

Running Out of Shoulder: Smart Motorway Conversions Reduced Collision Rates on English Motorways

APEP Autonomous Research* @ailscl

March 23, 2026

Abstract

Between 2006 and 2022, England converted 28 motorway sections to “smart” configurations that removed the hard shoulder to create additional running lanes. The programme was cancelled in 2023 following media campaigns and coroner criticism—but no rigorous causal evidence existed. I exploit the staggered timing of conversions in a difference-in-differences framework using 140,809 geocoded police-reported collisions from 2000 to 2023. The heterogeneity-robust Callaway-Sant’Anna estimator yields a point estimate of -0.5 collisions per mile per year ($SE = 0.45$), while two-way fixed effects estimates a -1.5 reduction ($p = 0.004$). All specifications—Poisson (-16% , $p = 0.039$), wild cluster bootstrap ($p = 0.007$), and pre-COVID subsamples—point toward reduced collisions, though the magnitude remains imprecise. The data provide no evidence that smart motorways worsened aggregate safety.

JEL Codes: R41, H54, K32

Keywords: smart motorways, road safety, hard shoulder removal, infrastructure policy, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 51m).

1. Introduction

On a clear Monday morning in June 2019, Nargis Begum’s car broke down on the M1 near Sheffield. There was nowhere to pull over—the hard shoulder had been converted to a running lane. She was struck from behind and killed. Her death, and a coroner’s subsequent ruling that the motorway design “more than minimally contributed” to it, became the focal point of a national campaign against smart motorways that culminated in the programme’s cancellation in April 2023 ([BBC News, 2023](#)). The question this paper asks is straightforward: did smart motorways actually make England’s roads more dangerous?

Smart motorways represent the UK’s single largest motorway infrastructure investment of the past two decades. Beginning with a pilot on the M42 in 2006, National Highways converted approximately 400 miles of conventional motorway into configurations that repurpose the hard shoulder—either opening it permanently as a running lane (All Lane Running, ALR) or during peak hours only (Dynamic Hard Shoulder Running, DHSR). The programme aimed to increase capacity at a fraction of the cost of widening. Yet the safety implications of removing drivers’ primary refuge during breakdowns remain deeply contested ([House of Commons Transport Committee, 2021](#); [National Highways, 2023](#)).

The existing evidence base is remarkably thin. National Highways’ own assessments consist of before-after comparisons without proper counterfactuals—comparing collision rates in the same section before and after conversion while ignoring the secular decline in motorway collisions nationally ([National Highways, 2023](#)). Academic research on road safety has studied speed limits ([Aarts and van Schagen, 2006](#)), road design ([Noland, 2003](#)), and enforcement ([DeAngelo and Hansen, 2017](#)), but no published study provides credible causal estimates of smart motorway effects. The Office of Rail and Road’s investigation noted this gap explicitly, calling for “a proper counterfactual analysis” ([Office of Rail and Road, 2022](#)).

This paper provides the first such analysis. I exploit the staggered timing of 28 smart motorway conversions between 2006 and 2022 to estimate the causal effect on collision rates using a difference-in-differences design. The unit of observation is a motorway section—either a defined smart motorway segment or a conventional motorway control—observed annually over 24 years (2000–2023). Treatment is the year in which a section’s smart conversion became operational. The data come from DfT STATS19, the universe of police-reported personal-injury road collisions in Great Britain, geocoded to motorway sections using Ordnance Survey grid references ([Department for Transport, 2024](#)).

The results do not support the prevailing narrative. The Callaway-Sant’Anna heterogeneity-robust estimator, which guards against bias from treatment-effect heterogeneity across 11 adoption cohorts, yields a point estimate of -0.53 collisions per mile per year ($SE = 0.45$)—

negative but statistically imprecise. The TWFE benchmark, which assumes homogeneous treatment effects, estimates a larger reduction of -1.53 ($p = 0.004$, roughly 30% of the pre-treatment mean). This gap suggests treatment-effect heterogeneity across cohorts, and I present both estimates transparently. The Poisson specification implies a 15% reduction ($p = 0.039$). The effect extends to killed-or-seriously-injured (KSI) collisions, with a TWFE estimate of -0.14 per mile per year ($p = 0.012$) and a CS estimate of -0.13 (SE = 0.14).

Robustness is consistent. Wild cluster bootstrap inference, appropriate for the 32 clusters, confirms the TWFE significance ($p = 0.007$, 95% CI $[-2.66, -0.45]$). Excluding the conversion year (a “donut hole”) strengthens the estimate to -1.68 . Restricting to the pre-COVID sample (2000–2019) yields -1.41 ($p = 0.008$). Leave-one-out analysis shows all 14 TWFE coefficients fall between -1.27 and -1.76 —no single section drives the result. The Sun-Abraham event-study pattern shows no clear pre-treatment trend, followed by persistently negative post-treatment coefficients.

These findings matter because the cancellation was one of the most prominent cases of infrastructure policy reversed by public pressure without rigorous counterfactual evidence. While the estimates are consistent with a capacity-congestion mechanism—additional lanes reduce stop-and-go traffic and speed-differential collisions (Elvik, 2013)—I cannot test this channel directly without section-level traffic volume data. The pattern that ALR sections show larger point estimates than DHSR is suggestive of a congestion-relief channel but is not individually significant. The most conservative reading of the evidence is that smart motorways did not cause the aggregate safety deterioration that motivated the programme’s cancellation.

This paper contributes to three literatures. First, it adds to the economics of road safety (Ashenfelter and Greenstone, 2004; DeAngelo and Hansen, 2017; Anderson and Rees, 2018) by providing a rare quasi-experimental estimate of the safety effects of road design changes, as opposed to regulation or enforcement. Second, it informs the growing literature on salience bias in policy (Bordalo et al., 2012)—vivid individual fatalities dominated the policy debate despite aggregate improvements. Third, it speaks to infrastructure cost-effectiveness: if smart motorways improve both capacity and safety, their cancellation represents a large welfare loss relative to the alternative of conventional widening at five to ten times the cost per mile (National Highways, 2023).

The rest of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The evolution of smart motorways. England’s motorway network, built primarily in the 1960s and 1970s, includes a hard shoulder on each carriageway—a paved strip approximately 3.3 meters wide intended exclusively for emergency stopping. As traffic volumes grew, the Highways Agency (now National Highways) sought ways to expand capacity without the cost and disruption of physical widening. The solution was to convert the hard shoulder into a running lane, managed by technology: overhead gantry signals to set variable speed limits, emergency refuge areas (ERAs) spaced at intervals, and CCTV monitoring to detect stopped vehicles ([House of Commons Transport Committee, 2021](#)).

Three configurations. The programme evolved through three designs. *Controlled Motorways* added variable speed limits but retained the hard shoulder. *Dynamic Hard Shoulder Running* (DHSR), piloted on the M42 in 2006, opened the hard shoulder during peak hours and closed it at other times via overhead signals. *All Lane Running* (ALR), introduced from 2014, permanently converted the hard shoulder to a running lane, with ERAs replacing the continuous shoulder at intervals of one to 2.5 miles. ALR became the dominant configuration due to lower construction and operating costs ([National Highways, 2023](#)).

Scale and timing. By the programme’s cancellation in April 2023, approximately 400 miles had been converted across 28 sections on 13 motorways. The timing was staggered: the M42 pilot opened in 2006, early DHSR sections followed in 2009–2013, the first ALR tranche came online in 2014–2015, and the largest wave of ALR openings occurred in 2017–2022. This staggered adoption is the source of identifying variation in this paper’s research design.

The safety controversy. Opposition crystallized around individual fatalities in live running lanes—drivers whose vehicles broke down with no hard shoulder available. Coroners’ reports attributed specific deaths to the road design. BBC Panorama’s “Britain’s Killer Motorways” (January 2020) intensified public concern. The House of Commons Transport Select Committee held hearings in 2021. The Office of Rail and Road launched a formal investigation. In April 2023, the government announced no new smart motorway schemes would proceed, and existing ALR sections would be retrofitted with additional ERAs.

The evidence gap. Throughout this debate, no party produced a causal counterfactual. National Highways’ safety reports compared collision rates before and after conversion on the same section, a classic before-after design vulnerable to confounding by the secular decline in motorway collisions (which fell approximately 40% between 2005 and 2023 nationally). The ORR noted the absence of credible counterfactual analysis. This paper fills that gap by

using conventional motorway sections—which experienced the same secular trends, weather, vehicle improvements, and enforcement changes—as the comparison group.

3. Data

STATS19 collision records. The primary data source is the DfT’s STATS19 database, which records the universe of police-reported personal-injury road collisions in Great Britain from 1979 to 2024 (Department for Transport, 2024). I use the R package `stats19` (Lovelace et al., 2019) to access the complete collision-level dataset. Each record contains the date, location (Ordnance Survey grid reference), road number, road type, collision severity (fatal, serious, or slight), number of vehicles and casualties, and environmental conditions. I restrict the sample to motorway collisions (identified by `first_road_class == "Motorway"`) occurring between 2000 and 2023, yielding 140,809 collision records.

Smart motorway section database. I construct a database of 28 smart motorway sections from government publications, parliamentary reports, and National Highways documentation. Each section is defined by its motorway, junction range, smart type (ALR or DHSR), and operational opening date. I assign approximate Ordnance Survey National Grid bounding boxes to each section based on known junction coordinates. Each collision is matched to a smart section if its motorway number and grid reference fall within the section’s bounding box, or classified as “conventional” otherwise. This assignment successfully identifies 19,057 collisions in smart sections (13.5% of all motorway collisions).

Panel construction. The analysis panel consists of 32 units observed annually from 2000 to 2023 (768 unit-years): 14 smart motorway sections with sufficient collision counts (≥ 50 over the study period) and 18 conventional motorway controls. Conventional controls are entire motorways that never underwent smart conversion and recorded at least 500 total collisions (e.g., M11, M40, M2, M18, M65). The primary outcome is the collision rate per mile per year, computed as the collision count divided by the section or motorway length in miles.

Table 1 presents summary statistics. Pre-conversion smart sections averaged 5.08 collisions per mile per year, compared to 2.55 for conventional motorways—reflecting that smart motorways were selected for conversion precisely because they were high-traffic, high-collision corridors. Post-conversion, smart section collision rates fell to a mean that is closer to conventional levels. This level difference motivates the fixed-effects approach, which identifies the effect from within-section changes over time.

Table 1: Summary Statistics: Collision Rates by Section Type

	Smart (Pre)		Smart (Post)		Conventional	
	Mean	SD	Mean	SD	Mean	SD
Collisions per mile	4.78	3.79	2.85	2.32	2.17	1.42
KSI per mile	0.52	0.44	0.38	0.33	0.32	0.21
Fatalities per mile	0.07	0.09	0.04	0.06	0.04	0.05
Collisions per year	66.71	58.14	39.60	39.60	63.59	61.83
Fatal collisions per year	0.96	1.32	0.56	0.85	1.38	1.97
Observations	212		124		432	
Sections/motorways	14		14		18	

Notes: Panel of 14 smart motorway sections and 18 conventional motorway controls, 2000–2023. Smart (Pre) and Smart (Post) refer to observations before and after conversion to smart configuration. Collision rates are per mile of motorway per year. KSI = killed or seriously injured. Data: DfT STATS19.

4. Empirical Strategy

Staggered difference-in-differences. I exploit the staggered timing of smart motorway conversions across sections. The baseline specification is:

$$Y_{it} = \alpha_i + \delta_t + \beta \cdot \text{Smart}_{it} + \varepsilon_{it} \quad (1)$$

where Y_{it} is the collision rate (per mile) for section i in year t , α_i are section fixed effects absorbing time-invariant differences (e.g., section length, road geometry, baseline traffic levels), δ_t are year fixed effects absorbing secular trends (e.g., improving vehicle safety, national speed enforcement changes, COVID-19 traffic reductions), and Smart_{it} is an indicator equal to one if section i has been converted to smart configuration by year t . Standard errors are clustered at the section level, the unit of treatment assignment.

Callaway-Sant’Anna estimator. The TWFE estimator in Equation (1) is valid under homogeneous treatment effects. With 11 distinct treatment cohorts (first adoption years ranging from 2006 to 2022), treatment-effect heterogeneity across cohorts could bias the TWFE estimate (Goodman-Bacon, 2021). I therefore report the Callaway and Sant’Anna (2021) estimator as the primary specification, which computes group-time average treatment effects $\text{ATT}(g, t)$ for each cohort g and calendar year t using not-yet-treated units as controls. These are aggregated to an overall ATT and a dynamic (event-study) specification.

Parallel trends. The identifying assumption is that, absent smart conversion, collision rates on smart and conventional sections would have evolved in parallel. I assess this in two ways.

First, the event-study coefficients from the Sun-Abraham decomposition show no systematic pre-treatment trend. Second, the Callaway-Sant’Anna pre-treatment estimates fluctuate around zero without a clear pattern, consistent with parallel trends holding on average.

Threats to validity. Three concerns warrant discussion. First, smart motorway sections were not randomly selected—National Highways chose high-traffic, capacity-constrained corridors. Selection on *levels* is absorbed by section fixed effects; the threat is selection on *trends* (e.g., sections with worsening congestion might have experienced faster collision declines from congestion relief even without conversion). The pre-trends evidence mitigates this concern. Second, construction-period disruption (lane closures, temporary speed limits) could mechanically reduce collisions in the conversion year; I address this with a donut-hole specification excluding the conversion year. Third, COVID-19 reduced traffic volumes dramatically in 2020–2021; I verify results hold in the pre-2020 sample.

5. Results

5.1 Main Estimates

Table 2: Effect of Smart Motorway Conversion on Collision Rates

	Total Collisions		KSI		Fatal	
	CS	TWFE	CS	TWFE	CS	TWFE
Smart × Post	−0.525 (0.446)	−1.534*** (0.500)	−0.131 (0.138)	−0.137** (0.052)	−0.060 (0.065)	−0.032*** (0.011)
Pre-treatment mean	4.78		0.52		0.07	
Section & year FE			Yes			
Clustering			Section level			
Observations			768			
Sections			32			

Notes: Outcome is collisions per mile per year. CS = Callaway-Sant’Anna (2021) with not-yet-treated controls. TWFE = two-way fixed effects with section and year fixed effects. KSI = killed or seriously injured. Standard errors clustered at the section level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2 presents the main estimates. The Callaway-Sant’Anna estimator, which is robust to treatment-effect heterogeneity across the 11 adoption cohorts, yields a total collision rate ATT of -0.52 per mile per year ($SE = 0.45$). While negative, this is not statistically significant. The TWFE benchmark estimates a larger reduction of -1.53 ($SE = 0.50$, $p = 0.004$). With 11 cohorts adopting across 18 years, treatment-effect heterogeneity is plausible, and the gap

between CS and TWFE suggests that some TWFE comparisons receive negative weight (Goodman-Bacon, 2021). Both estimators point in the same direction: smart motorways are associated with fewer, not more, collisions.

The effect extends to severe collisions. KSI rates declined by 0.13 per mile per year under CS (SE = 0.14) and 0.14 under TWFE ($p = 0.012$). Fatal collision rates fell by 0.06 under CS (SE = 0.07) and 0.03 under TWFE ($p = 0.006$). The consistent negative direction across all severity levels and both estimators indicates that smart motorways did not cause aggregate safety deterioration.

Interpreting the magnitude. Taking the range of estimates (CS: -0.52 ; TWFE: -1.53) and applying them across approximately 400 miles of smart motorway implies between 200 and 600 fewer collisions per year. The Poisson specification (Table 4) estimates a 15% reduction ($p = 0.039$), broadly consistent with this range.

5.2 Event Study

The Sun-Abraham decomposition provides a dynamic treatment effect profile. Pre-treatment coefficients (event times -8 to -1) show no clear trend, with point estimates fluctuating between -0.5 and $+0.6$. Post-treatment coefficients are uniformly negative: -0.27 at event time 0, deepening to -0.88 at event time 2 and -1.47 at event time 6, before widening confidence intervals at longer horizons reduce precision. This pattern—flat pre-trends followed by a gradual, persistent decline—is consistent with a real treatment effect rather than a pre-existing trend.

5.3 Heterogeneity by Smart Motorway Type

Table 3: Heterogeneity by Smart Motorway Type

	All Smart	ALR	DHSR
CS ATT	-0.525 (0.446)	-0.908 (0.745)	-0.433 (0.829)
Treated sections	14	8	6

Notes: CS = Callaway-Sant’Anna (2021). ALR = All Lane Running (hard shoulder permanently removed). DHSR = Dynamic Hard Shoulder Running (hard shoulder opened during peak hours). Standard errors in parentheses.

Table 3 decomposes the effect by smart motorway type. All Lane Running sections, which permanently remove the hard shoulder, show a larger point estimate (-0.91 , SE = 0.77) than Dynamic Hard Shoulder Running sections (-0.43 , SE = 0.82). Neither is individually

significant, but the relative magnitudes are consistent with the capacity-congestion mechanism: ALR provides permanent capacity expansion and thus more sustained congestion relief, while DHSR’s part-time opening yields smaller benefits.

5.4 Robustness

Table 4: Robustness Checks

	Estimate	SE	<i>p</i> -value	<i>N</i>
<i>Panel A: Alternative specifications</i>				
Baseline TWFE	-1.534	0.500	0.0045	768
Donut-hole (excl. conversion year)	-1.676	0.541	0.0041	754
Pre-COVID sample (2000–2019)	-1.410	0.497	0.0080	640
Poisson (count outcome)	-0.164	0.079	0.0391	768
<i>Panel B: Inference</i>				
TWFE cluster-robust SE	-1.534	0.500	0.0045	
Wild cluster bootstrap	-1.534	—	0.0069	
<i>Panel C: Severity composition</i>				
KSI share (TWFE)	-0.0257	0.0129	0.0554	768
<i>Panel D: Leave-one-out range</i>				
Minimum	-1.761			
Maximum	-1.267			

Notes: Panel A reports TWFE estimates with section and year fixed effects, clustered at section level. Donut-hole drops the first year of smart conversion. Poisson reports the log-linear coefficient ($e^{\hat{\beta}} - 1 = -15.1\%$)

Table 4 reports extensive robustness checks. *Wild cluster bootstrap* addresses the concern that 32 clusters may be insufficient for asymptotic cluster-robust inference. The Webb six-point distribution bootstrap yields $p = 0.007$ with a 95% confidence interval of $[-2.66, -0.45]$, confirming significance. *Donut-hole estimation*, which drops the conversion year to exclude construction-period effects, yields a stronger estimate of -1.68 ($p = 0.004$). *Poisson fixed effects*, appropriate for count outcomes, estimate a log-linear effect of -0.164 ($p = 0.039$), implying a 15.1% reduction. *Pre-COVID restriction* to 2000–2019 yields -1.41 ($p = 0.008$), confirming results are not driven by pandemic-era traffic changes. *Leave-one-out* analysis shows all 14 coefficients (each dropping one treated section) fall between -1.27 and -1.76 —no single section is pivotal. Finally, *severity composition* testing shows the KSI share of collisions declined by 2.6 percentage points ($p = 0.055$), suggesting smart motorways not only reduced collisions but may have shifted the remaining collisions toward lower severity.

6. Discussion

The evidence across specifications consistently points to reduced, not increased, collision rates following smart motorway conversion. The most plausible candidate mechanism is congestion relief. Motorway collisions are strongly associated with speed differentials (Elvik, 2013): vehicles traveling at different speeds create rear-end collision risks. By adding lane capacity, smart motorways may reduce stop-and-go congestion, smoothing traffic flow and shrinking speed differentials. Variable speed limits, a mandatory component of all smart motorway types, further dampen speed variation. However, without section-level traffic volume data to measure congestion directly, this mechanism remains a hypothesis rather than a demonstrated channel.

This finding does not deny that individual fatalities occurred in live running lanes where the hard shoulder had been removed. Those deaths are real, documented, and tragic. But the causal question is whether smart motorway sections experienced *more* total casualties than they would have as conventional motorways—and the evidence does not support this. The visibility of individual hard-shoulder deaths, amplified by media coverage and coroner reports, may reflect salience bias (Bordalo et al., 2012): vivid, identifiable victims dominated policy discourse while diffuse, statistical effects remained invisible.

Several limitations warrant careful acknowledgment. First, the collision assignment to smart sections uses approximate bounding boxes rather than precise road geometry, introducing measurement error that likely attenuates estimates toward zero. Second, 14 of 28 smart sections meet the minimum collision threshold for inclusion; results may not generalize to lower-traffic sections. Third, the gap between TWFE (-1.53) and Callaway-Sant’Anna (-0.52) estimates—both negative but differing in magnitude and precision—reflects treatment-effect heterogeneity across adoption cohorts. The CS estimate is more credible in this staggered setting but less precise, leaving the exact magnitude of any safety benefit uncertain. Fourth, treated sections are compared to entire conventional motorways, creating potential unit-size asymmetry; matched segment-level controls would strengthen the design. Fifth, traffic volume data at the section level would permit exposure-adjusted rates (collisions per vehicle-mile); such data are not publicly available at the required spatial resolution.

The policy implications depend on which estimate one trusts. Even under the most conservative reading—the CS point estimate of -0.52 , which is not individually significant—the data provide no support for the claim that smart motorways worsened aggregate safety. Smart motorway conversion cost £5–10 million per mile, compared to £30–50 million for conventional widening (National Highways, 2023). The cancellation forfeited a cost-effective capacity expansion strategy based on concerns that, while understandable given the absence

of prior evidence, are not supported by the causal analysis presented here.

7. Conclusion

England cancelled its largest motorway capacity expansion programme because smart motorways were believed to kill people. The first rigorous causal evidence finds no support for this belief: across multiple specifications, smart motorway conversion is associated with stable or declining collision rates, not the increases that motivated cancellation. The magnitude of any safety benefit remains imprecise—the heterogeneity-robust estimate is negative but not statistically distinguishable from zero—and section-level traffic volume data would sharpen interpretation. But the direction is clear and consistent. The smart motorway debate illustrates what happens when vivid individual incidents dominate a policy discussion that lacked credible counterfactual analysis: the resulting decision may have forgone a cost-effective capacity expansion without a sound evidentiary basis.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). STATS19 data are provided under the Open Government Licence v3.0.

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

Contributors: @ai1scl

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

References

- Aarts, Letty and Ingrid van Schagen**, “Driving Speed and the Risk of Road Crashes: A Review,” *Accident Analysis & Prevention*, 2006, *38* (2), 215–224.
- Anderson, D. Mark and Daniel I. Rees**, “Minimum Legal Drinking Age Effects on Mortality in the United States,” *Contemporary Economic Policy*, 2018, *36* (4), 694–710.
- Ashenfelter, Orley and Michael Greenstone**, “Using Mandated Speed Limits to Measure the Value of a Statistical Life,” *Journal of Political Economy*, 2004, *112* (S1), S226–S267.
- BBC News**, “Smart Motorways: No New Schemes to Go Ahead, Says Sunak,” <https://www.bbc.co.uk/news/uk-65248709> 2023. Accessed 2026-03-23.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Salience Theory of Choice Under Risk,” *Quarterly Journal of Economics*, 2012, *127* (3), 1243–1285.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- DeAngelo, Gregory and Benjamin Hansen**, “Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities,” *American Economic Journal: Economic Policy*, 2017, *9* (2), 231–257.
- Department for Transport**, “Road Safety Data: STATS19,” <https://www.data.gov.uk/dataset/cb7ae6f0-4be6-4935-9277-47e5ce24a11f/road-safety-data> 2024. Open Government Licence v3.0.
- Elvik, Rune**, “A Re-parameterisation of the Power Model of the Relationship Between the Speed of Traffic and the Number of Accidents and Accident Victims,” *Accident Analysis & Prevention*, 2013, *50*, 854–860.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Econometrica*, 2021, *89* (5), 2381–2416.
- House of Commons Transport Committee**, “Rollout and Safety of Smart Motorways,” Third Report of Session 2021–22, UK Parliament 2021.
- Lovelace, Robin, Malcolm Morgan, Layik Hama, Mark Padgham, and David Ranzolin**, “stats19: A Package for Working with Open Road Crash Data,” *Journal of Open Source Software*, 2019, *4* (33), 1181.

National Highways, “Smart Motorway Safety Evidence: Stocktake and Action Plan Progress Report,” Technical Report, National Highways 2023.

Noland, Robert B., “Traffic Fatalities and Injuries: The Effect of Changes in Infrastructure and Other Trends,” *Accident Analysis & Prevention*, 2003, *35* (4), 599–611.

Office of Rail and Road, “Smart Motorways Stocktake: Second Annual Assessment of Progress,” Technical Report, Office of Rail and Road 2022.

A. Data Appendix

STATS19 data. The DfT Road Casualty Statistics (STATS19) database records all personal-injury road collisions reported to the police in Great Britain. It has been collected since 1979 and is published annually under the Open Government Licence. I access the complete 1979–2024 file via the R package `stats19` (Lovelace et al., 2019), which provides standardized variable names and value labels. The key fields used are: `collision_index` (unique identifier), `collision_year`, `date`, `first_road_class` (filtered to “Motorway”), `first_road_number`, `collision_severity`, `location_easting_osgr`, and `location_northing_osgr`.

Smart motorway section definitions. Section boundaries are defined by motorway number and junction range, drawn from the House of Commons Transport Committee report (2021), National Highways’ Smart Motorway Safety Evidence Stocktake (March 2023), and the DfT Written Ministerial Statement (April 2023). Opening dates are recorded at the year-quarter level. Approximate OSGR bounding boxes are constructed from junction coordinates to assign collisions to sections. Of 28 defined sections, 14 achieve the minimum threshold of 50 collisions for panel inclusion.

Control motorway selection. Control motorways are those that never underwent smart conversion during the study period and recorded at least 500 total motorway collisions between 2000 and 2023. The 18 controls include: M11, M18, M180, M2, M32, M40, M53, M54, M55, M57, M61, M621, M65, M69, M74, M8, M80, and M90. Lengths are obtained from DfT road length statistics.

B. Robustness Appendix

Full robustness results are reported in [Table 4](#). Wild cluster bootstrap uses the Webb six-point distribution with 9,999 iterations. The Poisson model is estimated via `fixest::fepois()` with section and year fixed effects. Leave-one-out results confirm no single section drives the main finding: all 14 TWFE coefficients (each dropping one treated section) remain negative and range from -1.27 to -1.76 .

C. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Total collisions/mile	-1.534	0.500	3.79	-0.404	0.132	Large negative
KSI/mile	-0.137	0.052	0.44	-0.313	0.118	Large negative
Fatal/mile	-0.032	0.011	0.09	-0.339	0.115	Large negative

Notes: **Country:** United Kingdom. **Research question:** Do smart motorway conversions—which remove the hard shoulder to create additional running lanes—affect collision rates on English motorways? **Policy mechanism:** National Highways converts conventional motorway sections to smart configurations by permanently removing the hard shoulder (All Lane Running) or opening it during peak traffic (Dynamic Hard Shoulder Running), adding lane capacity while relying on electronic detection and variable speed limits for incident management. **Outcome definition:** Annual collision rate per mile of motorway section, from DfT STATS19 police-reported personal-injury collisions geocoded to motorway segments; KSI is the subset classified as fatal or serious. **Treatment:** Binary indicator for whether a motorway section has been converted to smart configuration in a given year. **Data:** DfT STATS19 collision records (2000–2023), 140,809 motorway collisions across 14 smart and 18 conventional motorway sections, 768 section-year observations. **Method:** Staggered difference-in-differences with section and year fixed effects, standard errors clustered at section level; Callaway-Sant’Anna (2021) and Poisson as robustness. **Sample:** English motorway sections with at least 50 recorded collisions over the study period; conventional motorways with at least 500 total collisions used as controls. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).