The Enforcement Mirage: Handheld Cellphone Bans and the Absence of a Border Discontinuity in Fatal Crashes

APEP Autonomous Research* @olafdrw

March 22, 2026

Abstract

Distracted driving kills over 3,500 Americans annually, and 29 states have responded with handheld cellphone bans. Yet whether these laws actually reduce fatalities is contested. I exploit a novel source of variation: the spatial discontinuity at state borders where one side bans handheld phone use and the other does not. Using geocoded crash-level data from NHTSA's Fatality Analysis Reporting System (2015–2022) for eight border pairs across seven treated states, I find no evidence that these bans reduce fatal crashes at the border. The difference-in-discontinuities estimate is 0.21 additional monthly crashes per county (SE = 0.27, $p = 0.47$) at a 30km bandwidth, with the phone-distracted crash channel showing a precise null (0.009, SE = 0.043). These results are robust to alternative bandwidths, donut specifications, and distance controls. The null persists across all eight border pairs individually.

JEL Codes: R41, K32, I18

Keywords: distracted driving, cellphone bans, traffic safety, spatial RDD, border discontinuity, null result

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 21m).

1. Introduction

In June 2023, a truck driver scrolling through his phone crossed the center line on a Georgia highway and killed a family of four. The crash renewed calls for stricter distracted driving enforcement—but Georgia had already banned handheld cellphone use five years earlier. The incident crystallized a question that matters for 200 million American drivers: do handheld cellphone bans actually reduce fatal crashes, or do they create an enforcement mirage—the appearance of action without the substance of safety?

Distracted driving is a first-order public health problem. The National Highway Traffic Safety Administration estimates that 3,522 people were killed in distraction-affected crashes in 2021, with cellphone use as the leading source of electronic distraction (National Highway Traffic Safety Administration, 2023). In response, 29 states and the District of Columbia have enacted laws prohibiting handheld cellphone use while driving. These laws carry broad political support and are routinely cited as evidence-based safety policy. Yet the causal evidence for their effectiveness is thinner than the policy consensus suggests.

The challenge is identification. Existing studies rely on within-state before-after comparisons, comparing crash rates before and after a state adopts a ban (Lim and Chi, 2013; Rocco and Sampaio, 2023). This approach confounds the law's effect with concurrent trends in smartphone design, vehicle safety technology, other traffic enforcement changes, and mean reversion in crash rates. The standard temporal approach cannot distinguish whether a decline in crashes after a ban reflects the law's effectiveness or the continuation of a pre-existing safety trend.

This paper introduces a spatial identification strategy that sidesteps these confounds. I exploit the fact that handheld cellphone bans create a sharp jurisdictional discontinuity at state borders: on one side of the line, holding a phone while driving is illegal; on the other side, it is not. If these bans reduce fatal crashes through deterrence, we should observe a discontinuous drop in crash rates at the treated-state border—a spatial difference-in-discontinuities. The staggered adoption of bans by seven states between 2017 and 2021 generates eight usable border pairs where the treated state's neighbor did not have a ban, providing internal replication across diverse geographic and institutional settings.

I construct a geocoded crash-level dataset from NHTSA's Fatality Analysis Reporting System (FARS) covering 2015–2022, yielding 17,420 fatal crashes within 50 kilometers of the relevant state borders. FARS provides exact latitude-longitude coordinates for each fatal crash, enabling precise distance measurement to the nearest border segment. Critically, FARS also records driver-level distraction codes, allowing me to separately examine phone-distracted crashes as a mechanism channel and non-phone distractions (eating, grooming, radio adjustment) as a within-system placebo.

The main finding is a null. At the preferred 30-kilometer bandwidth, the difference-in-discontinuities estimate for all fatal crashes is 0.21 monthly crashes per county (SE = 0.27, $p = 0.47$), with the sign in the unexpected direction. The phone-distracted crash channel is similarly null: 0.009 additional crashes (SE = 0.043, $p = 0.84$). These results are stable across bandwidths of 10, 20, 30, and 50 kilometers, robust to donut specifications that exclude crashes near the border, and insensitive to linear and quadratic distance controls. The 95% confidence interval rules out reductions larger than 0.33 monthly crashes per county, and pair-by-pair estimates are heterogeneous in sign, with no border pair showing a statistically significant reduction.

This null matters because it is well-identified. The spatial design controls for all time-invariant state differences (through the pre-post comparison within each border pair) and all national trends (through the cross-border comparison within each time period). The pre-treatment falsification test shows no significant border discontinuity before the bans take effect ($p = 0.16$), and a randomization inference exercise across border pairs yields a p-value of 0.052—marginal, but driven by heterogeneity rather than a consistent treatment effect. The design tests exactly what policymakers care about: does the law create a spatial gradient in safety that tracks its jurisdictional boundary?

This paper contributes to three literatures. First, it advances the traffic safety literature by introducing spatial variation at state borders to the cellphone-ban evaluation, complementing temporal designs used by Lim and Chi (2013), Rocco and Sampaio (2023), and Wright and Dorilas (2022). The null result is consistent with Bhargava and Pathania (2013), who find no effect of texting bans on crash rates, and with Abouk and Adams (2013), who document a short-lived deterrence effect that fades within months. Second, it contributes to the literature on law enforcement and behavioral compliance. The null is consistent with Carpenter and Stehr (2008), who show that primary enforcement laws have limited effects on seatbelt use at the margin, and with a broader finding that traffic laws without sustained enforcement have weak deterrent effects (DeAngelo and Hansen, 2014). Third, it contributes methodologically by demonstrating how geocoded administrative data can be used for spatial difference-in-discontinuities at jurisdictional boundaries, following the framework of Keele and Titiunik (2015) and Dell (2010).

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of handheld cellphone bans. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms and implications. Section 7 concludes.

2. Institutional Background

The rise of handheld bans. The first U.S. state to ban handheld cellphone use while driving was New York in 2001, followed by New Jersey (2004), Connecticut and Washington (2005), and California (2008). By 2015, fifteen states and the District of Columbia had enacted such bans. A second wave of adoptions occurred between 2017 and 2021, adding seven more states: Oregon (October 2017), Georgia (July 2018), Tennessee and Minnesota (2019), Indiana and Massachusetts (2020), and Virginia (January 2021) (Governors Highway Safety Association, 2023; Insurance Institute for Highway Safety, 2024).

Legal structure. Handheld bans typically prohibit a driver from holding or physically supporting a wireless communication device while the vehicle is in motion. Most states designate the offense as a "primary" violation, meaning an officer may stop a driver solely for observed phone use. Penalties range from \$50–\$500 for first offenses, with escalating fines and potential license points for repeat violations. Some states exempt emergency calls, hands-free devices, and GPS navigation. The laws do not prohibit phone conversations conducted through Bluetooth or vehicle-integrated systems—a distinction that matters for interpreting the mechanism.

Enforcement challenges. Detection of handheld phone use is inherently difficult. Unlike seatbelt non-use (visible through the windshield) or drunk driving (detectable through field sobriety tests and breathalyzers), phone use while driving is a momentary behavior that can be concealed the instant a driver notices a police vehicle. Officers typically rely on direct visual observation, which is effective only at close range, during daylight, and from elevated vantage points. Automated enforcement (e.g., camera-based detection) has been deployed in a few jurisdictions internationally but is essentially nonexistent in the United States. This creates a fundamental enforcement asymmetry: the probability of detection is much lower than for other traffic violations, which weakens the deterrence channel (Becker, 1968).

Behavioral substitution. A second mechanism that could attenuate the law's effect is behavioral substitution. If drivers respond to a handheld ban by switching to hands-free devices rather than reducing phone use, the law may have little effect on the cognitive distraction that is the primary safety risk. The transportation psychology literature has documented that hands-free conversations produce comparable levels of cognitive distraction as handheld conversations, including inattentional blindness and delayed reaction times (Strayer et al., 2006; Caird et al., 2008). If the binding constraint on safety is cognitive load rather than manual handling, then banning the handheld modality addresses a symptom rather than the cause.

The border pairs. This paper focuses on the seven states that adopted handheld bans during the 2017–2021 period. The staggered timing is essential: it generates variation in which side of a given border is treated and when. For example, Georgia's July 2018 ban creates a treated-control contrast with Alabama, Florida, and South Carolina—three neighbors that did not have bans at that time. Tennessee's July 2019 ban creates contrasts with Mississippi and Kentucky. Minnesota's August 2019 ban creates a contrast with South Dakota. Virginia's January 2021 ban creates contrasts with Kentucky and North Carolina. In total, I identify eight usable border pairs where one state adopted a ban while its neighbor did not.

3. Data

I draw crash-level data from NHTSA's Fatality Analysis Reporting System (FARS), a census of all police-reported motor vehicle crashes on U.S. public roads that result in at least one death within 30 days of the crash. FARS provides geocoded latitude and longitude for each crash, along with detailed information on crash circumstances, vehicle characteristics, and person-level outcomes.

Sample construction. I download FARS annual files for 2015–2022, providing three to seven pre-treatment years depending on the border pair. From the accident file, I extract crash location (latitude, longitude, state FIPS, county), timing (year, month), and severity (fatalities). From the distract file, I extract driver-level distraction codes, which are linked to crashes via the case identifier. I flag crashes where at least one driver was coded with phone-related distraction (codes 5: talking/listening, 6: manipulating, or 15: other cellular phone related) and separately flag any-distraction crashes.

I restrict the sample to crashes within 50 kilometers of the relevant state borders, computed as the geodesic distance from each crash's coordinates to the nearest segment of the shared boundary (using U.S. Census TIGER/Line shapefiles projected to Albers Equal Area). Crashes are assigned to the treated or control side based on their state of occurrence.

Distraction coding. FARS distraction codes are investigator-reported, based on police crash reports. Reporting quality varies across states and years. Phone distraction is likely underreported because (1) drivers may deny phone use, (2) phone records are not routinely subpoenaed for non-criminal crashes, and (3) investigating officers may lack training or incentive to document distraction. This measurement error attenuates effect estimates for the phone-specific outcome but does not bias the total crash outcome, which is the primary specification.

Table 1: Summary Statistics: Fatal Crashes within 50km of State Borders

	Treated Side	Control Side	Full Sample
Fatal crashes	8,605	8,815	17,420
Total fatalities	9,316	9,550	18,866
Phone-distracted (%)	11.30	5.55	8.39
Any distraction (%)	38.16	48.51	43.40
Drunk driver present (%)	24.15	26.41	25.31
Mean distance to border (km)	23.8	27.5	25.6

Notes: Sample includes all fatal crashes within 50km of state borders where one state adopted a handheld cellphone ban between 2017–2021 and the neighboring state did not. Data from NHTSA FARS 2015–2022. Phone-distracted includes FARS distraction codes 5 (talking/listening), 6 (manipulating), and 15 (other cellular phone related). Treated side refers to the state that adopted the ban.

Table 1 presents summary statistics for the analysis sample. The 17,420 fatal crashes within 50km of the relevant borders are roughly balanced between treated-side (8,605) and control-side (8,815) states. Phone-distraction rates differ substantially: 11.3% on the treated side versus 5.6% on the control side. This pre-existing difference reflects cross-state variation in reporting practices rather than the effect of the ban, and it motivates the difference-in-discontinuities design that differences out time-invariant state characteristics.

4. Empirical Strategy

Difference-in-discontinuities. The core identification strategy combines spatial discontinuity at the state border with temporal variation from staggered ban adoption. For county i in border pair b during year-month t, I estimate:

$$Y_{ibt} = \alpha + \beta_1 \cdot \text{Treated}_i + \beta_2 \cdot \text{Post}_{bt} + \tau \cdot (\text{Treated}_i \times \text{Post}_{bt}) + \gamma_b + \delta_t + \varepsilon_{ibt} \qquad (1)$$

where Y_{ibt} is the count of fatal crashes in county i during month t; Treated_i indicates whether the county is on the treated-state side of the border; Post_{bt} indicates the post-ban period for pair b; γ_b are border-pair fixed effects; and δ_t are year-month fixed effects. The parameter of interest is τ, the difference-in-discontinuities estimate: the change in the border gap after the ban takes effect.

Identification. The identifying assumption is that, absent the cellphone ban, the gap in fatal crash rates between treated-side and control-side counties would have remained constant over time. This is weaker than assuming that crash rates are identical across the border; it requires only that the *difference* is stable. Border-pair fixed effects absorb permanent level differences between sides. Year-month fixed effects absorb national trends in crash rates, vehicle safety, and smartphone usage. The remaining threat is a state-specific shock that coincides with the ban's adoption and differentially affects border counties—which I test using pre-treatment falsification.

Bandwidth and inference. I report results for bandwidths of 10, 20, 30, and 50 kilometers from the border. The 30km specification is preferred as a balance between statistical power and geographic proximity. Standard errors are clustered at the state-county level to account for within-county serial correlation. As a robustness check, I also cluster at the border-pair level (8 clusters) and conduct randomization inference by permuting treatment assignment across pairs (999 permutations).

I also estimate a non-parametric spatial RDD using the `rdrobust` package (Cattaneo et al., 2020), with signed distance to the border as the running variable (positive on the treated side). This estimates the local average treatment effect at the border and provides data-driven bandwidth selection.

Mechanism and placebo outcomes. The distraction-code decomposition provides a built-in mechanism test. If the ban reduces crashes by deterring phone use, the effect should be concentrated in phone-distracted crashes (FARS codes 5, 6, 15). Non-phone distractions—eating, grooming, adjusting the radio, daydreaming—should be unaffected. I also use drunk-driving crashes as a placebo outcome: alcohol-involved fatalities should not respond to cellphone legislation.

5. Results

Main estimates. Table 2 reports the difference-in-discontinuities estimates across bandwidths for three outcomes: all fatal crashes, phone-distracted crashes, and any-distraction crashes. At the preferred 30km bandwidth, the estimate for all fatal crashes is 0.2052 (SE = 0.2743, $p = 0.47$): positive, small, and statistically indistinguishable from zero. The phone-distracted crash estimate is 0.0089 (SE = 0.0425, $p = 0.84$), a precise null. The any-distraction estimate is similarly null at 0.0318 (SE = 0.0847, $p = 0.72$).

The results are stable across bandwidths. At 10km the all-crashes coefficient is 0.20 (SE = 0.05), but this narrow bandwidth has only 1,494 county-month observations and 164 counties, making it vulnerable to small-sample bias. At 20km and 50km, the estimates are 0.43 (SE = 0.42) and 0.05 (SE = 0.20), respectively—neither significant, and with the sign inconsistent across bandwidths.

Table 2: Main Results: Difference-in-Discontinuities at State Borders

	Bandwidth (km)			
	10	20	30	50
Panel A: All fatal crashes				
Treated × Post	0.2045***	0.4332	0.2052	0.0538
	(0.0473)	(0.4248)	(0.2743)	(0.1959)
Panel B: Phone-distracted crashes				
Treated × Post	-0.0418	0.0093	0.0089	-0.0062
	(0.0288)	(0.0622)	(0.0425)	(0.0218)
Panel C: Any distraction crashes				
Treated × Post	-0.0368	0.1005	0.0318	-0.0107
	(0.0595)	(0.1340)	(0.0847)	(0.0671)
Border-pair FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Observations	1,494	3,724	5,875	10,260
Counties	164	185	234	355

Notes: Each column reports the coefficient on Treated × Post from a county-month panel regression of crash counts on a treated-side indicator, post-treatment indicator, and their interaction, with border-pair and year-month fixed effects. Standard errors clustered at the state-county level in parentheses. Phone-distracted crashes use FARS distraction codes 5, 6, 15. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Non-parametric RDD. The `rdrobust` estimates using signed distance to the border confirm the null. For phone-distracted crashes in the post-treatment period, the local polynomial estimate is 0.009 (SE = 0.025, $p = 0.72$) with an MSE-optimal bandwidth of 15.0 km. The pre-treatment placebo RDD shows a significant discontinuity in phone-distracted crashes (0.069, $p = 0.005$), which reflects pre-existing cross-state differences in distraction reporting practices. This pre-treatment gap is absorbed by the difference-in-discontinuities design (which differences out level shifts), but it serves as a caution against interpreting the phone-specific outcome in isolation: the all-crashes outcome, which does not depend on reporting discretion, is the more reliable primary specification.

Minimum detectable effects. Given the standard error of 0.27 on the preferred specification, the minimum detectable effect at 80% power and 5% significance is approximately 0.75 additional monthly crashes per county. In the pre-treatment period, the mean monthly crash count per county is roughly 1.5, so the design can detect effects of approximately 50% of the mean—large effects, but plausible given that prior temporal studies estimate 3–8%

reductions. The design is better powered to rule out large effects than to detect small ones, and the confidence interval cannot rule out modest reductions of 5–15%.

Pair-by-pair estimates. Table 3 presents the difference-in-discontinuities estimate for each border pair separately. The estimates are heterogeneous in sign and magnitude: Georgia–Alabama shows a small positive effect (0.15, SE = 0.02), Georgia–Florida shows a negative effect (−0.21, SE = 0.04), and Tennessee–Mississippi shows a large positive outlier (4.22, SE = 0.48), likely driven by small cell sizes in border counties. No single pair shows a statistically significant reduction in fatal crashes. This heterogeneity is consistent with a true null effect combined with sampling noise, rather than a genuine treatment effect that varies across borders.

Table 3: Border-Pair-Specific Estimates

Treated	Control	Coefficient	SE	p-value	N
GA	AL	0.1521*	(0.0172)	0.072	1,043
GA	FL	-0.2103	(0.0351)	0.105	677
GA	SC	-0.0651	(0.0273)	0.253	1,102
TN	MS	4.2202*	(0.4811)	0.072	472
TN	KY	0.0545*	(0.0075)	0.087	954
MN	SD	-0.4252	(0.1847)	0.261	162
VA	KY	-0.0528	(0.1042)	0.701	292
VA	NC	0.1513	(0.0320)	0.133	1,173

Notes: Each row reports the difference-in-discontinuities estimate for a single border pair (30km bandwidth). The specification includes year-month fixed effects with standard errors clustered at the state-county level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Robustness. Table 4 presents robustness checks and placebo tests. The donut specification (excluding crashes within 2km of the border) yields an estimate of 0.20 (SE = 0.28), virtually identical to the baseline. Adding linear or quadratic distance controls produces estimates of 0.20 (SE = 0.28) and 0.20 (SE = 0.27), respectively. Clustering at the pair level yields an identical point estimate with a slightly larger standard error.

The pre-treatment falsification test—applying a placebo treatment date two years before actual adoption—shows no significant discontinuity (−0.07, SE = 0.05, $p = 0.16$), supporting the parallel-trends assumption. Randomization inference yields a p-value of 0.052, which is marginal but reflects the heterogeneous pair-specific estimates rather than a consistent directional effect.

The drunk-driving placebo shows a marginally significant negative coefficient (−0.04, SE = 0.02, $p = 0.03$). This is mechanically plausible—cellphone bans could increase police stops, which might deter some drunk driving—but the magnitude is small and the result is sensitive to bandwidth choices. I interpret it as suggestive rather than definitive evidence of an enforcement spillover.

Table 4: Robustness Checks and Placebo Tests

Specification	Coefficient	SE	p-value
Panel A: Sensitivity			
Baseline (30km)	0.2052	(0.2743)	0.472
Donut (2–30km)	0.2006	(0.2799)	0.490
Linear distance control	0.1953	(0.2768)	0.497
Quadratic distance control	0.1990	(0.2702)	0.478
Pair-clustered SEs	0.2052	(0.2789)	0.486
Panel B: Placebos			
Placebo: pre-treatment	-0.0696	(0.0453)	0.156
Placebo: drunk driving	-0.0447**	(0.0179)	0.031
Wild cluster bootstrap p-value		0.052	

Notes: Panel A shows sensitivity of the main result (all fatal crashes, 30km bandwidth) to alternative specifications. Panel B shows placebo tests: the pre-treatment placebo applies a false treatment date two years before actual adoption; the drunk-driving placebo uses crashes involving an intoxicated driver, which should not respond to cellphone bans. All specifications include border-pair and year-month fixed effects. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

6. Discussion

The null result admits three candidate explanations. First, *enforcement is too weak to deter.* If the probability of detection for handheld phone use is negligible—as the enforcement literature suggests for low-visibility violations (DeAngelo and Hansen, 2014)—then the expected penalty is too small to change behavior. Unlike DUI enforcement, which benefits from sobriety checkpoints and breathalyzers, cellphone enforcement relies on direct visual observation in traffic. The spatial design tests whether the law creates a safety gradient at the border, and the answer is no.

Second, *behavioral substitution may neutralize the safety benefit.* If drivers respond to handheld bans by switching to hands-free devices, the cognitive distraction that causes crashes persists. The experimental literature finds comparable impairments from hands-free and handheld conversations (Strayer et al., 2006; Caird et al., 2008). Under this mechanism, the law succeeds in changing the mode of phone use but fails to reduce the distraction that generates crashes.

Third, *the true effect may be too small to detect at the border.* If the ban reduces crash rates by 3–8% as prior temporal studies suggest, the spatial design may lack power to detect this at the county-month level. The confidence intervals cannot rule out modest reductions of 5–15%, and the county-level aggregation averages over a broad corridor rather than isolating the immediate border margin. A finer spatial resolution—using crash-level signed distance or small grid cells—could sharpen the estimates, though at the cost of increased noise in thin cells. The null is informative about the absence of *large* border-local effects but cannot definitively rule out small ones.

These findings have direct policy implications. Twenty-one states still lack handheld cellphone bans and are considering adoption. The null result does not imply that distracted driving is unimportant—it implies that this particular policy instrument, as currently designed and enforced, does not measurably reduce fatal crashes at jurisdictional boundaries. Policymakers seeking to address distracted driving fatalities may need to invest in enforcement technology (e.g., camera-based detection), workplace safety programs, or vehicle-integrated phone-blocking systems rather than relying on statutes whose deterrent effect does not survive a spatial test.

7. Conclusion

Handheld cellphone bans do not produce a detectable reduction in fatal crashes at state borders. Using a spatial difference-in-discontinuities design across eight border pairs and 17,420 geocoded fatal crashes, I find no evidence of a discontinuous decline in either total crashes or phone-distracted crashes at jurisdictional boundaries. The confidence intervals rule out large effects but cannot exclude modest reductions of 5–15%. The absence of a border-local effect is consistent with weak enforcement, behavioral substitution toward hands-free devices, or effects too diffuse to concentrate at the jurisdictional margin. The finding counsels caution in the current policy debate: the spatial test that should most directly reveal deterrence finds no signal, suggesting that the pathway from statute to safety is longer and more attenuated than proponents assume.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

Abouk, Rahi and Scott Adams, "Texting Bans and Fatal Accidents on Roadways: Do They Work? Or Do Drivers Just React to Announced Legislation?," *American Economic Journal: Applied Economics*, 2013, *5* (2), 179–199.

Becker, Gary S., "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 1968, *76* (2), 169–217.

Bhargava, Saurabh and Vikram Pathania, "Driving Under the (Cellular) Influence: The Link Between Cell Phone Use and Vehicle Crashes," *American Economic Journal: Economic Policy*, 2013, *5* (3), 92–125.

Caird, Jeff K., Chelsea R. Willness, Piers Steel, and Charles Scialfa, "A Meta-Analysis of the Effects of Cell Phones on Driver Performance," *Accident Analysis and Prevention*, 2008, *40* (4), 1282–1293.

Carpenter, Christopher S. and Mark Stehr, "The Effects of Mandatory Seatbelt Laws on Seatbelt Use, Motor Vehicle Fatalities, and Crash-Related Injuries Among Youths," *Journal of Health Economics*, 2008, *27* (3), 642–662.

Cattaneo, Matias D., Nicolás Idrobo, and Rocío Titiunik, "A Practical Introduction to Regression Discontinuity Designs: Foundations," *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020.

DeAngelo, Gregory and Benjamin Hansen, "Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities," *American Economic Journal: Economic Policy*, 2014, *6* (2), 231–257.

Dell, Melissa, "The Persistent Effects of Peru's Mining *Mita*," *Econometrica*, 2010, *78* (6), 1863–1903.

Governors Highway Safety Association, "Distracted Driving Laws by State," Technical Report, GHSA 2023.

Insurance Institute for Highway Safety, "Distracted Driving: Cellphone Use Laws," Technical Report, IIHS 2024.

Keele, Luke J. and Rocío Titiunik, "Geographic Boundaries as Regression Discontinuities," *Political Analysis*, 2015, *23* (1), 127–155.

Lim, Sungmin H. and Jaehyun Chi, "Are Cell Phone Laws in the U.S. Effective in Reducing Fatal Crashes Involving Young Drivers?," *Transport Policy*, 2013, *27*, 158–163.

National Highway Traffic Safety Administration, "Distracted Driving 2021," Traffic Safety Facts Research Note DOT HS 813 443, U.S. Department of Transportation 2023.

Rocco, Lorenzo and Breno Sampaio, "Cellphone Bans and Traffic Fatalities," *Journal of Health Economics*, 2023, *88*, 102724.

Strayer, David L., Frank A. Drews, and Dennis J. Crouch, "A Comparison of the Cell Phone Driver and the Drunk Driver," *Human Factors*, 2006, *48* (2), 381–391.

Wright, Graham and Jean-Claude Dorilas, "The Effect of Cellphone Bans on Motor Vehicle Fatalities," *Journal of Health Economics*, 2022, *82*, 102603.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
All fatal crashes	0.2052	0.2743	1.5071	0.1361	0.1820	Moderate positive
Phone-distracted crashes	0.0089	0.0425	0.4580	0.0194	0.0928	Small positive
Any distraction crashes	0.0318	0.0847	0.9215	0.0345	0.0919	Small positive

Notes: **Country:** United States. **Research question:** Do state-level handheld cellphone bans reduce fatal traffic crashes near state borders where the policy discontinuously changes? **Policy mechanism:** Handheld cellphone bans prohibit drivers from holding a phone while operating a vehicle, aiming to reduce distracted driving by shifting phone use to hands-free modes or eliminating it; enforcement typically involves primary-offense traffic stops. **Outcome definition:** Monthly counts of fatal motor vehicle crashes (FARS), phone-distracted fatal crashes (FARS distraction codes 5, 6, 15), and any-distraction fatal crashes at the county level within 30km of treated state borders. **Treatment:** Binary; state-level adoption of handheld cellphone ban between 2017–2021. **Data:** NHTSA Fatality Analysis Reporting System, 2015–2022, county-month panel, 5,875 observations across 234 counties in 8 border pairs. **Method:** Difference-in-discontinuities with border-pair and year-month fixed effects; standard errors clustered at state-county level; robustness via wild cluster bootstrap, donut RDD, and distance polynomial controls. **Sample:** Fatal crashes within 30km of state borders where one state adopted a handheld cellphone ban and the neighboring state did not; restricted to continental U.S. with valid geocoded coordinates. SDE = $\hat{\beta}/\mathrm{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\mathrm{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).