids the regressive burden of flat fines while not conditioning on income in a way that distorts labor supply. However, community service imposes time costs that may be more burdensome for low-income workers with less flexible schedules.

Non-monetary sanctions—point-based systems, license suspensions, and mandatory traffic safety courses—create penalties that are less directly tied to income. Bourgeon and Picard [2007] analyze point-record systems and show they can achieve effective deterrence through non-monetary channels. However, non-monetary sanctions may have their own distributional consequences if time costs or license dependency vary with income.

Finally, rather than linking fines to income, flat fines can be paired with enhanced transfers to low-income households. Kaplow and Shavell [2002] advocate separating the pricing function (deterrence through fines) from the redistributive function (transfers through the tax system). This approach avoids the implicit tax on labor while achieving distributional goals through a more efficient instrument. In the US context, expanding the EITC or other targeted transfers could offset the regressive impact of flat fines without distorting the fine-labor supply link.

5.7 Limitations

Our analysis rests on several simplifying assumptions that merit acknowledgment.

The model analyzes a single-period decision, abstracting from reputation effects, learning, and habit formation. Speeding behavior is likely persistent, and fines may have dynamic deterrent effects that our static model misses. Within each Monte Carlo draw, agents share common preference parameters (α, β), differing only in wages and marginal tax rates. Heterogeneity in risk attitudes, time preferences, or driving needs could affect optimal fine design; for instance, agents who must drive for work face different trade-offs than recreational drivers.

The model bases income-based fines on gross employment income ($y_i = w_i h_i$), whereas real-world systems vary in their income concept. Finland's day-fine system uses monthly net income minus a fixed deduction Kaila [2024]; San Francisco's pilot may use adjusted gross income or another tax-return-based measure. Using gross income in the model amplifies the double distortion because the fine compounds with existing taxes—the effective rate is $\mathrm{MTR}_i + \phi s$ on gross income, rather than ϕs applied to after-tax income. If a jurisdiction based fines on net (after-tax) income, the fine would not create an additional labor supply distortion beyond what the income tax already imposes, since earning more after-tax income would already reflect the agent's marginal rate. The choice of income concept is therefore a consequential design parameter that affects both the magnitude of the labor distortion and administrative feasibility.

We assume uniform detection probability across income levels. If wealthy individuals can better avoid detection—through legal representation, choice of routes, or vehicle technology—the effective deterrence of income-based fines may differ from what our model predicts.

The budget constraint uses each agent's marginal tax rate as a proportional rate, which is a linearization of the progressive tax schedule around the observed income level. This approximation overstates the average tax burden for agents in higher brackets (where the marginal rate exceeds the average rate) and understates it for agents on the EITC phase-in (where the marginal rate is negative). For the welfare *comparison* between fine systems, this bias is approximately symmetric: both fine regimes use the same linearized tax treatment, so the bias cancels to first order in the welfare difference $\Delta W = W_{\mathrm{IB}} - W_{\mathrm{flat}}$, which is our primary object of interest. The cancellation is not exact, because the behavioral responses to the two fine systems differ, and these responses interact differently with the linearization error. At empirically small fine rates ($\phi \approx 0.02$–0.05), the behavioral responses are small enough that this second-order interaction is negligible. The marginal rate is the correct object for analyzing agents' behavioral responses at the margin, which drive the labor distortion channel.

While our use of PolicyEngine-computed marginal tax rates from CPS microdata is a substantial improvement over a single scalar tax rate, the CPS itself has known limitations. High incomes are top-coded at varying thresholds (roughly \$200,000–\$300,000 depending on year and state), compressing the upper tail of the income distribution where income-based fines have the largest bite. Some benefit variables (Supplemental Nutrition Assistance Program, Medicaid, housing assistance) are statistically imputed, and these imputations drive the benefit phase-out MTR spikes. The MTR estimates also assume current-law tax policy and do not capture informal economy participation or tax noncompliance.

The model treats labor supply as a continuous intensive-margin choice, with hours bounded below at $h \geq 1$. If income-based fines push some agents below their participation threshold—inducing them to exit the labor force entirely—the efficiency costs would be larger than we estimate. This is particularly relevant for low-income workers on the EITC phase-out, who face the highest baseline marginal rates and for whom the additional fine-induced tax wedge may be most consequential.

The model assumes that fines are proportional to speeding intensity rather than conditioned on discrete enforcement events. In practice, speeding fines are imposed only upon detection, with some probability $\pi < 1$. Introducing detection probability would rescale the effective fine: the optimal flat fine of \$3,262 with $\pi = 0.05$ (a plausible estimate for speed cameras) implies an expected fine of approximately \$163, close to the current US nominal average. This reinterpretation makes the optimal fine levels more plausible and connects to the Becker-Polinsky-Shavell framework Becker [1968], Polinsky and Shavell [1979], which emphasizes the probability-magnitude trade-off.

5.8 Broader implications

The double distortion mechanism applies whenever penalties are linked to economic productivity. Criminal day-fines, income-contingent environmental penalties, and means-tested regulatory sanctions all create the same implicit tax on earnings that we identify for traffic fines. The general principle is that linking any cost to income adds to the effective marginal tax rate, with efficiency consequences that coexist with equity and deterrence benefits.

This connects to a broader theme in public economics: the Tinbergen principle Tinbergen [1952], which holds that achieving k policy objectives requires at least k independent instruments. Income-based fines attempt to serve two objectives—deterrence and redistribution—with a single instrument. The Tinbergen principle would suggest that separating these functions—using flat fines for deterrence and the tax-transfer system for redistribution—would achieve both objectives more efficiently. Our calibration shows, however, that the efficiency cost of bundling deterrence and redistribution through income-based fines is quantitatively small enough that income-based fines still dominate flat fines. This represents a case where the single-instrument approach succeeds despite the theoretical prescription, because the implicit tax on labor is modest relative to the deterrence equity gain.

As the United States moves toward greater experimentation with income-based fines—following San Francisco's 2025 pilot San Francisco Municipal Transportation Agency [2025] and building on earlier experiments like Staten Island Hillsman [1990]—the framework we develop here provides a principled basis for evaluating these programs. The key empirical inputs—the distribution of marginal tax rates across the income distribution, labor supply elasticities, and speeding behavior—are all measurable, making our welfare comparison empirically grounded rather than purely theoretical.

6 Conclusion

This paper has analyzed the welfare implications of income-based traffic fines through the lens of optimal taxation theory. Our central finding is that, despite the "double distortion" mechanism we identify, income-based fines generate higher welfare than flat fines under realistic US calibration. The deterrence equity gain from income-proportional penalties exceeds the labor distortion cost in the large majority of parameter configurations.

6.1 Summary of findings

We developed a model where heterogeneous agents jointly optimize labor supply and speeding intensity under flat or income-based fine structures, with fine and tax revenue redistributed as a uniform transfer. The model incorporates the Nilsson power model Nilsson [2004] for the speed-fatality relationship and is calibrated to the United States using per-agent marginal tax rates computed from Enhanced Current Population Survey (CPS) microdata via PolicyEngine microsimulation PolicyEngine [2024], US Census Bureau [2024]. Behavioral parameters are drawn from informative priors grounded in meta-analytic estimates of labor supply elasticities Chetty [2012], US regulatory values for the value of statistical life (VSL) US Environmental Protection Agency [2024], and traffic fatality data from the National Highway Traffic Safety Administration (NHTSA) Fatality Analysis Reporting System (FARS) National Highway Traffic Safety Administration [2024].

The key results are as follows.

First, income-based fines create an effective marginal tax rate of $\text{MTR}_i + \phi s$ on labor income, where MTR_i is the agent's pre-existing marginal tax rate from the tax-benefit system, ϕ is the fine rate, and s is speeding intensity. This implicit tax is absent under flat fines (Proposition 1). At the baseline calibrated rate ($\phi = 0.02$), the additional EMTR is 0.2–2 percentage points; at the welfare-maximizing rate ($\phi^* \approx 0.076$), it rises to 0.8–4 percentage points—still small relative to the 0–50 percentage point range of pre-existing MTRs across the US income distribution.

Second, income-based fines generate higher utilitarian welfare than flat fines in 95% of Monte Carlo draws. The mean welfare difference favors income-based fines ($\Delta W = 0.83$, 95% CI: $[-0.02, 3.22]$). This result reflects the distributional gain from income-proportional penalties: low-income agents face lower fines, improving their consumption and welfare through the concavity of log utility. Income-based fines also reduce consumption inequality (Gini 0.325 vs.\ 0.343).

Third, the welfare advantage of income-based fines strengthens under inequality-averse social welfare functions. Under Rawlsian evaluation, income-based fines dominate in virtually all draws.

Fourth, welfare decomposition reveals that the deterrence channel is the dominant component: income-based fines reduce under-deterrence of high-income agents, where the convexity of the fatality-speed relationship makes speeding particularly costly. The labor distortion channel works against income-based fines but is quantitatively smaller.

6.2 Policy implications

Under our calibration, income-based fines generate higher welfare than flat fines, with several caveats relevant to US jurisdictions moving toward greater experimentation—including San Francisco's 2025 speed camera pilot San Francisco Municipal Transportation Agency [2025] and building on earlier experiments such as the Staten Island day-fine project Hillsman [1990].

The finding that income-based fines dominate under realistic calibration should not be taken as blanket endorsement. The labor distortion we identify is real, and our model's quadratic disutility specification implies a Frisch elasticity of 1.0—four times the empirical consensus—meaning the labor distortion in our simulations is larger than it would be at realistic elasticities. Even so, higher fine rates, contemporaneous income assessment, or specific subpopulations with high participation elasticities could tip the balance. Backward-looking income assessment, as in Finland's day-fine system Kaila [2024], would attenuate the labor distortion channel by severing the contemporaneous link between work effort and fine liability.

The interaction between income-based fines and the EITC phase-out deserves attention. Workers in the phase-out region already face effective marginal rates near 40% Maag et al. [2012]; adding an income-based fine compounds this distortion. While our results show the net welfare effect is positive—because these workers also benefit from lower fines relative to a flat system—the labor supply response in this population remains a key empirical unknown for evaluating income-based fines.

More broadly, our framework demonstrates that income-based fines can achieve both deterrence and redistributive objectives through a single instrument, contrary to the Tinbergen principle Tinbergen [1952]'s prescription of one instrument per objective. The practical question is whether the efficiency cost of this bundling is large enough to justify maintaining separate instruments—flat fines for deterrence and the tax-transfer system for redistribution Kaplow and Shavell [2002]. Our calibration suggests it is not.

6.3 Future work

Several extensions would strengthen the analysis. Natural experiments from fine system reforms—including San Francisco's ongoing pilot—could provide causal estimates of labor supply responses to income-based fines. A multi-period extension would capture the distinction between backward-looking and contemporaneous income assessment, habit formation in speeding, and human capital accumulation. Allowing for heterogeneous risk attitudes and driving needs across agents would refine the welfare calculations. Incorporating external costs of speeding—harm to other road users, congestion, and environmental damage—would strengthen the case for deterrence and likely reinforce the welfare advantage of income-based fines. Finally, extending the model to include the extensive margin of labor supply would capture participation responses that may be important for low-income workers near the EITC phase-out.

6.4 Concluding remarks

As the United States considers expanding income-based fine programs, the framework developed here provides a principled basis for evaluating the equity-efficiency trade-off. Under our calibration, the labor distortion cost is an order of magnitude smaller than the deterrence equity gain for most plausible parameter configurations, even with a Frisch elasticity four times the empirical consensus.

References

H. Allcott, B. B. Lockwood, and D. Taubinsky. Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, 134(3):1557–1626, 2019. doi: 10.1093/qje/qjz017.

G. S. Becker. Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217, 1968. doi: 10.1086/259394.

J.-M. Bourgeon and P. Picard. Point-record driving licence and road safety: An economic approach. *Journal of Public Economics*, 91(9):1603–1629, 2007. doi: 10.1016/j.jpubeco.2007.02.007.

A. L. Bovenberg and R. A. de Mooij. Environmental levies and distortionary taxation. *American Economic Review*, 84(4):1085–1089, 1994.

A. L. Bovenberg and L. H. Goulder. Optimal environmental taxation in the presence of other taxes: General-equilibrium analyses. *American Economic Review*, 86(4):985–1000, 1996.

A. Chalfin and J. McCrary. Criminal deterrence: A review of the literature. *Journal of Economic Literature*, 55(1):5–48, 2017. doi: 10.1257/jel.20141147.

R. Chetty. Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018, 2012. doi: 10.3982/ECTA9043.

R. Chetty, A. Looney, and K. Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, 2009. doi: 10.1257/aer.99.4.1145.

G. DeAngelo and B. Hansen. Life and death in the fast lane: Police enforcement and traffic fatalities. *American Economic Journal: Economic Policy*, 6(2):231–257, 2014. doi: 10.1257/pol.6.2.231.

P. A. Diamond and J. A. Mirrlees. Optimal taxation and public production I: Production efficiency. *The American Economic Review*, 61(1):8–27, 1971.

R. Elvik. A re-parameterisation of the Power Model of the relationship between the speed of traffic and the number of accidents and accident victims. *Accident Analysis & Prevention*, 124:135–141, 2019. doi: 10.1016/j.aap.2018.11.014.

M. Feldstein. Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics*, 81(4):674–680, 1999. doi: 10.1162/003465399558391.

N. Garoupa. The theory of optimal law enforcement. *Journal of Economic Surveys*, 11(3):267–295, 1997. doi: 10.1111/1467-6419.00034.

U. Gneezy and A. Rustichini. A fine is a price. *Journal of Legal Studies*, 29(1):1–17, 2000. doi: 10.1086/468061.

B. Hansen. Punishment and deterrence: Evidence from drunk driving. *American Economic Review*, 105(4):1581–1617, 2015. doi: 10.1257/aer.20130189.

A. C. Harberger. The measurement of waste. *The American Economic Review*, 54(3):58–76, 1964.

A. Harris. *A Pound of Flesh: Monetary Sanctions as Punishment for the Poor*. Russell Sage Foundation, 2016.

S. T. Hillsman. Fines as Criminal Sanctions: The Staten Island Day-Fine Experiment. techreport, National Institute of Justice, 1990. Day-fine pilot in Richmond County Criminal Court, 1988–1990.

B. Jacobs and R. A. de Mooij. Pigou meets Mirrlees: On the irrelevance of tax distortions for the second-best Pigouvian tax. *Journal of Environmental Economics and Management*, 71:90–108, 2015. doi: 10.1016/j.jeem.2015.01.003.

M. Kaila. How do people react to income-based fines? Evidence from speeding tickets in Finland. Working Paper 11064, CESifo, 2024.

E. Kantorowicz-Reznichenko and M. G. Faure. Day-fines: Should the rich pay more? *Review of Law & Economics*, 17(3):657–688, 2021. doi: 10.1515/rle-2021-0041.

L. Kaplow and S. Shavell. *Fairness versus Welfare*. Harvard University Press, 2002.

M. P. Keane. Labor supply and taxes: A survey. *Journal of Economic Literature*, 49(4):961–1075, 2011. doi: 10.1257/jel.49.4.961.

A. E. Lerman and V. M. Weaver. *Arresting Citizenship: The Democratic Consequences of American Crime Control*. University of Chicago Press, 2014. doi: 10.7208/chicago/9780226137988.001.0001.

E. Maag, C. E. Steuerle, R. Chakravarti, and C. Quakenbush. How marginal tax rates affect families at various levels of poverty. *National Tax Journal*, 65 (4):759–782, 2012. doi: 10.17310/ntj.2012.4.02.

M. D. Makowsky and T. Stratmann. Political economy at any speed: What determines traffic citations? *American Economic Review*, 99(1):509–527, 2009. doi: 10.1257/aer.99.1.509.

J. A. Mirrlees. An exploration in the theory of optimum income taxation. *The Review of Economic Studies*, 38(2):175–208, 1971. doi: 10.2307/2296779.

National Highway Traffic Safety Administration. Fatality Analysis Reporting System (FARS). https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars, 2024. URL https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars.

G. Nilsson. Traffic safety dimensions and the Power Model to describe the effect of speed on safety. Bulletin 221, Lund Institute of Technology, 2004.

PolicyEngine. Policyengine US: Open-source microsimulation model of the US tax-benefit system. https://policyengine.org, 2024. URL https://policyengine.org. Enhanced CPS microdata with computed marginal tax rates.

A. M. Polinsky and S. Shavell. The optimal tradeoff between the probability and magnitude of fines. *The American Economic Review*, 69(5):880–891, 1979.

A. M. Polinsky and S. Shavell. A note on optimal fines when wealth varies among individuals. *American Economic Review*, 81(3):618–621, 1991.

A. M. Polinsky and S. Shavell. *The theory of public enforcement of law*, volume 1, pages 403–454. Elsevier, 2007. doi: 10.1016/S1574-0730(07)01006-7.

E. Saez. Using elasticities to derive optimal income tax rates. *The Review of Economic Studies*, 68(1):205–229, 2001. doi: 10.1111/1467-937X.00166.

E. Saez and S. Stantcheva. Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1):24–45, 2016. doi: 10.1257/aer.20141362.

E. Saez, J. Slemrod, and S. H. Giertz. The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of Economic Literature*, 50(1):3–50, 2012. doi: 10.1257/jel.50.1.3.

San Francisco Municipal Transportation Agency. Speed Safety Camera Enforcement Program. https://www.sfmta.com/projects/speed-safety-cameras, 2025. URL https://www.sfmta.com/projects/speed-safety-cameras. Income-based fine discounts for low-income drivers.

A. Sandmo. Optimal taxation in the presence of externalities. *The Swedish Journal of Economics*, 77(1):86–98, 1975. doi: 10.2307/3439329.

D. Taubinsky and A. Rees-Jones. Attention variation and welfare: Theory and evidence from a tax salience experiment. *Review of Economic Studies*, 85(4): 2462–2496, 2018. doi: 10.1093/restud/rdx069.

J. Tinbergen. *On the Theory of Economic Policy*. North-Holland, Amsterdam, 1952.

US Census Bureau. Current Population Survey, Annual Social and Economic Supplement. https://www.census.gov/programs-surveys/cps.html, 2024. URL https://www.census.gov/programs-surveys/cps.html.

US Environmental Protection Agency. Mortality Risk Valuation. https://www.epa.gov/environmental-economics/mortality-risk-valuation, 2024. URL https://www.epa.gov/environmental-economics/mortality-risk-valuation. Central VSL estimate of $11.6 million (2024 dollars).