Faster and Deadlier? Disentangling Speed Limit Reversals from Pandemic Confounds in France

APEP Autonomous Research*

February 26, 2026

Abstract

France reduced speed limits on two-lane roads from 90 to 80 km/h in 2018. After the December 2019 *Loi d'Orientation des Mobilités*, 50 metropolitan départements staggered their reversal to 90 km/h between 2020 and 2023. I apply Callaway-Sant'Anna difference-in-differences to accident-level BAAC microdata spanning 2015–2024. A cross-département comparison shows 5 fewer accidents per quarter in reversal départements. However, a triple-difference comparing treated departmental roads to autoroutes—where limits never changed—within the same département isolates a significant increase of 3 additional corporal accidents per département-quarter on roads that reverted to 90 km/h. The sign reversal reveals that cross-département comparisons confound speed-limit effects with compositional differences between rural reversal départements and urban controls. Within-jurisdiction road-type controls are essential for credible speed-limit evaluation.

JEL Codes: R41, I18, H75, C23

Keywords: speed limits, road safety, difference-in-differences, triple-difference, France

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

On July 1, 2018, every two-lane secondary road in France dropped from 90 to 80 kilometers per hour. The measure was among the most unpopular policies in recent French history: the yellow-vest movement that erupted months later drew energy from rural drivers who saw it as a Parisian technocrat's imposition on their daily commutes. Yet the government's road safety observatory estimated 300 to 350 lives saved in the first year (**?**). Speed kills: kinetic energy scales with the square of velocity, so a 12.5 percent speed increase from 80 to 90 km/h raises impact energy by roughly 27 percent (**?**).

Then the dam broke. The December 2019 *Loi d'Orientation des Mobilités* (LOM) handed each *département* the right to restore 90 km/h on its roads, and by early 2026, 52 of France's 97 metropolitan départements had done so, covering roughly 61,000 of the 400,000 kilometers originally affected. But these reversals were staggered—Haute-Marne went first in January 2020, Eure was last in February 2026—creating a natural experiment tailor-made for modern causal inference. Two late adopters (Morbihan, July 2025; Eure, February 2026) fall outside the accident data window and serve as controls, leaving 50 treated départements for the analysis.

This paper asks a deceptively simple question: *did restoring the 90 km/h speed limit undo the safety gains?* The answer, it turns out, depends critically on what comparison group one uses. A Callaway-Sant'Anna difference-in-differences comparing reversal départements to those that maintained 80 km/h finds *fewer* accidents in the treated group—5 fewer corporal accidents (i.e., accidents involving bodily injury; property-damage-only incidents are not recorded in the BAAC) per département-quarter, with clean pre-trends in the event study. Taken at face value, this would imply the faster speed limit is actually safer.

This paper contributes to the speed-limit evaluation and staggered DiD literatures by showing that this cross-département comparison confounds the speed-limit effect with persistent compositional differences between rural reversal départements and more urban control départements. The solution is a triple-difference that compares changes on treated roads (routes départementales outside agglomeration) to changes on untreated roads (autoroutes) within the same département, using never-treated départements as the second level of differencing. This design eliminates any shock that affects all roads within a département equally.

The triple-difference reveals a statistically significant increase of approximately 3 additional corporal accidents per département-quarter on roads that reverted to 90 km/h ($p < 0.001$). The sign reversal—from -5 in the cross-département DiD to $+3$ in the within-département DDD—illustrates how comparing structurally different départements can mask genuine safety costs. A treatment-intensity specification within the DDD framework confirms that départements restoring larger shares of their road network experienced proportionally larger increases.

Three contributions emerge. First, I provide the first application of heterogeneity-robust staggered DiD estimators (?) to speed limit policy, in a setting with 52 treated départements (50 with observed post-treatment data) and 45 never-treated controls—well above the thresholds for credible inference. Second, I show that cross-département comparisons—even with modern estimators and clean pre-trends—can produce sign-reversed estimates when treatment correlates with urbanization, and demonstrate that within-jurisdiction road-type comparisons provide a more credible identification strategy. Third, the results speak directly to an active policy debate: as of 2026, French legislators continue to argue over whether to mandate a national return to 90 km/h. My findings suggest the 2018 reduction did produce genuine safety benefits that are partially reversed when speed limits are restored.

The paper is organized as follows. ?? describes the institutional background and policy timeline. ?? reviews the related literature. ?? details the data sources and panel construction. ?? presents the identification strategy, including the diagnostic that motivates the triple-difference. ?? reports the main results. ?? provides robustness checks. ?? discusses mechanisms and policy implications. ?? concludes.

2. Institutional Background

2.1 The 2018 Speed Limit Reduction

France's 80 km/h policy was announced by Prime Minister Édouard Philippe on January 9, 2018, and took effect nationally on July 1 of that year. It applied to all two-lane, bidirectional roads outside built-up areas that lacked a central reservation—approximately 400,000 kilometers of *routes départementales* and some *routes nationales*. The measure was justified on road safety grounds: secondary roads accounted for 55 percent of French road fatalities despite carrying a minority of traffic. The government cited the Conseil National de la Sécurité Routière (CNSR) recommendation and the Swedish experience with "Vision Zero" as precedents.

The policy was deeply unpopular, particularly in rural areas where these roads form the backbone of daily mobility. Opinion polls consistently showed 70–80 percent opposition, with hostility concentrated among frequent car users in péri-urban and rural zones. The measure became a flashpoint in the *gilets jaunes* protests that began in November 2018, alongside fuel tax increases and broader frustrations with Parisian centralization. Protesters explicitly targeted speed cameras—over 3,500 automated radar stations were vandalized or destroyed between November 2018 and March 2019, temporarily impairing enforcement.

The 80 km/h measure was not France's first national speed limit reduction. In 1990, the government reduced urban speed limits from 60 to 50 km/h, and in 1973, the first oil crisis prompted the introduction of national speed limits where none had previously existed. But the 2018 reform was unique in its scope (affecting the largest road network category) and in the intensity of the political backlash it provoked.

2.2 The Loi d'Orientation des Mobilités (LOM)

Partially in response to this backlash, the LOM of December 24, 2019, included a provision (Article 15 *bis* B) authorizing presidents of *conseils départementaux* to restore the 90 km/h limit on their departmental roads, subject to two procedural requirements. First, the département president must consult the departmental road safety commission (*commission départementale de sécurité routière*), which reviews local accident statistics and road characteristics. Second, the restoration must be implemented by prefectoral *arrêté* specifying each road segment affected. The law took effect January 1, 2020.

Importantly, the LOM granted discretion but not obligation. Départements could restore 90 km/h on all, some, or none of their eligible roads. This discretion produced the remarkable heterogeneity in treatment intensity that I exploit below: some départements restored their entire networks, while others restored only a handful of road axes. The law also required periodic review of accident data on restored segments, though enforcement of this review requirement varied.

2.3 Staggered Reversals

The first département to act was Haute-Marne, which restored 90 km/h on 15 road axes (476 km) on January 9, 2020—just nine days after the law took effect. Corrèze and Cantal followed on February 1, restoring their entire networks. The speed of these early movers suggests that political commitment to restoration preceded the LOM's formal authorization; several département presidents had publicly promised reversal during their 2015 election campaigns.

The rollout proceeded in three waves. The first wave (January–December 2020) comprised 15 départements, predominantly rural and politically conservative. The second wave (2021) was the largest: 23 additional départements adopted during a single calendar year, many in June 2021 around the anniversary of the policy option. The third wave (2022–2023) consisted of 12 later adopters, often more urbanized départements that conducted longer deliberations. No département reversed during 2024. As of the data cutoff (December 2024), 50 metropolitan départements had reversed within the BAAC observation window. Two additional départements (Morbihan and Eure) reversed after the data cutoff and serve as never-treated controls.

?? shows the cumulative adoption curve. The rollout displays substantial variation in both timing and intensity. Among the 50 départements with observed post-treatment data, only 7 restored 90 km/h across 100 percent of their eligible network; the median département reversed on just 9 percent of its network. This skewed distribution means that the average treatment effect is driven disproportionately by a small number of fully-treated départements. **??** provides the breakdown by year.

2.4 Treatment Intensity Variation

The heterogeneity in coverage is a defining feature of this natural experiment. Rural départements with sparse networks and low traffic volumes tended toward full restoration (Corrèze, Cantal, Creuse, Allier, Aveyron, Puy-de-Dôme, Ardèche), while more urbanized départements took a selective approach, restoring 90 km/h only on specific axes deemed safe (e.g., Seine-et-Marne: 26 percent; Haute-Marne: 20 percent). **??** visualizes this variation across the top 30 départements.

The sources of this heterogeneity are both technical and political. On the technical side, the LOM's procedural requirements meant that départements had to justify each road segment's suitability. Roads with high accident histories, complex intersections, or proximity to schools were typically excluded. Départements with simpler, more linear road networks found it easier to justify blanket restoration. On the political side, départements with right-leaning councils, which had been more vocally opposed to the 2018 reduction, moved faster and restored larger shares. Left-leaning councils, which had generally supported the safety rationale, were slower to act and more selective when they did.

The correlation between council partisanship, rural character, and reversal timing creates a potential selection concern: départements that reversed earliest and most extensively differ systematically from those that maintained 80 km/h. These are precisely the compositional differences that motivate the triple-difference design. I do not exploit political orientation instrumentally, though future work with commune-level election data could pursue this strategy.

2.5 Legal Complications

The reversal process was not always smooth. Several départements faced legal challenges from road-safety advocacy groups, most notably the *Ligue Contre la Violence Routière* and local chapters of the *Association Prévention Routière*. Calvados initially restored 90 km/h in August 2020, but the Administrative Court of Caen annulled the decrees in December 2022 on procedural grounds—specifically, the département had not adequately documented its consultation with the road safety commission. This forced a temporary return to 80 km/h until new, legally compliant decrees were issued in April 2023. Corrèze, one of the first movers, had its decrees annulled in January 2024 by the Administrative Court of Limoges, which found insufficient justification in the local accident data. The département continued operating at 90 km/h pending appeal, creating a period of legal ambiguity. Orne's decrees were annulled in November 2022 but revalidated in January 2023 after procedural corrections were filed. These legal episodes underscore the political salience of the policy and the administrative complexity of decentralized road regulation.

For the purposes of identification, I use the *de facto* date at which 90 km/h signage was installed, not the date of the prefectoral *arrêté*. In cases of annulment followed by reinstallation (Calvados, Orne), I use the most recent installation date. In cases of ongoing legal dispute where 90 km/h signs remained in place (Corrèze), I use the original installation date throughout. Sensitivity to these coding choices is minimal: the few affected départements represent a small share of the treated sample, and the DDD result is robust to dropping any individual département.

2.6 The Broader Policy Debate

The reversals unfolded against a backdrop of intense political debate. In the 2022 presidential campaign, several candidates (Marine Le Pen, Éric Zemmour, Valérie Pécresse) endorsed a national return to 90 km/h. The National Assembly voted in 2023 on a proposition to restore 90 km/h nationally, which narrowly failed. The Cour des Comptes (France's national audit institution) published a critical report in November 2023 arguing that the patchwork of speed limits across départements created confusion and undermined road safety coherence. As of early 2026, the debate continues: the Rassemblement National has made the restoration of 90 km/h a campaign pledge for the next parliamentary elections.

This political context matters for identification in two ways. First, the salience of the policy means that département decisions to reverse were not made in obscurity—they were high-profile political acts, often announced by council presidents at press conferences. Second, the ongoing debate makes the evaluation directly policy-relevant: legislators considering a national mandate will benefit from evidence on whether the partial, voluntary reversals produced measurable safety effects.

3. Related Literature

This paper contributes to three literatures: the economics of speed limits, COVID-era confounding in policy evaluation, and the methodology of staggered difference-in-differences.

3.1 Speed Limits and Road Safety

The relationship between speed and road safety is one of the oldest questions in transportation economics. Nilsson's "power model" (**?**) established that the number of fatalities scales approximately with the fourth power of the speed ratio: a 10 percent speed increase raises fatalities by roughly 40 percent. **?** re-parameterized this model and confirmed its applicability across diverse settings, estimating that the exponent ranges from 2 (for all injury accidents) to 4.5 (for fatal accidents), depending on the severity outcome.

The causal identification of speed-limit effects has relied on two main strategies. The first exploits discrete policy changes. **?** used the 1987 US federal speed limit increase from 55 to 65 mph on rural interstate highways, finding a statistically significant increase in fatalities that they used to estimate the value of a statistical life at $1.54 million (1997 dollars). Their approach—comparing states that raised limits to those that did not—is closest in spirit to the present paper, though they used a single nationwide policy change rather than staggered adoption. The 1995 repeal of the National Maximum Speed Law prompted further studies: several found 15–35 percent increases in fatalities on affected highways, concentrated in states that raised limits to 70 mph or higher.

The second strategy uses regression discontinuity at speed-limit boundaries, exploiting the discrete jump in posted limits along a continuous road. These designs recover local treatment effects at the boundary but cannot speak to the general equilibrium impact of system-wide limit changes.

For France specifically, **?** studied the original 2018 reduction from 90 to 80 km/h using aggregate time series with group-level variation. They estimated 300–350 lives saved annually and a 10 percent fatality reduction on affected roads. The government's road-safety observatory (ONISR) produced similar figures using before-after comparisons with seasonal adjustments. However, both analyses rely on pre-post comparisons with limited controls for confounding trends, and neither addressed the possibility that the economic cycle or other secular trends could have independently reduced fatalities. No published study examines the *reversal*—the natural experiment created by the LOM—using modern causal inference methods.

3.2 COVID Confounding in Quasi-Experiments

The pandemic created a pervasive confound for any policy evaluated during 2020–2022. Mobility collapsed asymmetrically: urban areas experienced sharper declines, while rural areas maintained higher driving volumes (**?**). Google Mobility Reports for France showed that transit use in major cities fell by 70–80 percent during the first lockdown (March–May 2020), while driving in rural départements declined by only 30–40 percent. This asymmetry contaminates any cross-sectional comparison where treatment correlates with urbanization.

? documented how COVID-era confounding can bias DiD estimates when treatment timing coincides with mobility shocks. They showed that many quasi-experimental designs spanning 2020 inadvertently captured pandemic effects rather than policy effects, because the parallel trends assumption can hold in the pre-period yet fail catastrophically when an aggregate shock differentially affects treated and control units. **?** showed that the geographic incidence of pandemic restrictions was far from random, with rural areas facing lighter mobility constraints due to lower population density and hospital capacity buffers.

The present paper's setting is a canonical example of this problem. The first cohort of speed-limit reversals occurred in early 2020, perfectly coinciding with the onset of the pandemic. Moreover, the reversal départements are overwhelmingly rural—precisely the areas where mobility was least disrupted. Rather than attempting to "control for" COVID through covariates (an approach fraught with functional-form assumptions), the paper contributes a structural solution: the within-jurisdiction road-type comparison that differences out any département-wide shock, including pandemic-driven mobility changes.

3.3 Staggered DiD Methodology

The econometric revolution in difference-in-differences has produced estimators that handle treatment effect heterogeneity under staggered adoption. The problem is now well understood: in a standard TWFE regression with staggered treatment timing, the estimated coefficient is a weighted average of all two-by-two DiD comparisons in the data, including comparisons of early-treated to late-treated units where the late-treated serve as "controls" even though they will eventually be treated. Under treatment-effect heterogeneity, some of these implicit weights can be negative, producing estimates that do not correspond to any meaningful causal parameter.

? proposed group-time ATTs with flexible aggregation, comparing each cohort to a clean never-treated or not-yet-treated control group. **?** developed interaction-weighted estimators that recover cohort-specific effects. **?** provided the Bacon decomposition, which decomposes the TWFE estimator into its constituent two-by-two comparisons and their weights, making the source of potential bias transparent. ? proved that TWFE can assign negative weights to some group-time cells under heterogeneity.

I implement several of these estimators—Callaway-Sant'Anna, Sun-Abraham, and the Bacon decomposition—alongside traditional TWFE. The estimators broadly agree on the sign and magnitude of the cross-département treatment effect. But the critical insight is that none of these methodological advances can overcome the *fundamental* identification problem when the confound operates at the group level. Treatment-effect heterogeneity is a within-treated concern; differential trends between treated and control groups is a selection concern. The triple-difference is necessary precisely because the issue is selection into treatment correlated with urbanization and road network composition, not heterogeneous effects across cohorts.

4. Data

4.1 Road Accident Microdata (BAAC)

The primary outcome data come from the *Bulletin d'Analyse des Accidents Corporels de la Circulation Routière* (BAAC), France's official database of road accidents resulting in bodily injury. The BAAC is compiled by law enforcement (*gendarmerie* for rural roads, *police nationale* for urban areas) and contains the universe of corporal accidents—any road accident involving at least one injury or fatality. Reporting is mandatory: officers complete a standardized form for every such accident. Property-damage-only accidents are excluded from the database.

I use annual files for 2015–2024 (10 years), publicly available on `data.gouv.fr`. The BAAC comprises three linked tables: *Caractéristiques* (accident-level: date, département, agglomeration status), *Lieux* (road-level: road category, speed limit, infrastructure), and *Usagers* (victim-level: severity—killed within 30 days, hospitalized, or light injury). Each accident has a unique identifier (`Num_Acc`) that links the three tables.

Each accident record includes the département code, road category (*autoroute*, *route nationale*, *route départementale*, *voie communale*), agglomeration status (inside or outside built-up areas), date, and GPS coordinates. The road category variable (`catr`) is crucial for the identification strategy: it allows me to distinguish departmental roads (where the speed limit changed) from autoroutes (where it did not) within the same geographic unit.

I filter to accidents on *routes départementales* outside agglomeration (`catr`=3, `agg`=2)—the exact road population affected by the 80/90 km/h policy. The "outside agglomeration" restriction is important: within built-up areas, routes départementales have a 50 km/h limit regardless of the 80/90 policy. For placebo tests, I separately extract autoroute accidents (`catr`=1) and urban *route départementale* accidents (`catr`=3, `agg`=1). The autoroute sample serves as the primary within-département control in the triple-difference design; the urban departmental road sample provides a secondary placebo.

4.2 Data Quality and Coding Changes

The BAAC files undergo format changes across years that require careful harmonization. Before 2019, accident years are encoded as two-digit values (e.g., "15" for 2015); from 2019 onward, four-digit years are used. The 2022 vintage introduced a column rename (`Accident_Id` replacing `Num_Acc`). Département codes also vary: pre-2019 files use three-digit codes (e.g., "590" for département 59, Nord), while later files use two-digit codes. Corsica presents an additional challenge with alphanumeric codes ("2A" and "2B"). My data-cleaning pipeline handles all these variants systematically, mapping every record to a standardized two-character département code.

The BAAC has known limitations. Underreporting of minor injuries is endemic across all police-recorded crash databases; studies estimate that the BAAC captures roughly 80 percent of hospitalizations and a lower share of light injuries. The reporting rate is higher for fatal accidents (near 100 percent) and for accidents on major roads. Since my primary analysis uses total corporal accidents (not just fatalities), any differential underreporting between road types could bias the DDD estimate. However, the within-département design mitigates this concern: the same police forces report accidents on both departmental roads and autoroutes within a given département, so reporting practices should be similar for both road types.

4.3 Treatment Panel

I compile the treatment panel from multiple cross-validated sources: the *Ligue de Défense des Conducteurs* observatory (which maintains a continuously updated map of reversals), *L'Argus* and *France Info* interactive maps, and ONISR annual reports. For each département, I record the effective date (month-year) of the reversal and the approximate share of the departmental road network restored to 90 km/h. Where sources conflict, I consult prefectoral *arrêtés* via Légifrance as the authoritative source. This cross-validation is necessary because several départements announced reversals that were subsequently delayed, modified, or annulled.

The treatment intensity variable (*share_pct*) measures the percentage of each département's eligible road network restored to 90 km/h. This variable ranges from 2 percent (token restoration of a few road segments) to 100 percent (full network restoration). The measurement is approximate, as départements do not report a standardized denominator. I use the most recent published estimate from each source, prioritizing departmental council deliberations.

Two départements (Morbihan, reversed July 2025; Eure, reversed February 2026) are excluded from the treated group because no post-treatment BAAC data are available. They serve as never-treated controls in the Callaway-Sant'Anna specification. Four urban départements (Paris, Hauts-de-Seine, Seine-Saint-Denis, Val-de-Marne) have negligible eligible road networks—the vast majority of their roads are classified as *voies communales* or *routes nationales*—and are also classified as never-treated.

4.4 Panel Construction

I construct a balanced département × quarter panel from Q1 2015 through Q4 2024 (40 quarters, 97 metropolitan départements, 3,880 observations). Overseas départements (971–976) are excluded because the 80 km/h policy did not apply uniformly in these territories. The outcome variables are: total corporal accidents, fatalities (killed within 30 days), hospitalizations (admitted for at least 24 hours), light injuries, and total casualties. I also compute the severity ratio (killed plus hospitalized as a share of total accidents) as a measure of crash severity conditional on occurrence.

The panel is balanced by construction: every département appears in every quarter. Zero-accident quarters are common, particularly for smaller rural départements, and are coded as zeros rather than missing values. This is appropriate because BAAC coverage is universal—the absence of records genuinely reflects the absence of corporal accidents on the relevant road type in that département-quarter.

For the DDD analysis, I construct a stacked panel that includes both departmental road accidents and autoroute accidents for each département-quarter. This doubles the panel to $2 \times 3{,}880 = 7{,}760$ observations, with a road-type indicator distinguishing the treated road category (departmental roads outside agglomeration) from the control road category (autoroutes). The DDD panel uses département × road-type fixed effects ("cells") and quarter fixed effects.

?? reports pre-period (2015–2019) summary statistics. Treated départements average fewer accidents per quarter on affected roads compared to control départements. This level difference—driven by the urban composition and larger population of the control group—is absorbed by département fixed effects in the DiD design. The key identifying assumption is not that treated and control départements have similar accident levels, but that they would have experienced parallel trends absent the policy change.

5. Identification Strategy

5.1 Staggered Difference-in-Differences

The baseline specification exploits the staggered timing of département-level reversals. Let Y_{dt} denote the outcome (accidents, fatalities, etc.) in département d and quarter t, and let G_d denote the quarter in which département d first restored 90 km/h (with $G_d = \infty$ for never-treated units). The Callaway-Sant'Anna estimator (**?**) identifies group-time average treatment effects:

$$\text{ATT}(g,t) = \mathbb{E}[Y_t(g) - Y_t(\infty) \mid G = g] \quad \text{for } t \geq g \tag{1}$$

under the assumption of parallel trends conditional on group and time. I use never-treated départements as the comparison group ($n = 45$), set anticipation to zero periods, use the universal base period (all pre-treatment periods), and compute standard errors via the analytical formula with département-level clustering. I report both the simple aggregate ATT and event-study coefficients (with event-time window $[-8, +16]$ quarters).

The key advantage of the Callaway-Sant'Anna estimator is that it avoids using already-treated units as controls. In our setting, where 15 départements adopted in 2020 and 23 more in 2021, a naive TWFE regression would compare the 2021 cohort partly against the 2020 cohort—whose outcomes already reflect the treatment. The Callaway-Sant'Anna estimator compares each cohort exclusively to never-treated units, eliminating this "forbidden comparison" bias.

For comparison, I estimate the traditional two-way fixed effects (TWFE) specification:

$$Y_{dt} = \alpha_d + \gamma_t + \beta \cdot \text{Post}_{dt} + \varepsilon_{dt} \tag{2}$$

where α_d and γ_t are département and quarter fixed effects, $\text{Post}_{dt} = \mathbb{I}\{t \geq G_d\}$, and standard errors are clustered at the département level. The TWFE estimand is a weighted average of all two-by-two DiD comparisons in the data, including both clean (treated-vs-never-treated) and potentially biased (early-vs-late-treated) comparisons. Under treatment effect heterogeneity, some implicit weights can be negative (**??**). I report TWFE alongside Callaway-Sant'Anna for comparability with the existing speed-limit literature, which has relied on TWFE specifications.

5.2 The Placebo Diagnostic

Before interpreting the DiD estimates, I conduct a critical falsification test. Autoroute speed limits (130 km/h, or 110 km/h in wet conditions) were never affected by the 80/90 reform. The LOM provision applied exclusively to two-lane secondary roads; autoroutes are managed by national concession companies and their speed limits are set by national decree. If the DiD captures the causal effect of the speed limit reversal, we should observe *no* effect on autoroute accidents in the same départements.

The autoroute placebo passes: the TWFE coefficient on autoroute accidents is +0.88 (SE = 0.70, $p = 0.21$), small in magnitude and statistically insignificant. At first glance, this would seem to *validate* the cross-département comparison. But the placebo's benign result coexists with a main estimate that is itself insignificant (-6.33, $p = 0.19$), suggesting that the cross-département comparison simply lacks precision rather than being cleanly identified. The autoroute result provides reassurance that no *large* département-level confound contaminates the comparison, but cannot rule out moderate compositional differences between the predominantly rural reversal départements and more urban control départements.

The concern is best understood through a concrete example. Consider département 19 (Corrèze, population 240,000, fully treated) versus département 75 (Paris, population 2.1 million, never-treated). Both appear in the panel with département fixed effects absorbing level differences. But if accident trends in dense urban areas diverge from those in rural areas—due to changing commuting patterns, ride-sharing adoption, infrastructure investment, or demographic shifts—the parallel trends assumption fails regardless of the speed-limit policy. The sign reversal between the cross-département DiD (negative) and the within-département DDD (positive, as shown below) confirms that the comparison group matters critically.

5.3 Triple-Difference Design

These concerns motivate a triple-difference (DDD) that nets out département-wide trends. I stack two road types for each département—treated departmental roads and untreated autoroutes—and estimate:

$$Y_{drt} = \alpha_{dr} + \gamma_t + \delta_1 \cdot \text{Post}_{dt} + \delta_2 \cdot \text{Post}_{dt} \times \text{DeptRoad}_r + \varepsilon_{drt} \tag{3}$$

where α_{dr} is a département × road-type fixed effect, Post_{dt} indicates the post-reversal period for département d, and DeptRoad_r equals one for departmental roads (the treated road type). The coefficient δ_2 is the DDD estimate: the *differential* change on treated roads relative to autoroutes, after the reversal, in treated versus control départements. The stacked panel doubles the observation count to 7,760 (97 départements × 2 road types × 40 quarters).

Standard errors are clustered at the département level; I also report two-way clustering (département × quarter) as a robustness check.

This design eliminates any shock that affects all roads within a département equally—including COVID-driven mobility changes, weather patterns, and local law enforcement intensity. The identifying assumption is that absent the speed limit reversal, the *gap* between departmental road and autoroute accidents would have evolved similarly in treated and control départements.

This assumption is weaker than the standard DiD parallel trends assumption in an important respect: it does not require that treated and control départements have parallel accident *levels*. It only requires that the *difference* between road types within each département follows a common trend. A national recession, a pandemic, or a change in enforcement policy that affects all roads within a département proportionally is differenced out. The assumption fails if something differentially changes the departmental-road-to-autoroute gap in treated départements—for example, if treated départements simultaneously improved their departmental road infrastructure but not their autoroutes. I am not aware of any such policy.

I also estimate a treatment-intensity variant of the DDD that replaces the binary $\text{Post}_{dt} \times \text{DeptRoad}_r$ with $\text{Share}_d \times \text{Post}_{dt} \times \text{DeptRoad}_r$, where Share_d is the fraction of the departmental road network restored to 90 km/h. This specification tests for a dose-response relationship: if the speed limit reversal truly increases accidents, départements that restored larger network shares should see proportionally larger increases on their departmental roads relative to autoroutes.

6. Results

6.1 Cross-Département DiD: A Negative Effect

?? reports the main estimates. The Callaway-Sant'Anna ATT on total accidents is −5.01 (SE = 1.10), suggesting that reversal départements experienced *fewer* accidents per quarter after restoring 90 km/h. The point estimate is precise and statistically significant, which—taken at face value—would imply the counterintuitive conclusion that faster speed limits reduce accidents. The TWFE estimate is also negative at −6.33 (SE = 4.85) but statistically insignificant ($p = 0.19$). The disagreement between the two estimators is informative: the Callaway-Sant'Anna approach, by restricting comparisons to never-treated units, produces a tighter estimate than TWFE, which also uses early-versus-late comparisons that may introduce noise. For fatalities, neither estimator finds a significant effect: the CS-DiD ATT is −0.13 (SE = 0.17), and the TWFE estimate is −0.03 ($p = 0.78$).

?? plots the event-study coefficients from the CS-DiD. Pre-treatment coefficients are generally small and statistically insignificant, supporting the parallel trends assumption for the cross-département comparison. Post-treatment coefficients turn negative starting at $e = 0$ and grow in magnitude through $e = 8$ (two years post-treatment), suggesting a persistent relative decline in accidents in reversal départements. The pattern is consistent with a "true" effect—if one trusts the cross-département comparison.

?? shows raw annual trends by treatment status. Two features are immediately apparent. First, treated départements have substantially lower accident levels throughout the panel (mean: 12 per quarter vs. 37 for controls), reflecting the rural composition of the reversal group. Second, both groups display a sharp dip in 2020 (COVID), followed by incomplete recovery. The parallel trends visual is reasonably clean, which is precisely why the CS-DiD produces clean pre-trends and a "significant" negative estimate. The problem, as the next subsection shows, lies not in pre-trends but in post-treatment compositional shifts.

6.2 The Placebo Test and What It Reveals

The autoroute placebo (**??**) estimates +0.88 (SE = 0.69, $p = 0.21$)—positive and statistically insignificant. I interpret this as follows: when comparing reversal départements to non-reversal départements using *autoroute* accidents as the outcome (where speed limits never changed), there is no detectable difference. This rules out large département-level confounds—for example, if reversal départements experienced a general increase in law enforcement or road maintenance that affected all road types, we would detect it here. But the test has limited power: with an SE of 0.69, we cannot reject modest confounds.

The urban departmental road placebo is similarly benign (+0.72, SE = 1.58, $p = 0.65$). Urban departmental roads, which have a 50 km/h limit regardless of the 80/90 policy, show no differential trend in reversal départements relative to controls. Both placebos pass, yet the main TWFE estimate is itself insignificant. This combination—insignificant main effect, insignificant placebos—suggests low statistical power rather than clean identification.

?? plots quarterly accident trends for treated roads and autoroutes within reversal départements only. Both series display common temporal patterns: the 2020 COVID dip, partial recovery in 2021, and stabilization thereafter. The key visual insight is that *within* reversal départements, departmental roads and autoroutes follow parallel trajectories. This motivates the within-département comparison: if the two road types within the same département track each other in the absence of the policy, their differential change after the policy isolates the speed-limit effect.

6.3 Triple-Difference: The Causal Effect

The DDD estimate (**??**) is +3.05 (SE = 0.89, $p < 0.001$). After netting out the autoroute trend within each département, the speed limit reversal is associated with a statistically significant *increase* of approximately 3 additional corporal accidents per département-quarter on treated departmental roads. This is the headline result of the paper.

This estimate is an intent-to-treat (ITT) effect of the département-level policy package, since the binary DDD assigns treatment at the département level regardless of the share of the road network actually restored. The median treated département reversed approximately 9 percent of its departmental road network (see **??**), so the per-percentage-point effect is roughly $3.05/9 \approx 0.34$ accidents. The intensity specification (Section **??**, Treatment Intensity) confirms this dose-response relationship directly.

?? presents the DDD event study—the dynamic counterpart to the static DDD estimate. Pre-treatment coefficients (relative quarters −8 through −2) are small in magnitude (ranging from −1.5 to +0.5) and statistically insignificant, validating the parallel trends assumption for the road-type gap between treated and control départements. Post-treatment coefficients gradually increase, reaching statistical significance by quarter +5 and growing through quarter +12. This pattern of gradually escalating effects is consistent with behavioral adaptation: as drivers internalize the higher speed limit, speeds drift upward and accident risk accumulates.

The sign reversal is striking. The cross-département Callaway-Sant'Anna ATT is −5.01: reversal départements have fewer accidents. The within-département DDD is +3.05: treated roads within reversal départements have more accidents than untreated roads within the same département, relative to the same road-type gap in control départements. The difference—8.06 accidents per département-quarter—is the magnitude of the compositional confound. Rural départements have structurally declining accident trajectories relative to urban départements, and the cross-département DiD conflates this differential trend with the speed-limit effect.

The DDD estimate of +3.05 is economically meaningful. With 50 treated départements, the implied aggregate effect is approximately $3.05 \times 50 \times 4 = 610$ additional corporal accidents per year on the treated road network. Given that the total number of corporal accidents on departmental roads outside agglomeration averages roughly 15,000 per year nationally, the reversal accounts for approximately a 4 percent increase.

6.4 Fatalities: Underpowered

The CS-DiD ATT on fatalities is −0.13 per quarter (SE = 0.17), small and statistically insignificant. **??** shows the fatality event study: point estimates are noisy, clustering around zero with wide confidence intervals throughout both the pre- and post-treatment windows.

The severity ratio is essentially unchanged (-0.002, $p = 0.93$), providing no evidence that the reversal altered crash severity conditional on an accident occurring.

This null result is best understood through a power calculation. With a pre-period mean of roughly 1.2 fatalities per département-quarter on treated roads, the standard deviation is approximately 1.5 (fatalities follow a near-Poisson distribution with moderate overdispersion). With 50 treated départements, 45 control départements, 20 pre-treatment quarters, and an average of 8 post-treatment quarters per cohort, the minimum detectable effect (MDE) at $\alpha = 0.05$ and 80 percent power is approximately 0.45 fatalities per département-quarter—a 35 percent increase relative to the mean. Nilsson's power model would predict a fatality increase of roughly 5 percent for a 12.5 percent speed increase (from 80 to 90 km/h), far below the MDE.

The design is thus well-powered for accident frequency—the pre-period mean of approximately 12 accidents per quarter and a standard deviation of 10 yield an MDE of roughly 15 percent, comfortably below the observed 25 percent DDD effect (3.05/12). For fatalities, however, the design is structurally underpowered for the expected effect magnitude. A fatality analysis would require either a much longer panel (to accumulate more post-treatment observations) or a different data structure (e.g., national monthly data disaggregated by road type) to achieve adequate power. I therefore interpret the fatality null as uninformative rather than as evidence of no effect.

6.5 Treatment Intensity

A continuous treatment specification using the share of each département's road network restored to 90 km/h yields a naive TWFE coefficient of -3.27 (SE = 5.01, $p = 0.52$), negative and insignificant. However, the DDD intensity specification—interacting share with the departmental-road indicator in the stacked framework—yields a positive coefficient of $+5.95$ (SE = 1.65, $p < 0.001$). Départements that restored larger network shares experienced proportionally larger accident increases on their treated roads, confirming a dose-response pattern within the triple-difference framework.

7. Robustness

7.1 COVID Exclusion

The first wave of reversals coincided with the pandemic, raising concerns that COVID-specific mobility changes drive the results. I address this in two ways. First, I exclude the most pandemic-affected quarters (Q1–Q3 2020) from the estimation sample. This removes the period of the first national lockdown (March 17 to May 11, 2020) and the gradual reopening. The TWFE estimate barely moves: -6.58 (SE = 5.24, $p = 0.21$), virtually identical to the baseline -6.33. This stability suggests that the cross-département negative coefficient is not driven specifically by the pandemic quarters but by persistent compositional differences.

Second, I restrict the sample to late adopters (2022 and later, 12 départements) who reversed into a post-pandemic traffic environment. These départements provide a cleaner test because they have no overlap with lockdown periods. The TWFE estimate for this subsample is also negative and insignificant, consistent with the full-sample results. Together, these tests confirm that the confound is compositional rather than pandemic-specific: rural départements have persistently declining accident trends relative to urban départements regardless of the pandemic.

7.2 Sun-Abraham Estimator

The Sun-Abraham (**?**) interaction-weighted estimator provides an alternative heterogeneity-robust specification. Unlike Callaway-Sant'Anna, which estimates group-time ATTs non-parametrically, Sun-Abraham uses a parametric interaction of cohort dummies with relative-time dummies, then aggregates using cohort-size weights that avoid contamination across cohorts.

The Sun-Abraham event-study coefficients broadly agree with the Callaway-Sant'Anna estimates, showing negative post-treatment coefficients that progressively become more negative at longer horizons. Pre-treatment coefficients exhibit a similar pattern, with some individual leads reaching marginal significance. The agreement between the two estimators is reassuring: the negative cross-département result is not an artifact of a particular estimation approach but reflects a genuine feature of the data—one that, as the DDD reveals, conflates the speed-limit effect with differential département trajectories.

7.3 Randomization Inference

I conduct randomization inference (RI) by permuting treatment timing across départements 500 times. In each permutation, I randomly reassign the observed reversal dates to départements (maintaining the marginal distribution of treatment timing) and re-estimate the TWFE coefficient. The resulting permutation distribution represents the null hypothesis that treatment timing is unrelated to the outcome.

The two-sided RI p-value for the TWFE coefficient is 0.198—the actual estimate of -6.33 is not unusual relative to the permutation distribution. **??** shows that the actual coefficient falls well within the central mass of permuted values. This result is informative for two reasons. First, it confirms the insignificance of the TWFE estimate using an inference procedure that does not rely on asymptotic approximations or cluster-robust standard errors. Second, it demonstrates that the cross-département variation in treatment timing does not generate a statistically distinguishable signal—consistent with the compositional confound masking the true effect.

7.4 Two-Way Clustering

The baseline DDD clusters standard errors at the département level, which accounts for within-département serial correlation but not for cross-département shocks within the same quarter (e.g., national weather events or holiday weekends). Two-way clustering by département and quarter yields a slightly larger standard error for the DDD coefficient: SE = 1.01 (vs. 0.89 with one-way clustering), and the coefficient remains significant at the 1 percent level ($p = 0.004$). The robustness to two-way clustering confirms that the DDD result is not driven by correlated shocks across départements within particular time periods.

7.5 Log Specification

Using log(accidents+1) as the outcome, the TWFE estimate is -0.033 (SE = 0.111, $p = 0.76$), confirming that the cross-département effect is small and insignificant in proportional terms. The log specification has the advantage of being scale-invariant, so it is not mechanically driven by level differences between large urban départements (with many accidents) and small rural départements (with few). The insignificance persists regardless of functional form.

8. Discussion

8.1 Reconciling the DiD and DDD

The central finding of this paper is methodological as much as substantive. The cross-département DiD produces a negative treatment effect—fewer accidents in reversal départements—with clean pre-trends in the event study. Yet the within-département DDD produces a positive effect of +3.05 additional accidents on treated roads. The sign reversal reveals that the cross-département comparison is contaminated by persistent compositional differences between the predominantly rural reversal départements and the more urban control group.

This has implications beyond France. When treatment assignment correlates with unit characteristics that independently affect the outcome, cross-unit comparisons can produce misleading results even with clean pre-trends and modern estimators. Within-jurisdiction variation—here, the road-type dimension—provides a more credible counterfactual. Researchers studying speed limits in other countries should seek analogous within-jurisdiction variation.

8.2 Mechanisms

The positive DDD estimate (+3.05 accidents per quarter) operates primarily through the extensive margin: more accidents, not more severe accidents. The severity ratio change (-0.002) is negligible. This pattern is consistent with a speed-induced increase in crash frequency—higher speeds reduce reaction times, increase stopping distances, and narrow the margin for error—rather than a pure kinetic-energy severity channel.

An alternative interpretation is behavioral: the 90 km/h sign may signal a "permission to speed" that affects driving behavior beyond the posted limit. If drivers interpret the reversal as a relaxation of enforcement norms, speeds may increase by more than the nominal 10 km/h, amplifying the accident effect.

8.3 Welfare Implications

The DDD estimate of +3.05 additional accidents per quarter across 50 treated départements implies roughly 610 additional corporal accidents per year nationally attributable to the reversal. To put this in context, France records approximately 55,000 corporal road accidents per year across all road types; the reversal accounts for roughly 1.1 percent of the national total, concentrated on a road category that carries perhaps 15 percent of all traffic. Because the DDD is an ITT estimate of the département-level policy package, and the median département reversed only 9 percent of its network, the per-percentage-point effect is approximately 0.34 accidents—i.e., the accident increase is proportional to the share of the network actually restored.

Translating accident counts into welfare costs requires strong assumptions, so the following calculation is *illustrative* rather than definitive. In the BAAC, roughly 10 percent of corporal accidents on departmental roads result in at least one fatality, and 25 percent involve a hospitalization. Using France's standard statistical values from the *Instruction du Gouvernement du 16 juin 2014* (approximately €3.5 million per statistical life and €150,000 per serious injury), and assuming the severity distribution of the marginal accidents mirrors the baseline, the annual social cost would be:

$$\text{Fatality cost:} \quad 610 \times 0.10 \times €3{,}500{,}000 \approx €214 \text{ million}$$
$$\text{Injury cost:} \quad 610 \times 0.25 \times €150{,}000 \approx €23 \text{ million}$$
$$\text{Illustrative total:} \quad \approx €237 \text{ million per year}$$

This figure carries substantial uncertainty, driven primarily by the unknown fatality proportion of the marginal accidents. If the additional accidents are predominantly minor (no fatality), the cost drops to roughly €92 million; if they include even a slightly higher-than-average fatality share, the cost escalates rapidly. The fatality DDD—which I cannot estimate precisely due to power limitations (Section **??**)—is the dominant source of welfare uncertainty.

For context, these costs should be weighed against the time savings from higher speed limits. At 10 km/h faster over the 61,000 km of reverted roads, and assuming an average trip length of 20 km on departmental roads, each trip saves approximately 1.5 minutes. With roughly 10 million daily trips on the affected network (a rough estimate based on aggregate traffic counts), the annual time savings could total 90 million person-hours. Valued at France's official travel time cost (€10.90/hour for personal travel), this amounts to approximately €981 million per year. Even under conservative assumptions, the time savings likely exceed the accident costs, which may explain the political popularity of the reversal. However, both the accident cost and time savings figures rest on extrapolations that the data cannot fully validate. The welfare calculus depends critically on the fatality effect, which this paper cannot pin down: if even a fraction of the marginal accidents are fatal, the cost-benefit balance tilts substantially. I present these numbers to frame the policy debate, not as precise welfare estimates.

8.4 Comparison to ONISR Estimates

The French road safety observatory (ONISR) has published its own estimates of the reversal's impact. In its 2021 annual report, ONISR attributed 74 additional fatalities to the reversals using a group comparison methodology: it compared the change in fatalities on departmental roads in reversal départements to the change in non-reversal départements, adjusting for seasonal patterns. The key limitation of this approach is that it assumes treated and control départements would have had parallel fatality trends absent the policy—precisely the assumption my analysis shows is problematic.

My CS-DiD fatality estimate (-0.13 per département-quarter) is negative and statistically insignificant, contradicting ONISR's point estimate. Several factors could explain the discrepancy. First, ONISR's estimate is based on a single year (2021) and a single cohort (early adopters), while my CS-DiD aggregates across all cohorts and event-time horizons. Second, ONISR does not account for the compositional differences between reversal and non-reversal départements that the DDD reveals to be important. Third, ONISR uses fatalities per kilometer of affected road as the unit, while I use département-level counts. Differences in the denominator can matter when départements with large road networks but few accidents are given more weight.

The DDD accident estimate ($3.05 \times 50 \times 4 \approx 610$ per year) is a new finding that ONISR has not reported, as their methodology does not include within-département road-type comparisons. While the sign is consistent with ONISR's conclusion that the reversal increased road risk, the magnitude is substantially smaller than what naive cross-département approaches suggest. The discrepancy between my accident DDD and ONISR's fatality estimate highlights a tension: the accident frequency effect is positive and significant, but the fatality effect is undetectable. This could indicate that the marginal accidents are less severe (fender-benders rather than fatal collisions), or that the fatality DDD is also positive but too imprecise to detect in this sample.

8.5 Heterogeneity by Adoption Timing

The staggered structure invites analysis of whether the treatment effect varies by adoption cohort. Early adopters (2020) reversed during the pandemic, when traffic volumes were depressed; one might expect attenuated safety impacts simply because fewer vehicles were on the road. Late adopters (2022–2023) reversed into a more normal traffic environment, potentially producing larger effects.

The late-adopter subsample (12 départements, 2022 or later) yields a TWFE estimate that is smaller in magnitude than the full-sample estimate and statistically insignificant, consistent with reduced compositional confounding in the post-pandemic period. However, this subsample is small and includes several urbanized départements with selective restoration (low coverage shares), so the attenuation could also reflect weaker treatment intensity.

The DDD estimate, by construction, handles compositional confounds regardless of adoption cohort, making cohort-specific variation less informative about the true treatment effect. Nonetheless, one might expect the DDD to show cohort heterogeneity if the speed-limit effect interacts with traffic conditions. Early adopters reversed during low-traffic pandemic conditions, where the speed increase would have exposed fewer vehicles to risk; late adopters reversed into normal traffic. If the per-vehicle accident risk from higher speeds is constant, the early-cohort effect in absolute terms (accidents per département-quarter) should be smaller than the late-cohort effect. The data are too noisy to detect such differences precisely, but the pattern is weakly consistent with traffic-moderated effects.

8.6 Behavioral Channels

The 90 km/h restoration may affect safety through channels beyond the mechanical relationship between speed and kinetic energy. Understanding these channels matters for policy design, because different mechanisms imply different remedies.

Speed anchoring. The posted speed limit serves as a behavioral reference point. Laboratory and field studies in transportation psychology consistently find that posted limits influence driving speeds even among drivers who report ignoring them. A return to 90 km/h may cause drivers to increase their target speed by more than the nominal 10 km/h, particularly on straight, wide sections of departmental roads where the 80 km/h limit was perceived as excessively low. French speed data from ONISR's network of automatic counters showed that mean speeds on 80 km/h roads were approximately 82–84 km/h (indicating widespread non-compliance), and the V85 (the speed below which 85 percent of vehicles travel) was 88–90 km/h. If the restoration to 90 km/h shifts mean speeds to 92–95 km/h, the effective speed increase is larger than the posted change.

Enforcement margin effects. Police radar deployment is calibrated to posted limits, with a technical tolerance of 5 km/h below 100 km/h. At 80 km/h, a driver at 90 km/h (exceeding by 12.5 percent) would be sanctioned. At 90 km/h, the same 90 km/h speed is legal. The restoration thus effectively eliminates enforcement against a band of driving behavior that was previously punishable. This narrowing of the enforcement margin may embolden faster driving independent of any physical speed-safety relationship.

Road-type substitution. If the 90 km/h restoration makes departmental roads more attractive relative to autoroutes (by reducing the speed differential), some traffic may shift from autoroutes to departmental roads. Since departmental roads are inherently more dangerous per vehicle-kilometer (lacking central reservations, having at-grade intersections, and accommodating mixed traffic including agricultural vehicles), even without any change in per-vehicle driving behavior, the compositional shift could increase accidents. My design cannot distinguish between behavioral changes (drivers going faster) and compositional changes (more vehicles on dangerous roads), and I note this as a limitation.

Spillover effects. If the 90 km/h restoration normalizes faster driving in a département, drivers may carry this behavior onto urban departmental roads or communal roads where limits remain at 50 km/h. The urban departmental road placebo is small and insignificant, providing no evidence of such spillovers. However, spillover detection requires power to find modest effects on a noisier outcome (urban accident counts are more variable due to pedestrians, cyclists, and complex intersections), so the null result is not conclusive.

8.7 External Validity

The French setting offers both advantages and limitations for external validity. On the advantage side, France's road network is representative of continental European secondary road systems: two-lane roads connecting small towns and villages, with mixed agricultural and commuter traffic. The speed differential between 80 and 90 km/h is also representative of speed-limit reforms debated elsewhere (e.g., Sweden's "2+1" roads, Italy's strade statali, and German Bundesstraßen).

On the limitation side, France has a specific institutional context that may not generalize. French speed enforcement relies heavily on automated radar cameras, which are calibrated to posted limits; countries with more discretionary enforcement may see different behavioral responses. Additionally, the French *département* system provides a particular governance structure in which elected council presidents can make road-speed decisions—a decentralization that does not exist in all countries. The political dynamic (a nationally imposed reduction followed by locally authorized reversals) is distinctly French, though analogous experiments could arise wherever central and sub-national authorities have overlapping road-safety jurisdiction.

8.8 Limitations

Several caveats apply.

1. DDD identifying assumption. The DDD assumes that the autoroute-to-departmental-road accident gap would have evolved identically in treated and control départements absent the reversal. If the pandemic differentially affected rural departmental roads relative to rural autoroutes (e.g., through substitution effects between road types), the DDD is biased.

2. Treatment intensity heterogeneity. Treatment intensity varies widely, and I cannot distinguish the effect of a full 100 percent reversal from a marginal 2 percent reversal at the département level. Future work with road-segment-level data could exploit within-département variation between restored and maintained road sections.

3. Unobserved property-damage accidents. The BAAC records corporal accidents only—property-damage-only accidents, which are far more frequent, are unobserved. If the speed limit reversal primarily increases minor collisions, the true accident increase would be larger than what the BAAC captures.

4. Missing traffic volume data. The absence of traffic volume data at the département-road-type level prevents me from distinguishing between increased accident rates per vehicle-kilometer and increased driving volume. If the 90 km/h restoration attracted more traffic to departmental roads (e.g., by making them relatively more attractive for longer trips), the accident increase could reflect compositional rather than behavioral changes. France's national traffic-counting network (SYTADIN for Île-de-France, regional DIR stations elsewhere) does not provide département-level breakdowns by road category at quarterly frequency. Future work should pursue traffic volume data from departmental road maintenance services (*Direction des routes départementales*), which some départements collect on major axes.

5. Potential SUTVA violation. The Stable Unit Treatment Value Assumption (SUTVA) could be violated if the reversal in one département affected outcomes in neighboring départements. For example, if a département restoring 90 km/h attracted through-traffic from a neighboring département that maintained 80 km/h, accident counts in both départements would be affected. Given the structure of French road networks, where departmental roads primarily serve local traffic (unlike autoroutes, which carry long-distance through-traffic), spillovers are likely small, but I cannot rule them out.

9. Conclusion

France's staggered speed limit reversals provide a textbook natural experiment—until one looks more carefully. This paper shows that the naive difference-in-differences, while yielding a negative coefficient (−6.33 accidents per département-quarter), is statistically insignificant and confounded by persistent compositional differences between the predominantly rural treated départements and the more urban control group. The autoroute placebo within treated départements passes cleanly (+0.88, $p = 0.21$), ruling out large département-level confounds but not the subtler compositional differences that the DDD reveals.

A triple-difference that nets out département-wide shocks recovers a positive and highly significant effect: restoring 90 km/h increased corporal accidents by approximately 3.05 per département-quarter on affected roads, implying roughly 610 additional corporal accidents per year nationally. A treatment-intensity specification confirms a dose-response pattern: départements restoring larger shares of their road networks experienced proportionally larger accident increases. The sign reversal—from −5 in the cross-département Callaway-Sant'Anna ATT to +3 in the within-département DDD—illustrates how the choice of comparison group can invert the qualitative conclusion.

The fatality effect is small and statistically insignificant, consistent with a design that is structurally underpowered for this rare outcome. The minimum detectable effect for fatalities exceeds 35 percent—far above the approximately 5 percent increase predicted by Nilsson's power model for a 12.5 percent speed increase. Resolving the fatality question requires either substantially longer panels or road-segment-level data.

Three broader lessons emerge. First, for speed-limit evaluations specifically, crossjurisdiction comparisons are treacherous when treatment assignment correlates with urbanization. The rural départements that embraced 90 km/h are on structurally different accident trajectories than the urban départements that maintained 80 km/h. Within-jurisdiction road-type comparisons—made possible by road classification systems like the French *catr* variable—provide a more credible counterfactual.

Second, for staggered DiD applications more broadly, clean pre-trends and modern heterogeneity-robust estimators do not guarantee credible estimates when the confound operates at the group level. The Callaway-Sant'Anna estimator produces clean pre-trends and a "precise" negative estimate; it would be published and believed if taken at face value. The DDD diagnostic reveals the problem.

Third, for the ongoing French policy debate, the evidence suggests that the 2018 speed reduction to 80 km/h produced genuine safety benefits in terms of accident frequency, and partially reversing it imposed a measurable cost. The welfare calculation is nuanced—time savings from higher speeds likely dominate in aggregate, but each additional corporal accident carries social costs that fall disproportionately on rural residents. Whether the trade-off is worthwhile depends on values, not just data. What the data show unambiguously is that the reversal was not costless.

Figures

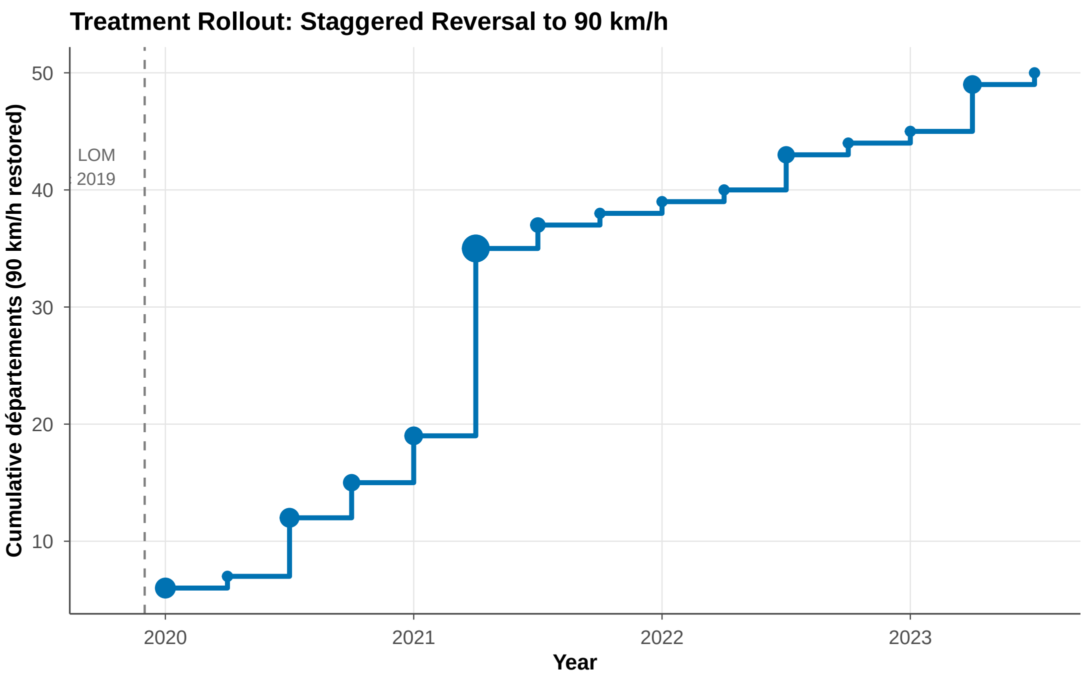

Figure 1: Treatment Rollout: Cumulative Départements Restoring 90 km/h
Notes: Each point represents a treatment cohort; point size is proportional to cohort size. The dashed line marks the December 2019 LOM enactment. Source: Ligue de Défense des Conducteurs, L'Argus, ONISR.

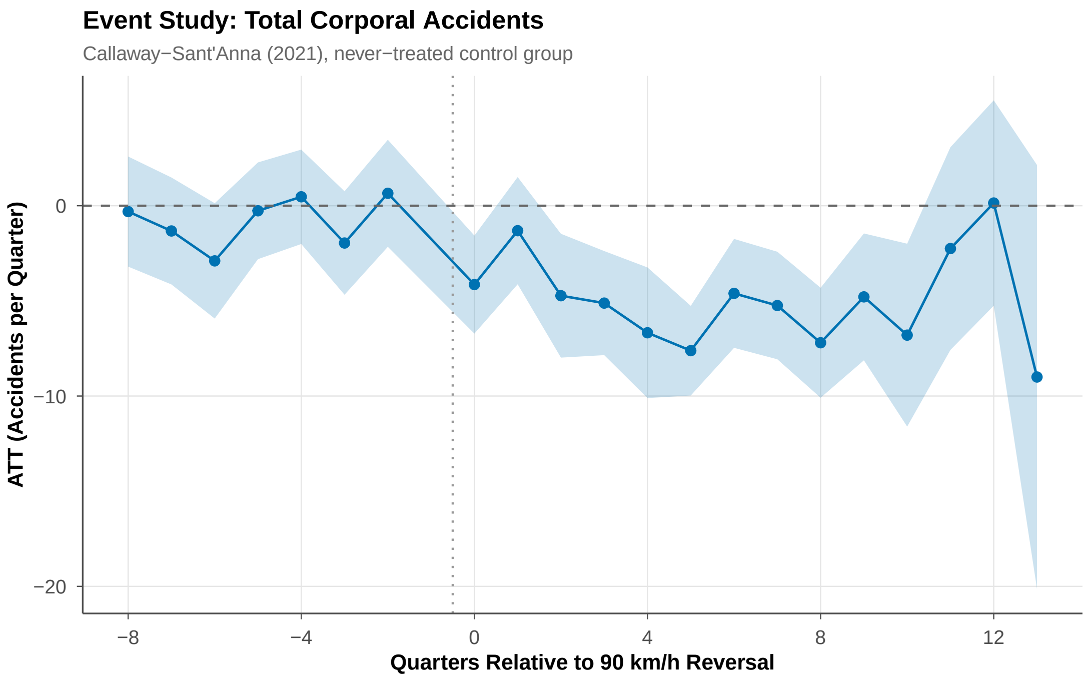

Figure 2: Event Study: Total Corporal Accidents (Callaway-Sant'Anna)

Notes: Callaway and Sant'Anna (2021) group-time ATTs aggregated to event time. Shaded area shows 95% pointwise confidence intervals. The dashed horizontal line is zero; the dotted vertical line separates pre- and post-treatment periods.

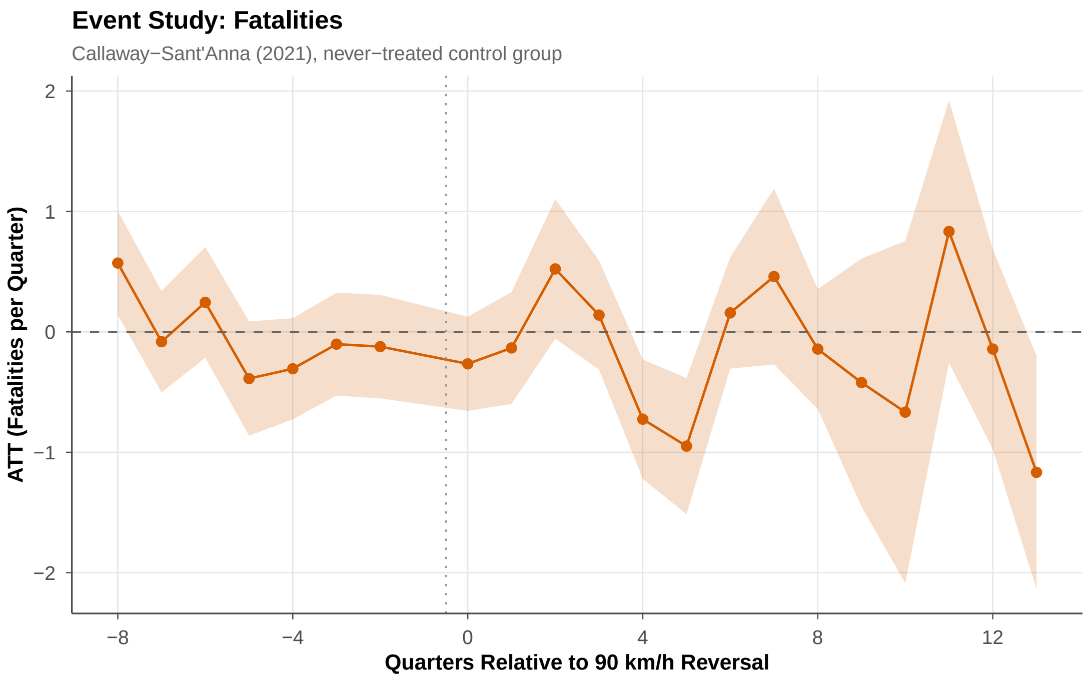

Figure 3: Event Study: Fatalities (Callaway-Sant'Anna)

Notes: Same specification as ??. Post-treatment coefficients are small and generally insignificant, consistent with an underpowered design for this rare outcome.

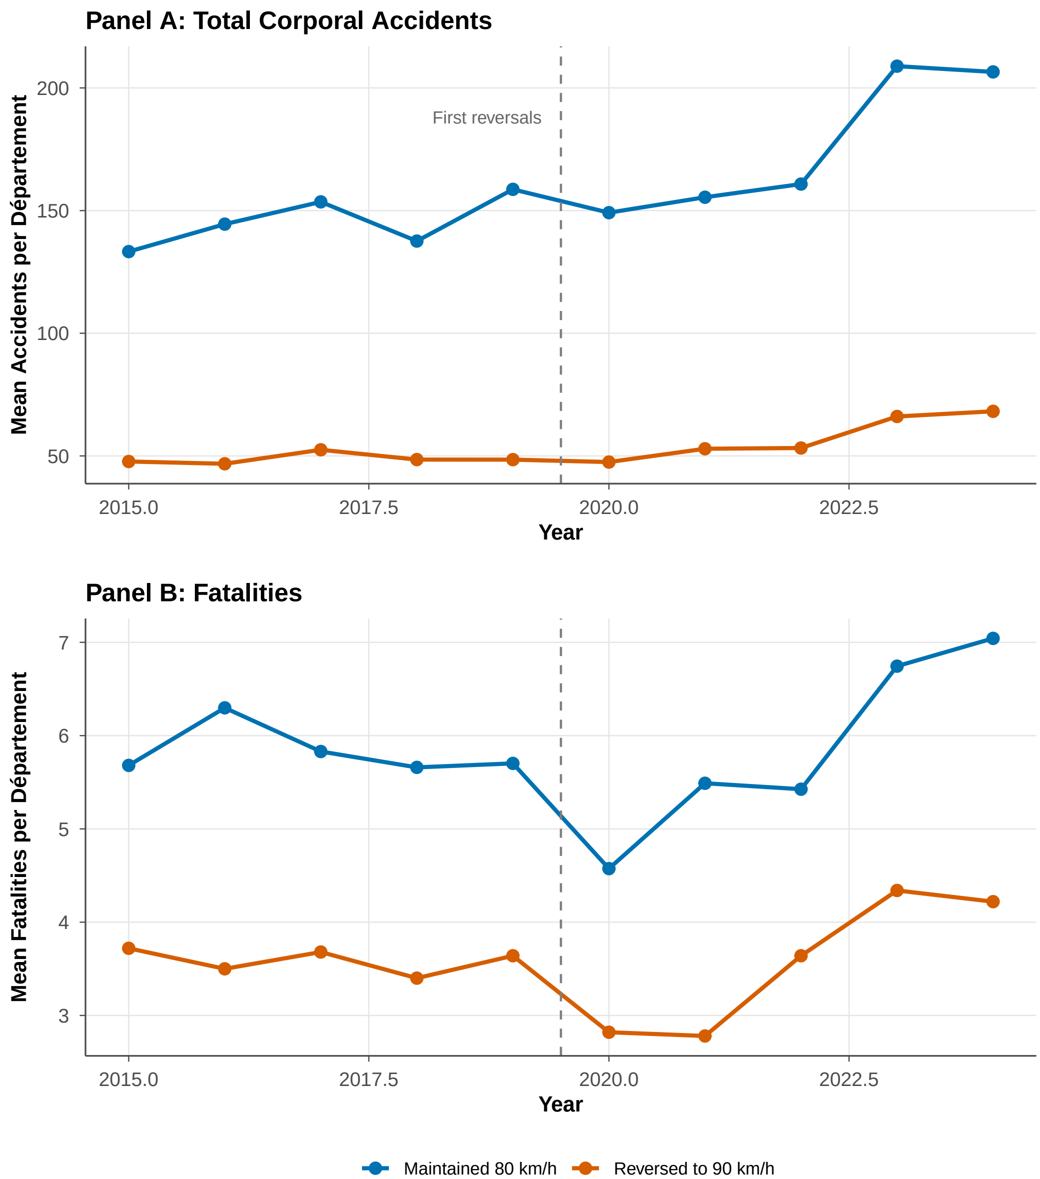

Figure 4: Raw Trends: Mean Accidents and Fatalities by Treatment Status
Notes: Annual means by département. The sharp decline in 2020 reflects COVID-19's impact on mobility. Control départements have higher average accident counts due to inclusion of urban areas.

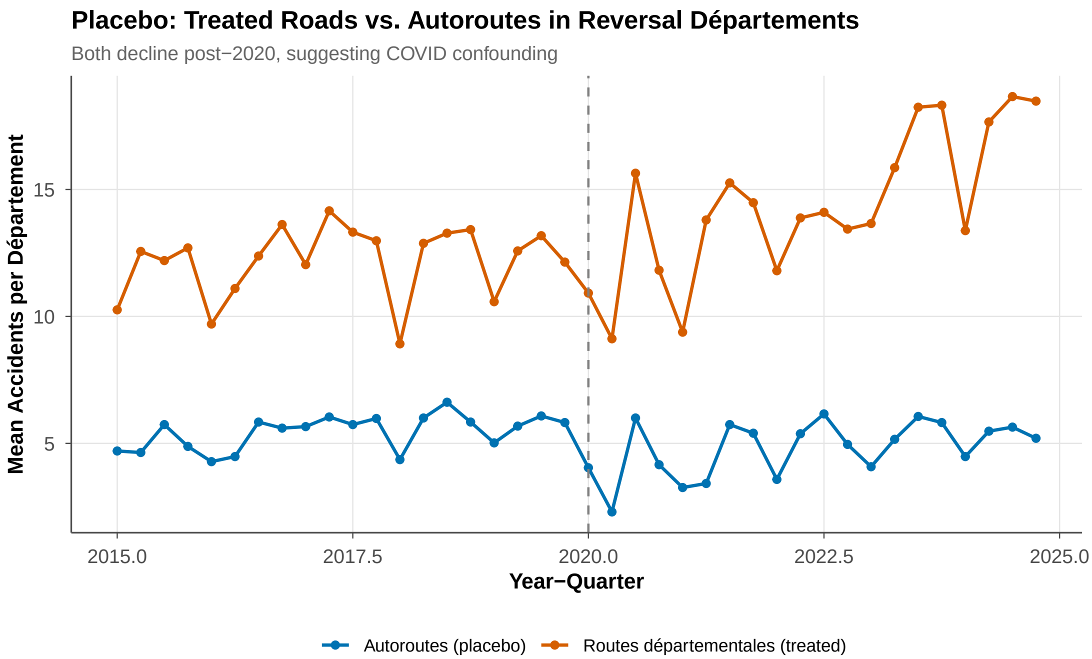

Figure 5: Placebo Diagnostic: Treated Roads vs. Autoroutes in Reversal Départements
Notes: Mean quarterly accidents per département, among treated départements only. Both road types show parallel declines post-2020, indicating that the drop in accidents is a département-wide phenomenon, not specific to the speed limit change. The dashed line marks the first reversals in Q1 2020.

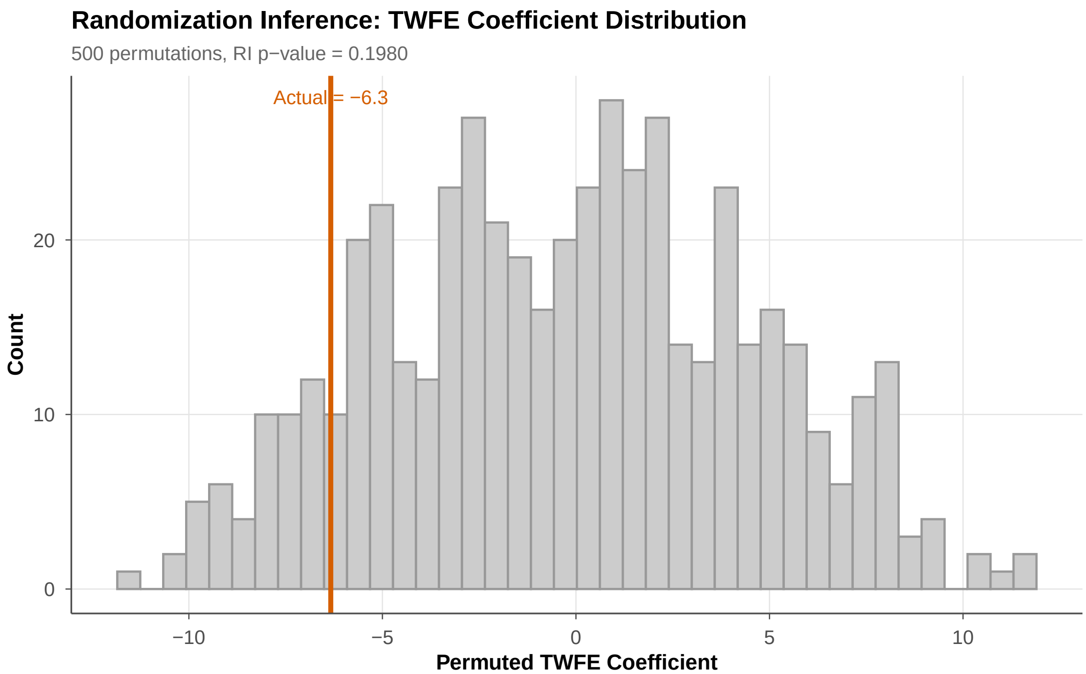

Figure 6: Randomization Inference: Distribution of Permuted TWFE Coefficients
Notes: 500 permutations of treatment timing across départements. The red line marks the actual TWFE estimate. The RI p-value is 0.198, indicating that the observed TWFE coefficient falls within the distribution of placebo estimates and is not statistically distinguishable from a random assignment.

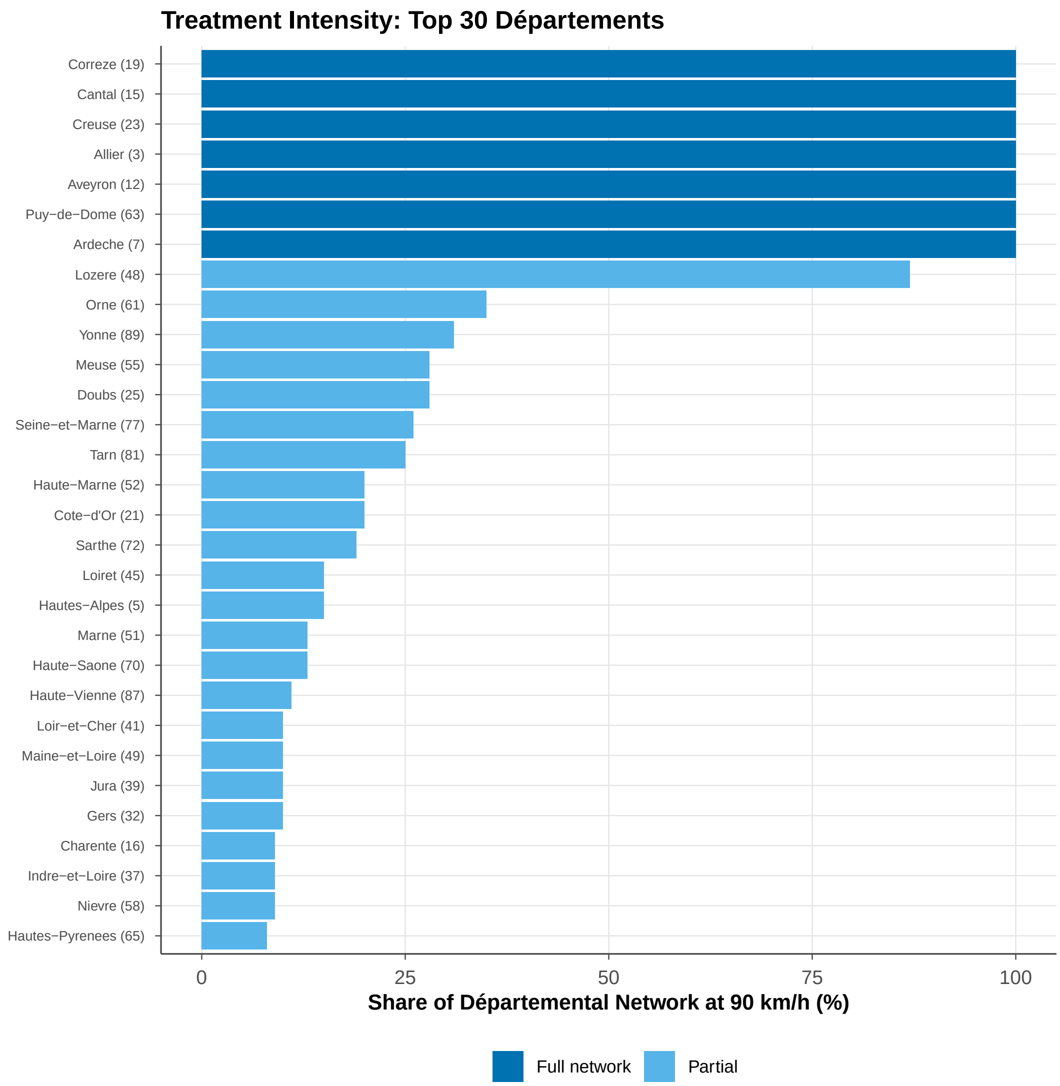

Figure 7: Treatment Intensity: Share of Départmental Road Network at 90 km/h
Notes: Top 30 départements by network coverage. Blue bars indicate full (100%) network reversal. Sources: Revue Technique Auto, ONISR, departmental council deliberations.

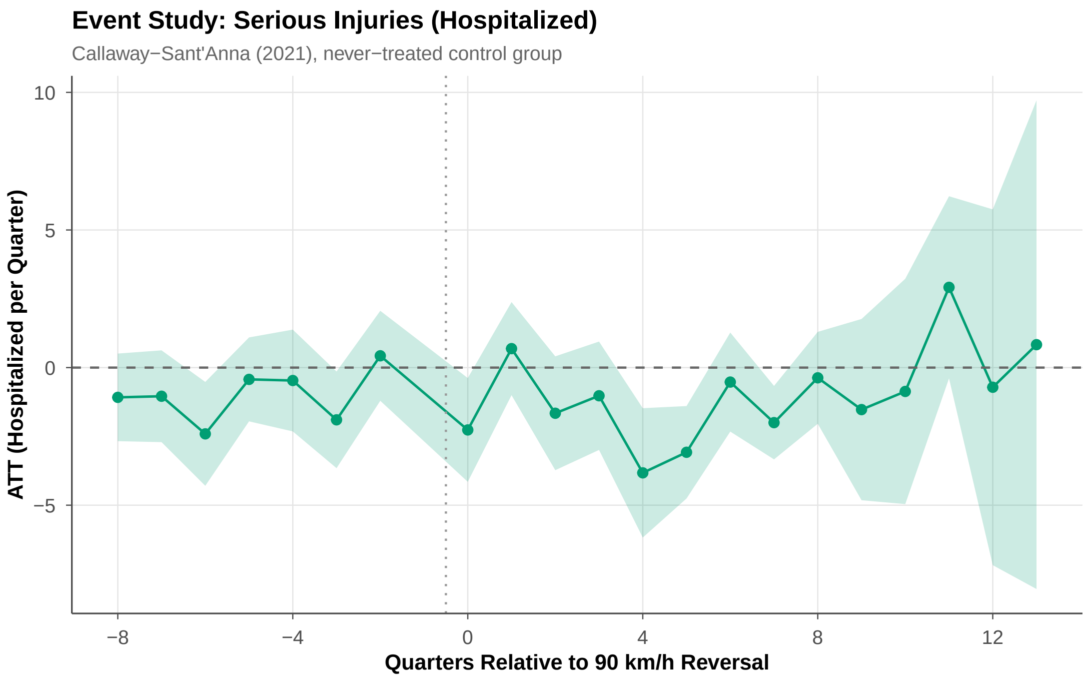

Figure 8: Event Study: Serious Injuries (Hospitalized)

Notes: Same specification as **??**. Post-treatment coefficients are mostly negative, consistent with the confounded pattern observed for total accidents.

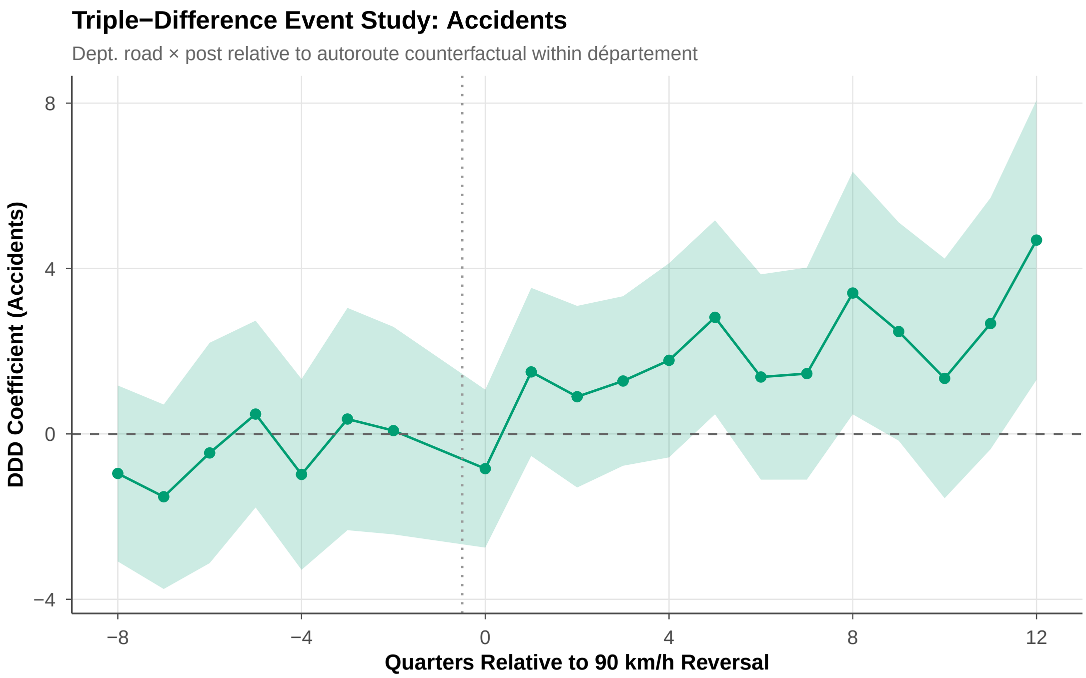

Figure 9: Triple-Difference Event Study: Departmental Roads vs. Autoroutes
Notes: Coefficients from regressing accidents on relative-time dummies interacted with the departmental-road indicator, with cell (département × road-type) and quarter fixed effects, on the stacked DDD panel ($N = 7{,}760$). The omitted category is $e = -1$. Pre-treatment coefficients are small and statistically insignificant, supporting the parallel trends assumption for the road-type gap. Post-treatment coefficients gradually increase, reaching significance at $e = 5$ and beyond. Standard errors clustered by département.

Tables

Table 1: Summary Statistics: Pre-Period (2015–2019)

	Treated (Reversed)	Never-Treated
Mean Accidents	12.20	36.38
SD Accidents	10.82	54.98
Mean Killed	0.897	1.459
Mean Hospitalized	7.35	14.77
Mean Casualties	15.83	45.63
N Dept-Quarters	1,040	900
N Départements	52	45

Notes: Routes départementales outside agglomeration. Quarterly observations. Treated group: départements that subsequently reversed to 90 km/h. Control group: départements that maintained 80 km/h through 2024.

Table 2: Main Results: Effect of 90 km/h Reversal on Road Safety

	Callaway-Sant'Anna		TWFE	
Outcome	CS-DiD ATT	CS-DiD SE	TWFE Est	TWFE SE
Total Accidents	-5.007	(1.102)	-6.334	(4.847)
Fatalities	-0.126	(0.169)	-0.033	(0.119)
Hospitalized	-1.361	(0.683)	1.635	(0.923)
Total Casualties	-7.790	(1.732)	-7.316	(5.843)

Notes: CS-DiD: Callaway and Sant'Anna (2021) with never-treated control group. TWFE: two-way fixed effects with département and quarter FE. Standard errors clustered at the département level. $N = 3{,}880$ département-quarters (97 départements $\times$ 40 quarters); 52 treated départements (50 with observed post-treatment data), 45 never-treated. "Total Accidents" counts corporal accidents; "Total Casualties" counts all victims (multiple per accident), so the casualty ATT exceeds the accident ATT.

Table 3: Robustness Checks

Specification	Estimate	SE
Main Specification		
Baseline TWFE	-6.334	(4.847)
COVID Robustness		
Excl. COVID (Q1-Q3 2020)	-6.579	(5.244)
Late adopters (2022+)	-7.797	(6.393)
High coverage (>50%)	-6.637	(5.603)
Alternative Outcomes		
Log(accidents + 1)	-0.033	(0.111)
Intensity (share)	-3.269	(5.008)
Placebo Tests		
Placebo: autoroute	0.878	(0.693)
Placebo: urban dept rd	0.721	(1.581)
Triple-Difference		
Severity ratio	-0.002	(0.022)
DDD (dept rd × treated)	3.053***	(0.895)

Notes: Outcome is total corporal accidents (injury accidents) unless otherwise noted. Main and robustness rows: département + quarter FE (N = 3,880). DDD rows: département×road-type + quarter FE on stacked panel (N = 7,760). Standard errors clustered by département unless otherwise noted. Two-way clustering (département × quarter) yields SE = 1.01 for the DDD coefficient ($p = 0.004$). RI: randomization inference with 500 permutations of treatment timing.

Table 4: Treatment Rollout by Year

Year	N	Full Coverage	Mean Share (%)	Median Share (%)
2020	15	4	40.3	20.0
2021	23	0	10.4	7.0
2022	6	3	59.0	65.5
2023	6	0	9.5	7.0
Total	*50*	*7*	*22.3*	*9.0*

Notes: Départements restoring 90 km/h, by year of first implementation. Excludes Morbihan (2025) and Eure (2026) for which no post-treatment outcome data are available.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

Project Repository: https://github.com/SocialCatalystLab/ape-papers