Externality of Driving Luxury Vehicles and Optimal Taxation

Jaewon Lee
Econic Partners

Suk Joon Son
University of Tokyo

Jintaek Song
University of Florida

Korea Insurance Research Institute

19 August 2026 · Seoul

Son gratefully acknowledges support from the Japan Science and Technology Agency (ERATO 230700000164). All views are the authors' own.

Introduction

Motivation: A Ferrari Is a Moving Disaster

- Two vehicles collide when one runs a red light and the other proceeds on a green light.
- Under tort law, the at-fault driver pays the damages to innocent driver .
- Suppose the innocent driver's car was…
  - **Case 1**: Ferrari ($200K car) with the repair cost ($100K).
  - **Case 2**: Civic ($20K car) with the repair cost ($10K).
- **Externality**: The *innocent driver*' s choice (Ferrari vs. Civic) affects how much the *at-fault driver*' s insurer pays.
- Ordinary actuarial pricing does not make the innocent driver pay for the social cost created by your vehicle value.

Research Questions and Relevance

Main Questions:

- How large is the externality from driving luxury vehicles?

Research Questions and Relevance

Main Questions:

- How large is the externality from driving luxury vehicles?
- How does a luxury vehicle tax affect equilibrium automobile demand/supply?

Research Questions and Relevance

Main Questions:

- How large is the externality from driving luxury vehicles?
- How does a luxury vehicle tax affect equilibrium automobile demand/supply?
- What is the welfare gain from an *optimal* tax?

Research Questions and Relevance

Main Questions:

- How large is the externality from driving luxury vehicles?
- How does a luxury vehicle tax affect equilibrium automobile demand/supply?
- What is the welfare gain from an *optimal* tax?

Research Questions and Relevance

Main Questions:

- How large is the externality from driving luxury vehicles?
- How does a luxury vehicle tax affect equilibrium automobile demand/supply?
- What is the welfare gain from an *optimal* tax?

Policy Context:

- In Korea, imported ($\simeq$ luxury) cars are 15% of insured vehicles, yet 33% of insurer costs.
- British Columbia and other regions have introduced or considered *value-based* premiums or higher luxury taxes.
- Global auto sales exceed $4 trillion/year (약 5,600조 원); thus, any small percentage of externality can translate into *very large* absolute costs.

What We Do

- **Estimate the size of the externality**:
  - Use model-level repair cost data (HLDI) to link vehicle value and insurer payouts.
- **Measure the equilibrium demand response to a luxury tax**:
  - Exploit BC’s 2018 tax reform to conduct a non-linear DiD (Athey & Imbens, 2006).
  - Novel transaction-level data to identify how higher taxes shift demand from expensive cars.
- **Quantify welfare gains and optimal policy**:
  - Use the DiD results above to estimate a model of automobile demand/supply.
  - Calculate how optimal taxes can reduce overconsumption of high-value cars.

What We Find

- **Externality Estimate**:
  - Each \$1 rise in car value $\rightarrow$ \$0.10 higher repair costs *paid by someone else' s insurer.*
- **Luxury Tax Impacts**:
  - Shifts consumers away from pricier cars, lowering total repair costs.
  - Estimated elasticity of vehicle value with respect to price $\approx$ 1.3.
- **Welfare Gain from an Optimal Tax**:
  - A simple 10% linear tax roughly internalizes the external cost.
  - This boosts welfare by $\sim 0.73\%$ of car value ($\approx \$7.3$B/year in the U.S., 약 10조 원).

A Toy Example

Toy Model with Two Cars

- Two cars: **luxury** (\$200K) and **economy** (\$20K), and everyone buys one
- Everyone values **economy** car at \$20K
- $\alpha \sim F(\alpha)$: valuation for the **luxury** car
- p: crash probability
- a: repair cost per dollar of vehicle value

Driving luxury rather than economy therefore adds an externality of

$$\text{Ext} = \underbrace{p}_{\text{crash prob.}} \cdot \underbrace{a}_{\text{repair share}} \cdot \underbrace{(\$200\text{K} - \$20\text{K})}_{\text{price gap}} = p \cdot a \cdot \$180\text{K}.$$

Privately, buy luxury if $\alpha > \$200\text{K}$,

Socially, buy luxury only if $\alpha > \$200\text{K} + p\,a\,\$180\text{K}$

Collision exposure is taken to be the same for both cars, so p does not itself depend on what you drive.

Deadweight Loss (Inefficiency) from Over-Consumption

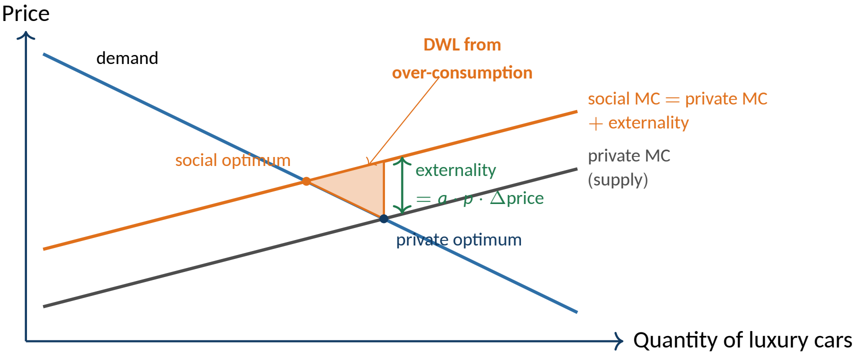

Privately, buyers equate demand with *private* MC: consumption exceeds the social optimum, and the **orange triangle** is the deadweight loss.

Pigouvian Tax Restores the Social Optimum

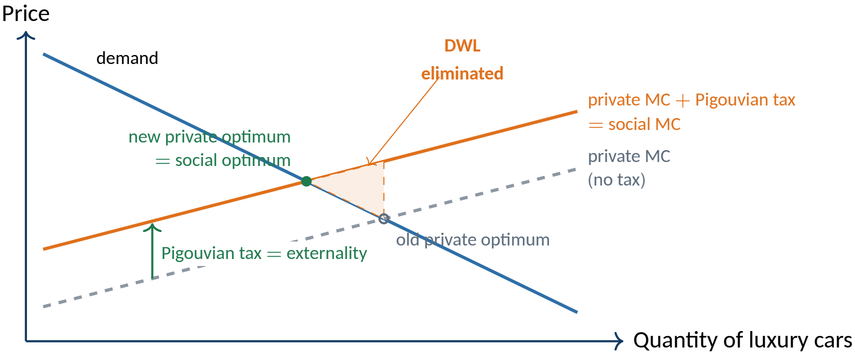

A tax equal to the externality shifts private MC up to social MC: buyers now face the full social cost, the two optima coincide, and the welfare gain is the recovered triangle.

Measuring the DWL: Two Quantities Are Sufficient

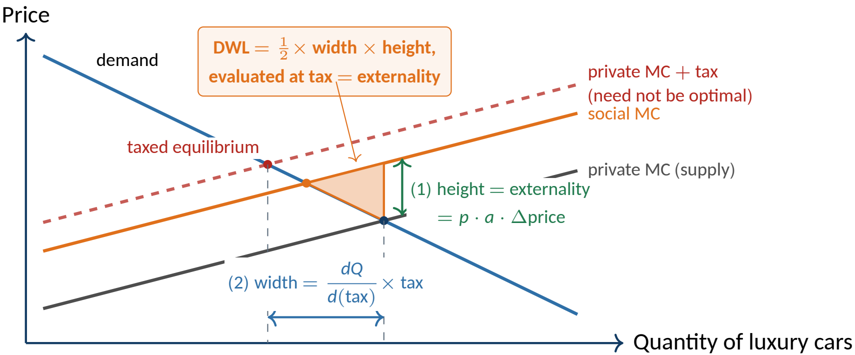

Two measurements suffice:

1. **size of externality** (height)
2. **quantity response per unit of (any) tax:** $\frac{dQ}{d(\textbf{tax})}$ (height/width)

Equilibrium Demand/Supply Response to Tax

British Columbia Tax Reform

2016: luxury-vehicle insurance charge, $+10$pp above \$150K

2018: luxury *sales tax*, flat 10% becoming 15% over \$125K and 20% over \$150K

The effective wedge τ a buyer faces:

- no reform: 10%, then $\mathbf{20}\%$ above \$150K
- after 2018: $10/15/\mathbf{30}\%$ (tax $+$ premium)

$\Rightarrow$ **notches** at \$125K and \$150K: the total bill *jumps*.

Want to know: how equilibrium quantities change with the 2018 reform.

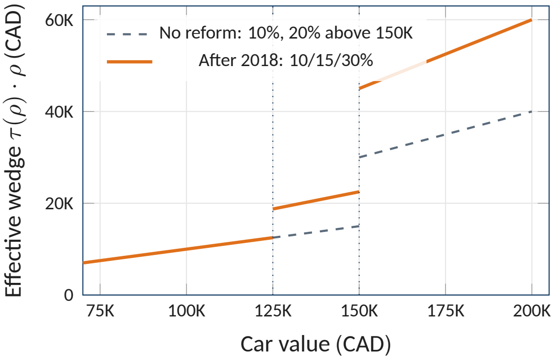

Sales-tax schedule plus the luxury-vehicle insurance charge.

Data: 26,867 Luxury Transactions

Records from an auto-financing platform in Canada.

- **26,867** completed deals above CAD 80K; BC and Quebec, 2016–2020
- Cash-equivalent price, MSRP / fair value, make–model–year, age, date, province
- Only transactions processed through the platform; all-cash purchases unobserved

The $80K floor is unlikely to bind for the substitution margin we study: the tax bites above $125K, and a buyer avoiding it is unlikely to drop below $80K.

The Tax Pushed Choices toward Cheaper Cars

Horizontal axis: what a car would have cost in a market with no luxury tax.

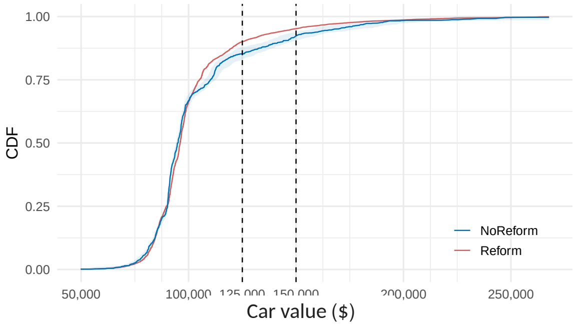

- Want to know: how equilibrium quantities change with the 2018 reform.
- We recover the *counterfactual* no-reform CDF of car values (how: later).
- The no-reform CDF **first-order stochastically dominates** $\Rightarrow$ the tax pushes choices toward **lower-value, cheaper** cars

Equality of the two distributions is rejected: $129.86 \sim \chi^2(9), p < 0.001$.

Shaded band is a 95% pointwise bootstrap CI. A placebo using 2016 vs. 2017, both pre-reform, gives two curves that nearly coincide: we find no comparable pre-reform shift. [Appendix]

Recovering the No-Reform CDF

- Simple before and after CDF: would reflect other factors that changed between two periods
- Need to *net out the trends* other than the reform.

Assumption: Distributional Parallel Trends (Athey and Imbens, 2006)

A BC buyer and a QC buyer at the *same rank* of their province's distribution at $t = 0$ would have picked the same car at $t = 1$ *absent* treatment.

Similar to DiD, but the tax hits fewer than 10% of cars, so a simple mean-DiD would miss it.

	Before reform	**After (w/o reform)**
BC's top-10% car	θ^* (BMW)	
QC's top-5% car	θ^* (BMW)	

- Same car at $t = 0$: BC's top-10% rank $=$ QC's top-5% rank

Recovering the No-Reform CDF

- Simple before and after CDF: would reflect other factors that changed between two periods
- Need to *net out the trends* other than the reform.

Assumption: Distributional Parallel Trends (Athey and Imbens, 2006)

A BC buyer and a QC buyer at the *same rank* of their province's distribution at $t = 0$ would have picked the same car at $t = 1$ *absent* treatment.

Similar to DiD, but the tax hits fewer than 10% of cars, so a simple mean-DiD would miss it.

	Before reform	**After (w/o reform)**
BC's top-10% car	θ^* (BMW)	
QC's top-5% car	θ^* (BMW)	θ^{**} (Ferrari)

- Same car at $t = 0$: BC's top-10% rank $=$ QC's top-5% rank
- We *observe* QC's top-5% buying θ^{**} at $t = 1$

Recovering the No-Reform CDF

- Simple before and after CDF: would reflect other factors that changed between two periods
- Need to *net out the trends* other than the reform.

Assumption: Distributional Parallel Trends (Athey and Imbens, 2006)

A BC buyer and a QC buyer at the *same rank* of their province's distribution at $t = 0$ would have picked the same car at $t = 1$ *absent* treatment.

Similar to DiD, but the tax hits fewer than 10% of cars, so a simple mean-DiD would miss it.

	Before reform	**After (w/o reform)**
BC's top-10% car	θ^* (BMW)	θ^{**} (Ferrari)
QC's top-5% car	θ^* (BMW)	θ^{**} (Ferrari)

- Same car at $t = 0$: BC's top-10% rank $=$ QC's top-5% rank
- We *observe* QC's top-5% buying θ^{**} at $t = 1$
- $\Rightarrow$ absent the tax, BC's top-10% would also have bought θ^{**}

Recovering a Tax-Invariant Vehicle Value Index

- $\theta \equiv$ inherent vehicle value, invariant to tax
- Normalization: vehicle's pre-tax price in BC in the **after-reform** period, but **w/o reform**
- Never observed; we use a DiD-style regression to **predict** it **(assumption: parallel trends)**
- $\theta =$ **predicted pre-tax price**, from a regression on MSRP and used value interacted with make, age, year, month and province
- Fit only on **untreated cells**: BC-before, and Quebec in both periods. The one tax-contaminated cell, BC-after, is dropped, so the coefficients never see the tax [Appendix]

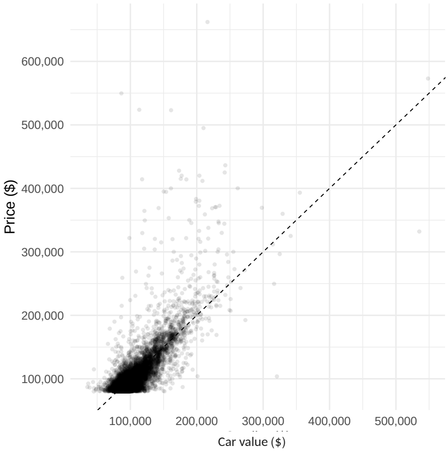

Predicted vs. actual pre-tax price, pooled over untreated cells.

Size of the Externality

Externality = 10% of Car Value

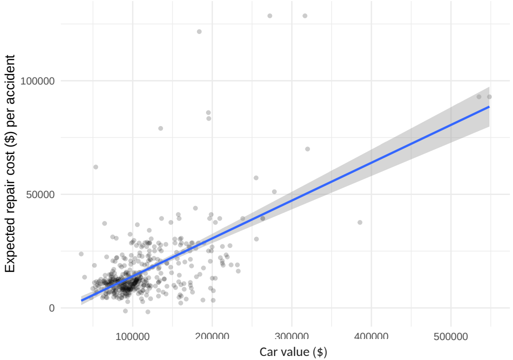

HLDI reports average *collision* loss by model: what the insurer pays to repair the policyholder's own car.

Regress that on vehicle value θ:

- incremental severity $= 0.1354$ per \$1 of value
- $\times$ annual collision probability 0.0609
- $\times$ average car life 12.1 years

$$\boxed{0.1354 \times 0.0609 \times 12.1 \approx \mathbf{0.0997}}$$

\$1 more of car value $\Rightarrow$ **\$0.10** of repair-cost externality.

Why collision can stand in for liability: a given car costs the same to repair regardless of whose insurer pays. Korean aggregates, where both sides are observed, give an imported/domestic payout ratio per accident of 2.67 under collision and 2.54 under liability. Average collision frequency is held fixed across vehicle value; dropping five outliers gives 0.0964.

What We Really Want vs. What We Can Measure

What we really want: the *liability* side — how much an extra $1 of vehicle value adds to the repair bills that **other drivers' insurers** pay.

What we actually compute: the *collision* side in HLDI data — how much an extra $1 of value adds to the owner's **own** repair losses: **0.0997**.

The two numbers are the same, under the following assumptions:

- A car's repair cost depends on its **own value**, NOT on the **car that hit it**.
- Fault share (과실비율) is unrelated to vehicle value.
- Crash probability does not depend on vehicle value.

Of course, we would love better data! Claim-level liability payouts would let us measure the externality directly $\to$ our question at the end.

Korean aggregates support the mapping *jointly*, not assumption by assumption: the imported/domestic severity ratio is 2.67 under collision vs. 2.54 under liability, and the aggregate liability gradient is 0.12 $\approx$ 0.0997. [Appendix]

Model of Demand/Supply & Welfare Effects

Why Need a Model of Demand/Supply?

Toy example suggests: we *seem to* know everything to compute

- Optimal (Pigouvian) taxation = externality (10% of car value)
- Height & Width of the deadweight loss triangle

But:

- in reality, choices b/w **multiple** differentiated vehicles
- we actually know response to BC's **nonlinear tax schedule** reform, not really $\frac{dQ}{d(\text{tax})}$
- toy assumed **linear demand/supply** curves

Buyers Trade Value against Price with an Elasticity of 1.3

Consumer i picks a vehicle value θ to maximize

$$\frac{\alpha_i}{r}\,\theta^r \;-\; \big(1+\tau\big)\big(p(\theta)+e_i\big)$$

- α_i: taste for luxuriousness
- e_i: bargaining ability, so identical cars sell at different prices
- r: curvature, $\sim$ **demand curve slope**
- $p(\theta)$: dealer prices, θ with no reform and $k\theta$ once taxed, so **k** $\sim$ **pass-through parameter**

Supply: Dealers set prices to max profit (Bertrand-Nash competition) [Appendix]

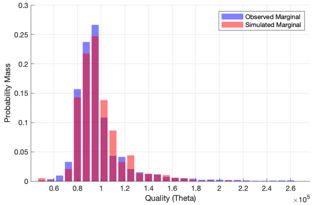

Simulated vs. observed θ, actual tax regime.

Estimates

$$\hat{r} = 0.21 \qquad \hat{k} = 0.97$$

- elasticity of chosen vehicle value $\approx$ **1.3**
- pre-tax prices sit 3% below the no-reform schedule

Bargaining dispersion: the bunching evidence behind e_i is in the [Appendix].

10% Tax Raises Welfare by CAD 733 per Car (Relative to No Tax)

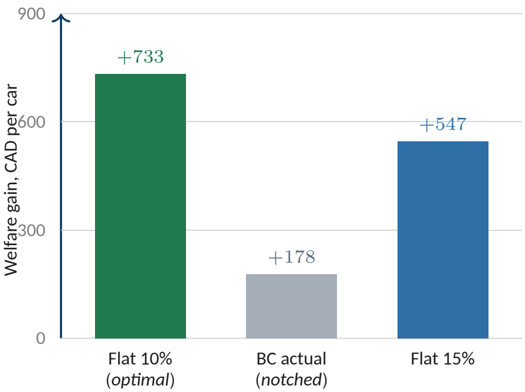

- The flat 10% benchmark gives the largest welfare gain: **CAD 733 per car**, or **0.73% of vehicle value** in our estimated BC high-end market
- BC's actual schedule delivers CAD 178 per car

Where does the CAD 733 come from?

Repair-cost externality avoided	$+1{,}237$
Other surplus effects	-504
Net welfare gain	$\mathbf{+733}$

Tax payments themselves are transfers: the CAD 10,500 increase in tax revenue largely offsets the measured loss in consumer surplus.

CAD per car, relative to no charge. Total surplus $= CS + PS + TR - RC$, with $RC = 0.0997 \times \mathbb{E}[\theta_i]$. Full table with bootstrap standard errors in the [Appendix].

Why Does the Flat 10% Outperform BC's Actual Schedule?

- Our estimated repair-cost externality is about **10% of vehicle value**
- A flat 10% charge therefore **aligns the private cost** of a more expensive vehicle with the additional repair cost it creates
- BC's actual schedule rises to an effective **30%** at the top. In our model, it reduces repair costs *more* than the flat 10% policy, but it also induces **too much substitution** away from high-value vehicles

	BC actual		**Flat 10%**
Repair-cost reduction	$1{,}776$	$>$	$1{,}237$
Welfare gain	CAD 178	$<$	CAD 733

Conclusion

- Under tort law, driving an expensive car creates a **repair-cost externality** that flat liability premiums fail to correct
- It is **\$0.10 per \$1** of vehicle value (U.S. collision data; Korean liability data agree at 0.12)
- BC's 2018 luxury tax **shifted demand toward cheaper cars**
- The externality estimate implies a **10% Pigouvian rate** under our benchmark; feeding the BC demand response into the model gives a welfare gain of **0.73% of vehicle value**, against $+$CAD 178 for BC's actual schedule

Policy implications for Korea: I will leave the discussion panelists.

Q: Better Repair Data from Korea?

What we currently do

- HLDI observes **own-vehicle collision claims**, not the liability cost imposed on the *other* driver
- We infer the latter assuming that repair technology, fault allocation, and accident exposure do not systematically vary with vehicle value
- Korean aggregate data provide a useful cross-check, but **not a direct test**

The ideal Korean data would observe, for each two-vehicle accident:

- vehicle value / model of **both** cars
- fault shares
- repair costs or insurance payouts for **both** cars
- ideally, whether each payment comes through collision or property-damage liability

Do such claim-level data exist in Korea?

Even data on liability payouts indexed by the *damaged* vehicle's model or value would be extremely useful.

Thank you

Suk Joon Son · sjson@e.u-tokyo.ac.jp

Appendix

Appendix: Korean Liability Data Give Almost the Same Number

Korea is the one place where we can see the *liability* side directly.

- Imported cars: 13.8% of registered passenger cars; payout ratio 2.42 vs. 0.78 for domestic (2019)
- With an average liability premium of about KRW 220,000, implied annual payouts are KRW 170,000 (domestic) vs. 525,257 (imported)
- Average vehicle price: KRW 33M (domestic) vs. 68M (imported); 12-year life

$$12 \times \frac{\mathbb{E}[\text{Payout} \mid \text{Imported}] - \mathbb{E}[\text{Payout} \mid \text{Domestic}]}{\mathbb{E}[\text{Price} \mid \text{Imported}] - \mathbb{E}[\text{Price} \mid \text{Domestic}]} = \mathbf{0.12}$$

Two very different calculations, the same order of magnitude

HLDI (U.S. collision, model-level): **0.0997**. Korea (liability, aggregate): **0.12**.

Source: Board of Audit and Inspection report; tables compiled by KIDI. A joint validation of the overall mapping, not a separate test of each assumption.

↩ back

Appendix: Korean Aggregates behind the 0.12

Insurance payout ratios (%), 2019

Vehicle type	Domestic	Imported	Total
Small	73.4	204.7	78.6
Medium	73.5	261.4	100.5
Large	96.9	232.4	133.7
Total	**78.4**	**241.8**	**101.2**

Share of registered vehicles (%), 2019

Vehicle type	Domestic	Imported	Total
Small	35.7	1.4	37.1
Medium	32.7	6.0	38.7
Large	17.8	6.4	24.2
Total	**86.2**	**13.8**	**100.0**

Payout ratio $= 100\times$ (payout / premium). Personal passenger cars. Source: report of the Board of Audit and Inspection of Korea; tables compiled by the Korea Insurance Development Institute from insurance data.

Collision as a proxy for liability

Per-accident payout, imported vs. domestic: **2.67** under collision (KRW 3.279M vs. 1.230M) and **2.54** under liability (KRW 2.891M vs. 1.140M).

Appendix: Welfare Table with Standard Errors

Scenario	ΔCS	ΔPS	ΔTR	ΔRC	Δ**Total**
Linear tax (10%)	$-11{,}091.84$	$+87.24$	$+10{,}500.48$	$-1{,}236.89$	$+733.10$
	(102.31)	(78.38)	(106.96)	(99.76)	(138.80)
Nonlinear (actual)	$-12{,}798.94$	$+36.29$	$+11{,}164.25$	$-1{,}776.12$	$+177.71$
	(1219.28)	(106.14)	(1286.15)	(116.31)	(174.40)
Linear tax (15%)	$-16{,}232.03$	$+58.12$	$+14{,}950.03$	$-1{,}769.86$	$+546.98$
	(149.19)	(114.63)	(198.65)	(190.36)	(99.16)

CAD per car unit, relative to the no-tax scenario. $TS = CS + PS + TR - RC$. Standard errors from 100 bootstrap samples in parentheses.

- Note the large standard errors on the *actual* nonlinear scheme: the notched schedule is estimated far less precisely, because the counterfactual price schedule must be solved at the discontinuities.

↩ back

Appendix: What Goes into θ, and What Stays Out

Governing idea: $\theta =$ the price a car *would* fetch in a world without the luxury tax. Two *separate* design choices follow.

A. Which rows to fit on?

Keep every untaxed cell; drop only cells contaminated by *this* tax.

	before	after
BC	✓	× (taxed)
QC	✓	✓

B. Which variables on the right-hand side?

In: luxury proxies (MSRP / used value, make, age); non-tax price shifters (province, year, month $\times$ MSRP).

Out: the tax rate, post-tax price, and any equilibrium *response* to the tax (quantities, BC $\times$ post).

Test: "would this value change *only because the tax exists*?" If yes, out.

The distinction that resolves the confusion

Invariance applies to the *mapping* car $\rightarrow \theta$, not to the *distribution* of θ, which should move post-tax. That is the effect. So province is a regressor (B), yet only BC-after is dropped (A).

↩ back

Appendix: Is Building θ the Same Assumption as CIC?

Building θ extrapolates BC-after from QC's time change, which *rhymes* with changes-in-changes. Same spirit, not the same assumption.

	CIC (outcome layer)	**θ extrapolation (measurement)**
Object	choice *distribution* (behavior)	price *level* of a fixed car
Form	non-parametric, rank-based	parametric linear (additive)
Trend on	the *distribution* (quantiles)	the conditional *mean* (DiD-style)

- Shared DNA: both say "absent the tax, BC would have followed the over-time change we see in QC."
- Irony worth acknowledging: CIC was *designed* to avoid parametric parallel trends on the distribution, yet θ quietly reintroduces a weaker parametric one, one layer down.
- Different layers $\Rightarrow$ not circular. But the result needs *both* to hold.

Appendix: Placebo Test (2016 vs. 2017)

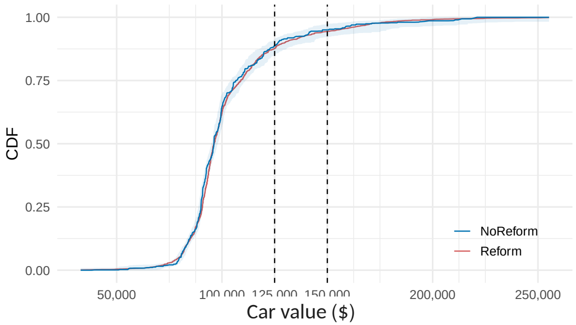

- Treat 2016 as “pre” and 2017 as a fake “post”, both before the real tax.
- The two curves nearly coincide.
- $\Rightarrow$ we find **no comparable pre-reform shift**.

⟵ back

Appendix: Bunching of Prices, Not of Underlying Values

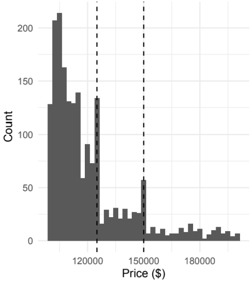

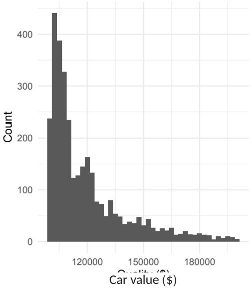

- Strong **bunching of pre-tax prices** at the \$125K / \$150K cut-offs $\cdots$
- $\cdots$but **no bunching in the underlying value θ**.
- One explanation: cars of the *same underlying value* sell at *different transaction prices*.
- This motivates modeling price dispersion within the same underlying value as bargaining heterogeneity e_i.

A buyer with more bargaining power gets a better car for the same notch-avoiding price.

↩ back

Appendix: Supply Side

Each value level θ is supplied by a dealer with constant marginal cost $c(\theta)$; dealers play Nash–Bertrand in list prices:

$$p(\theta) \;=\; \arg\max_{p(\theta)} \left(\underbrace{\mathbb{E}\left[p(\theta) + e_i \mid \theta_i = \theta; p\right]}_{\text{expected payment}} - c(\theta) \right) \underbrace{q(\theta; p)}_{\text{demand}} .$$

- Lets the model capture **tax pass-through** into pre-tax prices, without which we would confuse a price movement with a demand movement.
- In estimation: no-reform pricing $p(\theta) = \theta$ (a normalization, since θ *is* the no-reform predicted price); actual-economy pricing $p(\theta) = k\theta$, where k measures pass-through.
- θ is discretized on 15 grid points, chosen large enough that dealer profits are near zero, since we want to isolate the externality distortion from markup distortions.

⟵ back

Appendix: Matching Two Economies

Two worlds, both BC at $t = 1$, differing only in the tax schedule:

$$\tau^{\text{actual}} = \{.10,\ .15,\ .30\}, \qquad \tau^{\text{no reform}} = \{.10,\ .20\} \quad \text{(incl. the luxury premium).}$$

Minimise the squared distance between simulated and observed densities, jointly:

$$\min_{r,\sigma_e,k,f_\alpha} \sum_{\theta} \left(f^{\text{no reform}}_{\text{A-I}} - f^{\text{no reform}}_{\text{sim}} \right)^2 + \sum_{\theta,\rho} \left(f^{\text{actual}}_{\text{obs}} - f^{\text{actual}}_{\text{sim}} \right)^2 .$$

- $f^{\text{no reform}}_{\text{A-I}}$ is the changes-in-changes target from the reduced form; $f^{\text{actual}}_{\text{obs}}$ is the observed *joint* distribution of value and price.
- Solved with MATLAB's genetic algorithm with multi-start (global search over a nearly nonparametric f_α).

Appendix: Welfare Accounting

Given parameters and any tax τ, solve the dealer first-order conditions for prices, simulate choices, and add up:

$$CS = \int \max_{\theta} \left[\frac{\alpha_i}{r}\theta^r - (1+\tau)(p(\theta)+e_i) \right] \mathrm{d}F, \qquad TR = \int \tau\big(p(\theta_i)+e_i\big)\,\mathrm{d}F,$$

$$PS = \int \big(p(\theta_i)+e_i-c(\theta_i)\big)\,\mathrm{d}F, \qquad RC = 0.0997 \times \mathbb{E}[\theta_i].$$

Total surplus $TS = CS + PS + TR - RC$, compared across tax schemes against the no-tax world.

Optimal policy

With an externality of 0.0997 per $1 of value and near-perfect competition, the optimal corrective instrument is a **flat 10% of vehicle value**: textbook Pigou, now with the number attached.

Appendix: Why a Mean Difference-in-Differences Fails

Goal: how much did the 2018 tax shift demand toward cheaper cars?

- The tax is **highly nonlinear**: fewer than 10% of cars, at the very top of the distribution, are affected at all
- A standard difference-in-differences on the *mean* misses this, for two distinct reasons:
  - **Dilution and power.** An effect on $<10\%$ of cars shrinks to $\sim 0.1\times$ in the mean, and is swamped by the enormous dispersion of high-end prices
  - **Wrong object.** The tax reshapes the *upper tail* (truncation $+$ bunching); it does not shift the *location*. A constant-effect, central-tendency estimator is simply misspecified

Fix: compare the entire distribution, not the mean

Use **changes-in-changes** (Athey–Imbens 2006) to construct the *counterfactual* BC sales distribution had there been no tax, then compare it with what actually happened.

Appendix: The Changes-in-Changes Counterfactual

Stitching the rank-matching together across all quantiles gives the no-reform BC distribution:

$$F_{BC,1}^{\text{no reform}}(\theta) \;=\; F_{BC,0}\Big(F_{QC,0}^{-1}\big(F_{QC,1}(\theta)\big)\Big).$$

- Read right to left: take θ's rank in QC-after, map it to the QC-before value of that rank, then to BC-before's rank for that value
- Compare $F_{BC,1}^{\text{no reform}}$ (no tax) with the observed $F_{BC,1}$ (with tax)

What the method needs. A vehicle-value index θ that is (i) ordered by luxuriousness and (ii) *invariant to the tax*. Constructing it is the delicate step.

Appendix: Identification

Object	Variation that identifies it
$f_\alpha(\cdot)$	the distribution of vehicle value actually purchased
r	**how much the tax shifts the value distribution**, i.e. the CIC result
σ_e^2	price dispersion among cars of the *same* underlying value (the bunching fact)
$p(\theta)$, k	the observed value $\rightarrow$ price relation, and how it moves with the tax
$c(\theta)$	backed out from dealers' Nash–Bertrand first-order conditions

The link between the two halves of the talk

The changes-in-changes result is *exactly* the variation that identifies r. The structural model is not a separate exercise; it is the device that converts that one distributional shift into a welfare number.

Appendix: Estimation Details

- Discretization: 30 grid points each for θ, e_i and $p(\theta)$; 80 grid points for α_i; 15 grid points for θ in the supply/counterfactual step.
- For each (α, e) on the grid we solve the consumer problem under the relevant tax–price regime and aggregate into model-implied distributions.
- Objective: sum of squared deviations across both regimes (no-reform marginal $+$ actual joint), minimized with MATLAB's genetic algorithm and multi-start.
- Supply: under a *linear* tax, $\mathbb{E}[e_i \mid \theta_i = \theta] = 0$ because e_i enters additively and drops out of the optimal choice. This gives the simple first-order condition $q(\theta; p) + (p(\theta) - c(\theta))\, \partial q/\partial p(\theta) = 0$, from which $c(\theta)$ is backed out. Under nonlinear schemes this no longer holds, which is why the no-reform world is approximated as a 10% linear regime.