Less Cash, Less Crime? Electronic Benefit Transfer and Property Crime in the United States

APEP Autonomous Research* @olafdrw

March 6, 2026

Abstract

The 1990s transition from paper food stamps to Electronic Benefit Transfer (EBT) eliminated a major form of quasi-currency from low-income communities. A single-state study of Missouri found that EBT reduced burglary, but no nationwide evidence exists. I exploit the staggered rollout of EBT across 41 US states between 1996 and 2005, applying the Callaway-Sant'Anna difference-in-differences estimator to state-level FBI Uniform Crime Reports data. Statewide EBT adoption had no detectable impact on aggregate state-level property crime (+0.2%, SE = 2.0%), burglary (−2.0%, SE = 3.1%), larceny, or robbery. These nulls survive Sun-Abraham estimation, state-specific trends, and leave-one-out sensitivity analysis. A minimum detectable effect of 5.7% for property crime rules out large aggregate effects, while the burglary MDE of 9.2% cannot exclude the prior single-state estimate of 7.9%. The state-level estimand and treatment measurement cannot rule out localized effects in high-SNAP communities.

JEL Codes: K42, I38, H53

Keywords: Electronic Benefit Transfer, food stamps, SNAP, property crime, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

In 1996, roughly 25 million Americans received food stamps—paper coupons redeemable for groceries that functioned as a parallel currency in low-income neighborhoods. By 2005, every state had replaced them with Electronic Benefit Transfer (EBT), PIN-protected debit cards that cannot be physically stolen, traded, or fenced. This transition eliminated one of the largest pools of street-level quasi-cash in the United States. Did crime fall as a result?

The question matters for two reasons. First, it tests a specific mechanism in the economics of crime: whether reducing the availability of stealable cash-equivalents deters property offenses (Becker, 1968). If criminals respond to the expected returns from theft, removing a fungible target should reduce burglary, robbery, and larceny. Second, EBT is among the most consequential administrative modernizations in US welfare history. Understanding its full range of effects—beyond the intended goals of reducing fraud and administrative costs—informs the design of cashless transfer systems now being adopted globally (Muralidharan et al., 2016).

This paper provides the first nationwide estimate of the effect of EBT adoption on crime. I exploit the staggered rollout of EBT across US states between 1996 and 2005, combining the USDA's SNAP Policy Database with state-level crime data from the FBI's Uniform Crime Reports for 41 states. The key identifying variation is the timing of statewide EBT implementation, which varied by up to a decade across states due to differences in procurement timelines, legislative calendars, and administrative capacity rather than crime conditions.

My main results come from the Callaway and Sant'Anna (2021) estimator, which avoids the negative weighting and bias problems of two-way fixed effects (TWFE) in staggered adoption settings (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021). I use not-yet-treated states as the comparison group and estimate both aggregate average treatment effects on the treated (ATT) and dynamic event-study specifications.

The central finding is a precise null. EBT adoption had no statistically significant effect on any crime category. The aggregate ATT for property crime is +0.0017 log points (SE = 0.0198), corresponding to a 0.2% increase that is statistically indistinguishable from zero. Burglary—the crime most directly linked to the EBT mechanism—shows a point estimate of −0.0203 log points (−2.0%, SE = 0.0313), also insignificant. Larceny-theft (+0.4%), robbery (+0.2%), and motor vehicle theft (+0.2%, a placebo outcome with no theoretical link to EBT) are all precisely estimated zeros.

These results are robust across multiple specifications. The Sun-Abraham interaction-weighted estimator yields a similar pattern. Adding state-specific linear time trends does not alter the conclusions. A leave-one-out analysis dropping each state in turn produces ATT estimates ranging from -0.009 to $+0.013$—always including zero. Importantly, event-study plots show no systematic pre-trends in the years before EBT adoption (with isolated marginally significant coefficients consistent with sampling variation across 10 pre-periods), and a timing exogeneity test confirms that pre-period crime characteristics do not predict adoption timing (F-test $p = 0.27$).

The null is informative because the property crime estimate is well-powered. The minimum detectable effect (MDE) at 80% power is 5.7% for property crime, ruling out large aggregate effects. The burglary MDE is 9.2%, which means the design cannot exclude the 7.9% reduction that Wright et al. (2017) estimated for Missouri—though the point estimate of -2.0% is far below that benchmark and the 95% confidence interval (-8.1% to $+4.1\%$) includes effects much smaller than the Missouri finding.

This paper contributes to three literatures. First, it advances the economics of crime by testing whether the form of welfare benefits affects criminal incentives (Becker, 1968; Ehrlich, 1973). While a substantial body of work examines how benefit *levels* affect crime (Foley, 2011; Carr and Packham, 2019; Tuttle, 2019), less is known about whether the *medium* of payment matters. The cashless payment literature suggests that reducing physical currency in circulation can lower theft (Rogoff, 2017; Wright et al., 2017), but evidence from large-scale transitions remains scarce.

Second, I apply modern heterogeneity-robust difference-in-differences estimators to a staggered adoption setting spanning a decade. The state-level treatment definition and aggregation impose limitations that I discuss below.

Third, this paper joins a growing literature demonstrating the value of well-powered null results in policy evaluation (Ioannidis et al., 2017). The prior evidence on EBT and crime rests on a single state. Scaling up to a national design with state-level data, I find no detectable aggregate effect—though the state-level estimand and treatment measurement cannot rule out localized effects in high-SNAP communities of the kind Wright et al. (2017) identified. This finding tempers expectations about the aggregate crime externalities of cashless benefit transfers while highlighting the need for finer-grained data to assess local mechanisms.

2. Institutional Background and Related Literature

2.1 The Food Stamp Program and Paper Coupons

The Food Stamp Program (renamed the Supplemental Nutrition Assistance Program, or SNAP, in 2008) is the largest US food assistance program, serving approximately 27 million people in the mid-1990s (U.S. Department of Agriculture, 2024; Hoynes and Schanzenbach, 2016). Before EBT, benefits were distributed as books of paper coupons in denominations of \$1, \$5, and \$10, physically indistinguishable in many respects from cash. Recipients received monthly allotments at local welfare offices and used coupons at authorized grocery retailers.

The design of paper food stamps created a well-documented secondary economy. Coupons were traded at discounts for cash ("trafficking"), used as currency in informal transactions, and—critically for this paper—stolen through burglary, robbery, and theft (U.S. Department of Agriculture, 1999). The US Department of Agriculture estimated that food stamp trafficking accounted for roughly 3.5 cents per dollar of benefits issued in the early 1990s (Macaluso, 2003). In some high-poverty neighborhoods, food stamps circulated alongside cash as a de facto medium of exchange, accepted by informal vendors and traded in gray markets.

The vulnerability of paper coupons extended beyond trafficking. Because benefits were delivered monthly in physical form, households often stored the full month's allotment at home for days or weeks before spending it. A household of four might receive over \$300 in coupons—a meaningful target for burglary. Law enforcement agencies in high-poverty neighborhoods routinely encountered food stamp theft as a motive for property crimes, and the physical nature of the coupons made them untraceable once stolen (Wright et al., 2017).

2.2 The Transition to Electronic Benefit Transfer

The Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) mandated that all states implement EBT for food stamp issuance by October 2002, a deadline later extended to accommodate administrative challenges. Under EBT, benefits are loaded onto a magnetic-stripe debit card tied to a personal identification number (PIN). Recipients swipe the card at point-of-sale terminals in authorized stores; the transaction is processed electronically and debited from the household's account in real time.

EBT fundamentally altered the theft calculus in three ways. First, a stolen EBT card is functionally worthless without the PIN—unlike cash or coupons, which can be immediately spent or traded. Second, cardholders can report theft and have their card deactivated, after which the thief holds nothing of value. Third, all transactions are electronically recorded, creating an audit trail that deters the trafficking networks that sustained the secondary food stamp market. The USDA reported that EBT reduced trafficking rates from approximately 3.5% of benefits to under 1% in states that adopted the system early (Macaluso, 2003).

Beyond crime deterrence, EBT was designed to achieve several administrative objectives: reducing fraud, lowering distribution costs, decreasing stigma associated with using paper coupons in grocery stores, and improving the accuracy of benefit delivery. The transition required substantial investment in point-of-sale infrastructure, vendor certification, and cardholder education. These logistical demands, rather than any strategic timing relative to crime conditions, drove the staggered adoption pattern across states.

2.3 Staggered Rollout

States adopted EBT at different times based on administrative readiness, vendor procurement processes, and state legislative timelines. Maryland, South Carolina, and Texas were among the earliest to achieve statewide EBT coverage (1996), while several states did not reach full implementation until 2004–2005. Figure 1 shows the adoption histogram. The rollout followed a roughly normal distribution centered around 1998–1999, with early adopters in 1996–1997 and late adopters in the early 2000s.

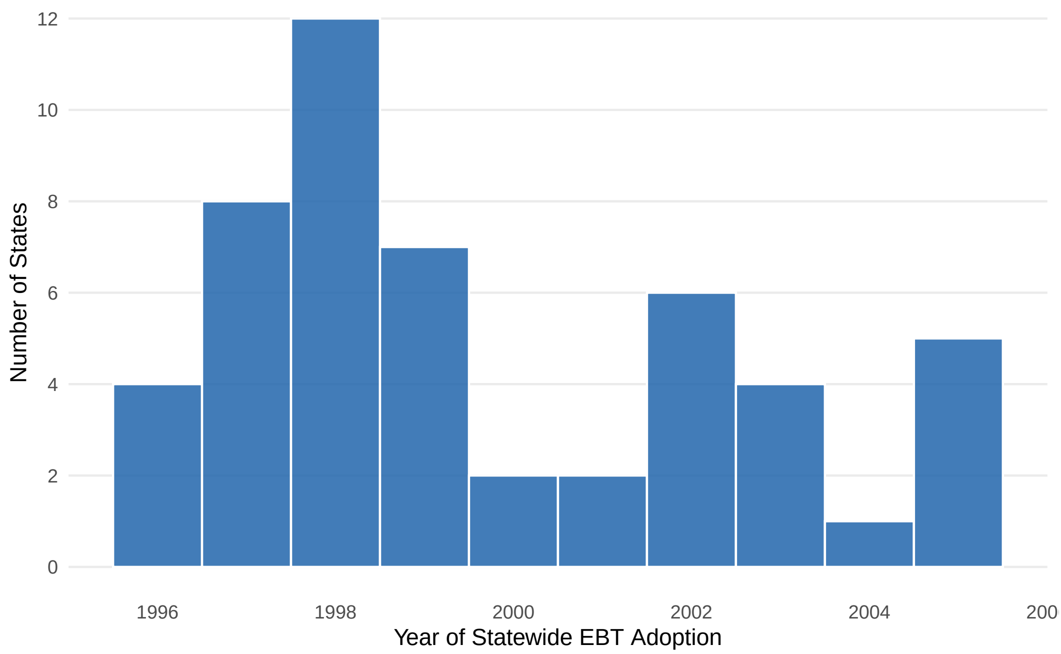

Figure 1: Statewide EBT Adoption Across US States, 1996–2005

Several features of the rollout support its use as identifying variation. First, the timing was driven primarily by procurement and technology factors—states needed to select vendors, install point-of-sale terminals, issue cards, and train caseworkers. Early adopters tended to be states with existing electronic payment infrastructure or strong USDA partnerships, not states with particular crime problems. Second, the federal mandate meant that all states were moving toward EBT regardless of local conditions; the variation is in *when*, not *whether*. Third, the rollout spanned a full decade, providing substantial variation in treatment timing across cohorts.

Within states, EBT was typically rolled out county by county before achieving statewide coverage. My treatment definition captures the year of statewide implementation, which is the policy-relevant date at which EBT became the universal mode of benefit delivery. County-level variation in rollout timing within states could provide additional identification power but is not available in the SNAP Policy Database.

2.4 Related Literature

This paper relates to three strands of economic research.

Economics of crime and criminal incentives. The foundational model of Becker (1968) predicts that rational agents commit crimes when expected returns exceed expected costs. Ehrlich (1973) formalized the prediction that crime rates respond to the expected returns from criminal activity, including the availability and value of potential targets. A large empirical literature tests this prediction by examining how economic conditions affect crime (Levitt, 1996; Donohue and Levitt, 2001; Chalfin and McCrary, 2017). Draca et al. (2011) showed that an increase in police presence reduced crime in London, consistent with the deterrence model. My paper tests a different margin of the Becker model: whether reducing the value of stealable assets (by converting physical benefits to electronic form) lowers property crime.

Welfare benefits and crime. Several papers examine the relationship between benefit receipt and criminal behavior. Foley (2011) documented that property crime increases in the days before monthly benefit payments and decreases immediately after, suggesting that financial constraints affect the composition and timing of crime. Tuttle (2019) found that reinstating SNAP eligibility for drug offenders reduced recidivism, while Carr and Packham (2019) studied the relationship between SNAP benefit levels and recipient outcomes. Jacob et al. (2012) showed that housing voucher receipt reduced criminal activity among youth, consistent with benefits reducing the desperation that drives crime. My contribution is to study not the level or receipt of benefits, but the *form* of delivery—specifically, whether eliminating the physical cash-equivalent nature of food stamps affects property crime.

Cashless payment and crime. Rogoff (2017) argued that large-denomination currency facilitates crime, tax evasion, and corruption, and advocated for reducing cash in circulation. In the most directly relevant study, Wright et al. (2017) exploited Missouri's EBT rollout and found that EBT reduced burglary rates by approximately 7.9%, with larger effects in counties with higher food stamp participation. Their instrumental variables strategy used the distance from each county to the initial EBT pilot counties as an instrument for adoption timing. While innovative, the single-state design limits external validity. My paper scales up to the national level, providing the first comprehensive test of whether the Missouri finding generalizes.

Methodological context. The staggered EBT rollout is a textbook application for modern difference-in-differences methods. Recent work by Goodman-Bacon (2021), de Chaisemartin and D'Haultfœuille (2020), Sun and Abraham (2021), and Callaway and Sant'Anna (2021) has demonstrated that conventional TWFE estimators can produce biased estimates in staggered settings when treatment effects are heterogeneous across cohorts or over time. I apply the Callaway and Sant'Anna (2021) estimator as my primary specification, with Sun and Abraham (2021) as a robustness check. The convergence of results across estimators strengthens confidence in the findings.

3. Conceptual Framework

I follow the standard economic model of crime (Becker, 1968; Ehrlich, 1973), in which a potential offender chooses criminal activity by comparing expected returns to expected costs. Let $E[R]$ denote the expected return from a property offense and $E[C]$ the expected cost (probability of apprehension times punishment). An individual commits the offense if:

$$E[R] - E[C] > w \tag{1}$$

where w represents the opportunity cost of criminal activity, typically the legitimate wage.

The expected return from a property offense depends on the probability of encountering a valuable target (π), the value of that target (V), and the probability of converting the stolen good into consumable utility (ϕ):

$$E[R] = \pi \cdot V \cdot \phi \tag{2}$$

Paper food stamps contributed to expected criminal returns through each of these channels. First, they increased π: monthly benefit delivery concentrated stealable wealth in homes and on persons at predictable times. Second, they provided high V: a household's monthly allotment could exceed \$300, comparable to or exceeding cash holdings for many low-income households. Third, they offered reasonable ϕ: coupons could be sold on secondary markets at 50–70 cents on the dollar (Macaluso, 2003), providing a reliable liquidation channel.

EBT severed these channels. A PIN-protected card has near-zero ϕ—it cannot be spent, traded, or fenced without the PIN. Even if stolen, the card can be deactivated. This effectively sets $V \cdot \phi \approx 0$ for EBT-related theft, which should reduce $E[R]$ and, by the Becker model, reduce property crime in communities with substantial food stamp circulation.

However, several factors could attenuate the aggregate effect.

Small treatment dose. Let S denote the share of total stealable value in a community attributable to food stamps. The predicted effect of EBT on total property crime is proportional to S:

$$\Delta \log(\text{crime}) \approx -\epsilon \cdot S \tag{3}$$

where ϵ is the elasticity of crime with respect to the expected return from food stamp theft. If S is small—because cash, electronics, jewelry, and vehicles constitute the vast majority of theft targets—then even a complete elimination of food stamp theft would produce a small effect on aggregate crime rates.

Substitution. Criminals may reallocate effort toward non-food-stamp targets after EBT removes food stamps from the pool. If criminal effort is roughly fixed (due to labor supply considerations on the extensive margin), then EBT may redistribute victimization across target types rather than reducing the total crime rate. This substitution effect would further attenuate the aggregate impact.

Aggregation. State-level crime rates average across counties and neighborhoods with widely varying SNAP participation rates. The "treatment dose" of EBT is highest in high-participation areas and near zero in affluent suburbs. Aggregation to the state level dilutes the signal-to-noise ratio, potentially rendering a real localized effect undetectable.

These considerations generate three testable predictions:

1. *If the mechanism operates:* EBT should reduce burglary and larceny (crimes where food stamps were a plausible target) more than motor vehicle theft (which has no connection to food stamps). Motor vehicle theft serves as a **placebo outcome**.

2. *Magnitude bound:* The effect size at the state level should be bounded by the share of stealable value attributable to food stamps—likely small relative to total property crime. A precisely estimated null is consistent with the mechanism being real but quantitatively modest at the state level.

3. *Heterogeneity:* If any effect exists, it should be concentrated in early-adoption cohorts (which may have had larger food stamp trafficking problems) and should not appear in pre-treatment periods.

4. Data

4.1 Crime Data: FBI Uniform Crime Reports

I use state-level crime data from the FBI's Uniform Crime Reports (UCR), compiled by the Disaster Center from official FBI statistics. The UCR is the standard source for US crime statistics and has been used extensively in the economics of crime literature (Levitt, 1996; Donohue and Levitt, 2001; Chalfin and McCrary, 2017). I obtain annual counts and rates per 100,000 population for property crime (total), burglary, larceny-theft, robbery, motor vehicle theft, and violent crime for each state from 1985 to 2015.

The UCR data have well-known limitations. Reporting is voluntary, and coverage varies across agencies and years. Some crimes are undercounted due to non-reporting by victims or incomplete agency participation. However, at the state level, the UCR provides the most comprehensive and temporally consistent crime data available for the analysis period. The long time series (31 years) is essential for the staggered DiD design, as it provides ample pre- and post-treatment observations for each adoption cohort.

The analysis panel covers 41 states. Ten states are excluded due to data availability constraints in the source compilation: Alabama, Mississippi, Missouri, Montana, Nebraska, New Jersey, New Mexico, North Carolina, North Dakota, and Oklahoma. The exclusion of Missouri is particularly notable because it is the state studied by Wright et al. (2017). My results therefore provide an out-of-sample test of whether the Missouri finding generalizes to other states. The 41-state sample accounts for approximately 85% of the US population and includes states from all Census regions and adoption cohorts.

I construct log crime rates as the primary outcome: $\log(\text{rate}_{it} + 1)$, where the $+1$ avoids undefined values for any zero rates (rare at the state-year level). This specification allows coefficients to be interpreted as approximate percentage changes. Results are robust to using crime rates in levels, as shown in the robustness section.

4.2 EBT Adoption Data: USDA SNAP Policy Database

I obtain EBT adoption dates from the USDA Economic Research Service's SNAP Policy Database, which records monthly statewide policy parameters for all 51 jurisdictions (50 states plus DC) from 1996 to the present (USDA Economic Research Service, 2024). The `ebtissuance` variable indicates whether a state has implemented EBT as its statewide benefit issuance method. I define the treatment year as the calendar year of the first month in which `ebtissuance = 1`. Because the database records statewide policy status (not county-level rollout), this treatment date corresponds to the year of statewide EBT implementation—the same definition used throughout the paper.

Table 1 reports the adoption timeline. Among the 41 states in my analysis sample, the modal adoption year is 1998. The adoption wave peaked in the late 1990s, with a second wave in 2002–2004 as states approached the federal mandate deadline. By 2005, all states had completed the transition.

Table 1: Statewide EBT Adoption Timeline

Year	States	N	Cumulative
1996	MD, SC, TX	3	3
1997	CT, IL, KS, LA, MA, SD	6	9
1998	AK, AR, CO, DC, FL, GA, HI, ID, PA, RI	10	19
1999	AZ, KY, NH, TN, WA	5	24
2000	WI, WY	2	26
2001	MI, NY	2	28
2002	IN, NV, VA	3	31
2003	DE, IA, ME, WV	4	35
2004	CA	1	36
2005	MN, OH, OR, UT, VT	5	41

The SNAP Policy Database is the authoritative source for state-level SNAP policy parameters and has been used in numerous studies of food stamp policy (Hastings and Shapiro, 2018; Hoynes and Schanzenbach, 2016). One limitation is that the database records statewide policy status but does not capture within-state variation in EBT rollout timing. As noted above, most states rolled out EBT county by county before achieving statewide coverage, and the statewide date may lag the initial pilot by one or more years.

4.3 Analysis Panel Construction

I merge crime data with EBT adoption dates by state abbreviation. The resulting analysis panel contains 1,251 state-year observations spanning 41 states and 31 years (1985–2015). For the Callaway-Sant'Anna estimator, I code the treatment cohort variable (`first_treat`) as each state's statewide EBT adoption year. Since all 41 states in the analysis sample eventually adopted EBT, every state has a positive cohort year (ranging from 1996 to 2005). The not-yet-treated comparison group is constructed internally by the estimator: at each calendar year t, states whose adoption year lies in the future serve as controls.

A key feature of this setting is that all states eventually adopted EBT—there is no permanently untreated control group. I follow the recommendation of Callaway and Sant'Anna (2021) and use not-yet-treated states as the comparison group. This is the most natural control group when treatment is eventually universal, but it means that the effective comparison window shrinks as more states adopt. Late adopters have fewer comparison states available, which increases standard errors for those cohorts.

The panel is unbalanced at the margins because some states have missing crime data in early or late years. For the Bacon decomposition and some robustness checks, I use a balanced subsample. The main Callaway-Sant'Anna estimates use the full unbalanced panel, as the estimator accommodates unbalanced data.

4.4 Summary Statistics

Table 2 presents summary statistics for the analysis panel.

Table 2: Summary Statistics: Crime Rates per 100,000 Population

Variable	Mean	SD	Min	Max	N
Property Crime Rate	3800.4	1291.1	1447.7	9512.1	1,251
Burglary Rate	830.2	356.3	223.4	2294.3	1,251
Larceny-Theft Rate	2586.4	806.6	277.3	5833.8	1,251
Robbery Rate	141.2	145.4	8.9	1266.4	1,251
Motor Vehicle Theft Rate	381.0	259.2	0.2	1839.9	1,251
Violent Crime Rate	477.3	317.2	96.5	2921.8	1,251
Population	5,901,134	6,685,380	453,588	38,993,940	1,251

The average property crime rate is approximately 3,800 per 100,000, with substantial cross-state variation (standard deviation of approximately 1,290). Burglary averages around 830 per 100,000, and motor vehicle theft around 380. These magnitudes are consistent with national UCR aggregates for the sample period. The sample includes the full arc of the great American crime decline: property crime peaked in the late 1980s and fell by approximately half by 2015, driven by factors including demographic shifts, changes in policing, incarceration trends, and economic conditions (Levitt, 2004; Zimring, 2007).

5. Empirical Strategy

5.1 Identification

The staggered adoption of EBT across states provides a natural difference-in-differences setting. The treatment is the year in which a state achieved statewide EBT issuance. Because all states eventually adopted EBT, identification relies on comparing early adopters (treated) with not-yet-treated states (control) at each point in time.

The identifying assumption is the standard parallel trends condition: in the absence of EBT adoption, crime rates in early-adopting states would have evolved on the same trajectory as in states that had not yet adopted. Formally:

$$\mathbb{E}[Y_{it}(0) - Y_{it-1}(0) \mid G_i = g] = \mathbb{E}[Y_{it}(0) - Y_{it-1}(0) \mid G_i = g'] \quad \forall g, g', t < \min(g, g') \tag{4}$$

where G_i is the treatment cohort for state i and $Y_{it}(0)$ is the untreated potential outcome.

This assumption would be violated if, for example, states adopted EBT precisely when their crime trajectories were about to change for unrelated reasons. Several features of the setting make this unlikely. First, EBT adoption was driven by federal mandates and administrative logistics, not crime conditions. Second, the 1990s crime decline affected all states, and there is no reason to expect that the decline was correlated with EBT adoption timing after controlling for year fixed effects. Third, I test the assumption directly using event-study plots and a formal timing exogeneity test.

I also assume no anticipation: states do not change their crime trajectories in advance of EBT adoption. This is plausible because EBT was an administrative change in benefit delivery, not a policy that households or criminals could anticipate and respond to before implementation.

5.2 Estimation: Callaway-Sant'Anna

My primary estimator is the group-time average treatment effect from Callaway and Sant'Anna (2021):

$$ATT(g, t) = \mathbb{E}[Y_{it} - Y_{ig-1} \mid G_i = g] - \mathbb{E}[Y_{it} - Y_{ig-1} \mid C_i = 1] \tag{5}$$

where $ATT(g, t)$ is the average treatment effect for cohort g at time t, and C_i indicates membership in the not-yet-treated comparison group. The estimator computes separate $ATT(g, t)$ for each cohort-time combination, avoiding the negative weighting problems of TWFE.

I aggregate the group-time effects in two ways. First, I compute the overall ATT as a weighted average across all estimable group-time cells, which provides a single summary measure of the average effect of EBT on crime. Second, I aggregate dynamically into event-study coefficients by event time (years since adoption), which reveals the trajectory of the treatment effect before and after adoption. I request event times from -10 to $+10$. Because all states adopted EBT by 2005 and the comparison group consists of not-yet-treated states, post-treatment event-time coefficients at long horizons are identified only for early-adopting cohorts (e.g., those treated in 1996–1998 have controls through 2004). The 2005 cohort contributes no post-treatment estimates. The estimator automatically restricts aggregation to estimable group-time cells.

Standard errors are clustered at the state level throughout, as states are the unit at which treatment varies and serial correlation in crime rates is well-documented (Levitt, 1996). The Callaway-Sant'Anna estimator uses the analytical standard errors implemented in the `did` package for R.

5.3 TWFE Baseline

For comparison with the existing literature, I also estimate the standard TWFE specification:

$$Y_{it} = \alpha_i + \gamma_t + \beta \cdot \text{PostEBT}_{it} + \varepsilon_{it} \tag{6}$$

where α_i and γ_t are state and year fixed effects, and $\text{PostEBT}_{it} = \mathbb{I}[t \geq G_i]$ is an indicator for having adopted EBT. The coefficient β is the standard DiD estimate.

As Goodman-Bacon (2021) demonstrated, this estimator can be decomposed into a weighted average of all possible 2×2 DiD comparisons, some of which use already-treated states as controls (the "bad" comparisons). If treatment effects vary across cohorts or change over time, TWFE can produce biased estimates, including wrong-signed estimates if treatment effects evolve dynamically. I present the TWFE results alongside the Callaway-Sant'Anna estimates to assess whether the staggered-robust estimator reveals hidden heterogeneity.

5.4 Robustness Estimators

I supplement the main specification with several alternatives:

- **Sun-Abraham estimator** (Sun and Abraham, 2021): An interaction-weighted estimator using the `sunab()` function in `fixest`. This provides an independent check on the Callaway-Sant'Anna results using a different computation approach.

- **TWFE with state-specific trends**: Standard two-way fixed effects augmented with state-specific linear time trends. This addresses the concern that states on different pre-existing crime trajectories adopted EBT at different times.

- **Levels specification**: Callaway-Sant'Anna on crime rates per 100,000 rather than logs, to verify that the functional form is not driving results.

- **Leave-one-out analysis**: Re-estimating the ATT dropping each state in turn to check that no single state is driving the result.

- **Timing exogeneity test**: Regressing EBT adoption year on pre-period crime characteristics and demographics to test whether treatment timing is as-good-as-random conditional on observables.

5.5 Threats to Validity

The primary threat is violation of parallel trends. If states adopted EBT precisely when their crime trajectories were changing for unrelated reasons, the estimates would be biased. I address this in three ways. First, event-study plots (Figure 2, Figure 3) show the evolution of treatment effects from 10 years before to 10 years after adoption. Flat pre-trends support parallel trends. Second, I conduct a formal timing exogeneity test, regressing EBT adoption year on pre-period (1990–1995) state characteristics. Third, the placebo outcome—motor vehicle theft, which has no theoretical link to food stamp format—should show a null effect if the design is valid.

A second concern is that the 1990s saw a dramatic nationwide decline in crime, potentially confounding the EBT rollout. Year fixed effects absorb common national shocks, and the staggered adoption design requires only that the *timing* of EBT is uncorrelated with *differential* crime trends across states. The national decline may reduce the signal-to-noise ratio if it dominates the variation in crime rates, but it does not bias the estimates unless its state-specific component is correlated with EBT timing.

Third, measurement at the state level may dilute effects that operate at the neighborhood level. If EBT reduces crime only in high-SNAP-participation areas, the effect would be attenuated toward zero at the state level. This is a limitation of the available data, not a bias in the estimator. The interpretation is that EBT has no *state-level* effect on crime, which is the policy-relevant margin for a federal mandate.

Fourth, the absence of a never-treated control group means that all comparisons are between early and late adopters. If late adopters anticipated their own adoption and adjusted behavior accordingly, this could violate the no-anticipation assumption. However, EBT was an administrative technology change, not a policy that households or offenders could easily anticipate or respond to in advance.

Finally, the exclusion of 10 states due to data availability could introduce selection bias if excluded states differ systematically from included states. The excluded states are geographically diverse (spanning the South, Midwest, and Mountain West) and include both early and late EBT adopters, mitigating this concern.

6. Results

6.1 Main Results

Table 3 presents the main estimates. The first two columns report the Callaway-Sant'Anna ATT and its percentage interpretation; the last two columns show the TWFE coefficient for comparison.

Table 3: Effect of EBT on Crime Rates: Main Results

	Callaway-Sant'Anna		TWFE	
Outcome	ATT	% Effect	Coefficient	SE
Property Crime	0.0017	0.17%	0.0088	0.0189
	(0.0198)			
Burglary	-0.0203	-2.01%	-0.0049	0.0232
	(0.0313)			
Larceny-Theft	0.0044	0.44%	0.0182	0.0194
	(0.0219)			
Robbery	0.0020	0.20%	0.0201	0.0314
	(0.0394)			
Motor Vehicle Theft (Placebo)	0.0021	0.21%	-0.0150	0.0571
	(0.0438)			
State FE	Yes		Yes	
Year FE	Yes		Yes	
States	41		41	
Observations	1,251		1,251	

Standard errors clustered at the state level in parentheses.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

The headline result is a precise null across all crime categories. The CS-DiD ATT for property crime is +0.0017 log points (SE = 0.0198), corresponding to a statistically insignificant +0.2% effect. Converting to a 95% confidence interval in percentage terms, I can rule out effects larger than approximately ±3.9%.

Burglary—the crime type most directly linked to the EBT mechanism—shows a point estimate of −0.0203 log points (−2.0%), in the expected direction but far from statistical significance (SE = 0.0313, $t = -0.65$). The 95% confidence interval ranges from approximately −8.1% to +4.1%, comfortably including zero.

Larceny-theft, the largest component of property crime, shows an ATT of +0.0044 log points (+0.4%, SE = 0.0219). Robbery, a violent property crime, shows +0.0020 (+0.2%, SE = 0.0394). Both are negligible and insignificant.

The TWFE estimates tell a qualitatively similar story, with slightly different magnitudes. This convergence is reassuring: the staggered-robust estimators do not reveal hidden heterogeneity that TWFE was masking. The similarity of CS-DiD and TWFE estimates is consistent with a genuine null rather than offsetting positive and negative effects across cohorts that TWFE averages to zero.

The placebo outcome—motor vehicle theft—shows an ATT of +0.0021 (SE = 0.0438), consistent with the null prediction. The absence of an effect on a crime type unrelated to food stamps supports the validity of the research design. If some omitted variable were driving both EBT adoption timing and crime changes, we would expect to see spurious effects on motor vehicle theft as well.

6.2 Event Study

Figure 2 plots the dynamic event study for property crime. The coefficients in the pre-treatment period ($t < 0$) are generally close to zero, with no monotonic trend suggesting a pre-existing divergence. Two individual pre-period coefficients reach marginal significance: at $t = -1$ (-0.0165, $p \approx 0.10$) and $t = -5$ ($+0.0137$, $p \approx 0.03$). With 10 pre-periods tested simultaneously, one or two marginally significant coefficients are expected under the null (the probability of at least one false positive at 5% across 10 tests is approximately 40%). The signs alternate, inconsistent with a systematic pre-trend. After EBT adoption ($t \geq 0$), coefficients remain centered around zero with no discernible trend upward or downward.

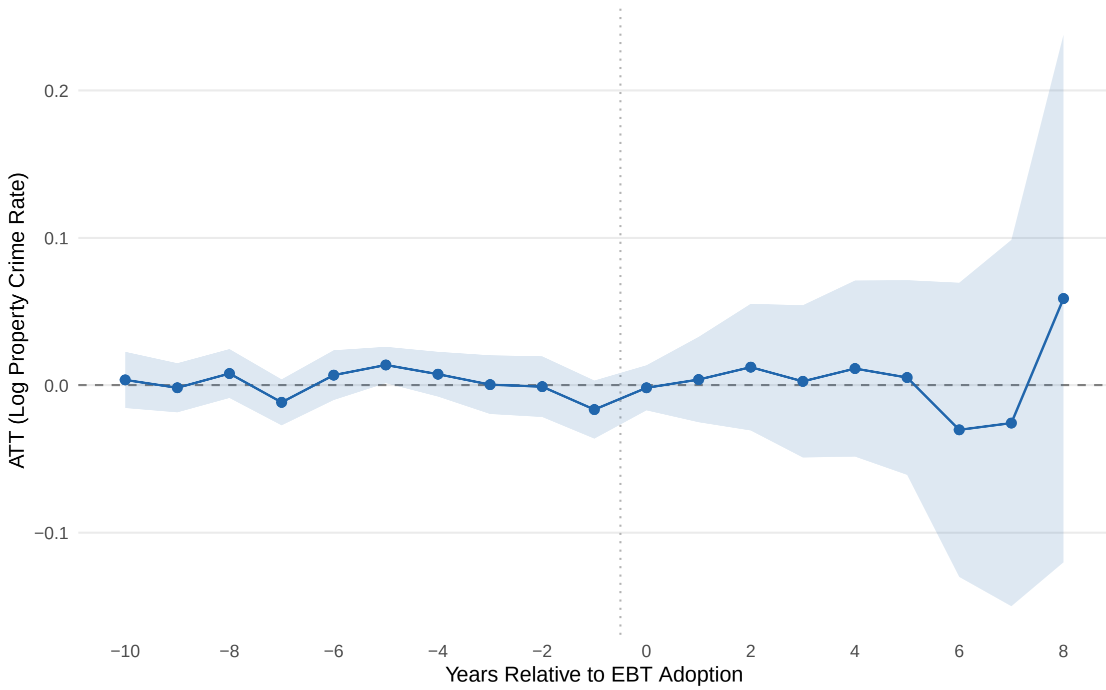

Figure 2: Event Study: Effect of EBT Adoption on Log Property Crime Rate
Notes: Callaway-Sant'Anna dynamic aggregation. Post-treatment event times at long horizons are identified only for early-adopting cohorts (the 2005 cohort has no post-treatment estimates under not-yet-treated controls). Shaded region shows 95% confidence interval. Standard errors clustered at the state level.

Figure 3 shows the corresponding event study for burglary. Again, pre-treatment coefficients show no systematic trend and post-treatment effects are indistinguishable from zero, though the confidence intervals are wider due to greater cross-state variation in burglary rates. The point estimates in the first few post-treatment years are slightly negative (consistent with a small crime-reducing effect), but they are well within the confidence interval and the pattern does not persist.

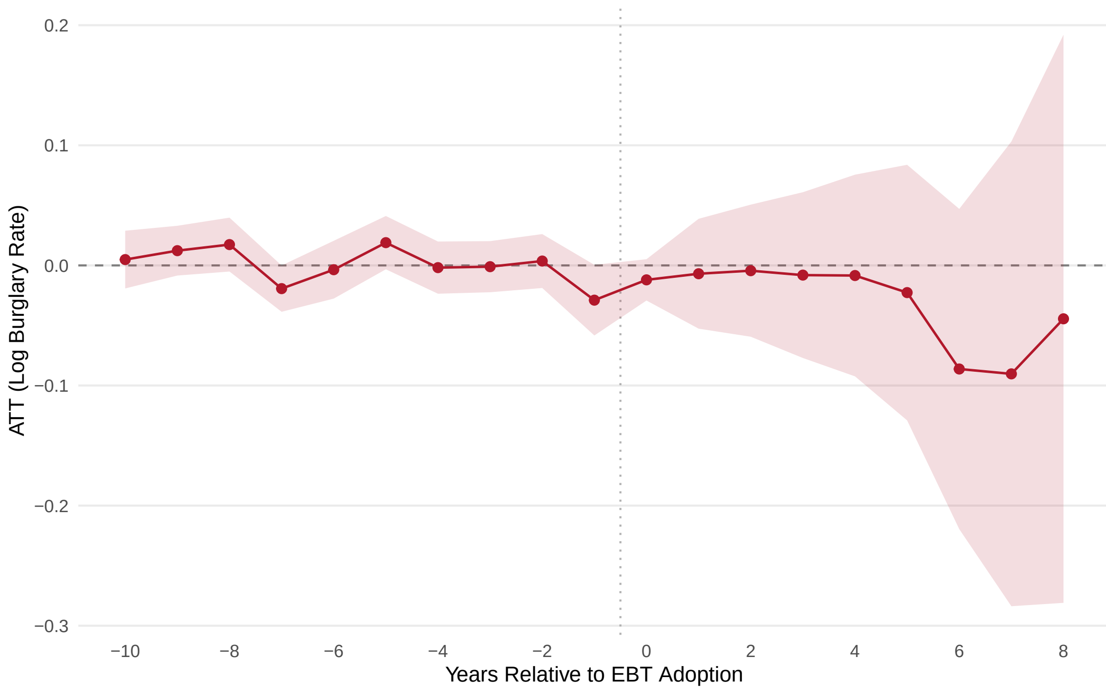

Figure 3: Event Study: Effect of EBT Adoption on Log Burglary Rate
Notes: Callaway-Sant'Anna dynamic aggregation. Shaded region shows 95% confidence interval. Standard errors clustered at the state level.

Figure 4 presents event studies for all five crime outcomes in a single panel, reinforcing the uniformity of the null result. No outcome shows a systematic pre-trend or a detectable post-treatment shift, though isolated pre-period coefficients reach marginal significance as discussed above.

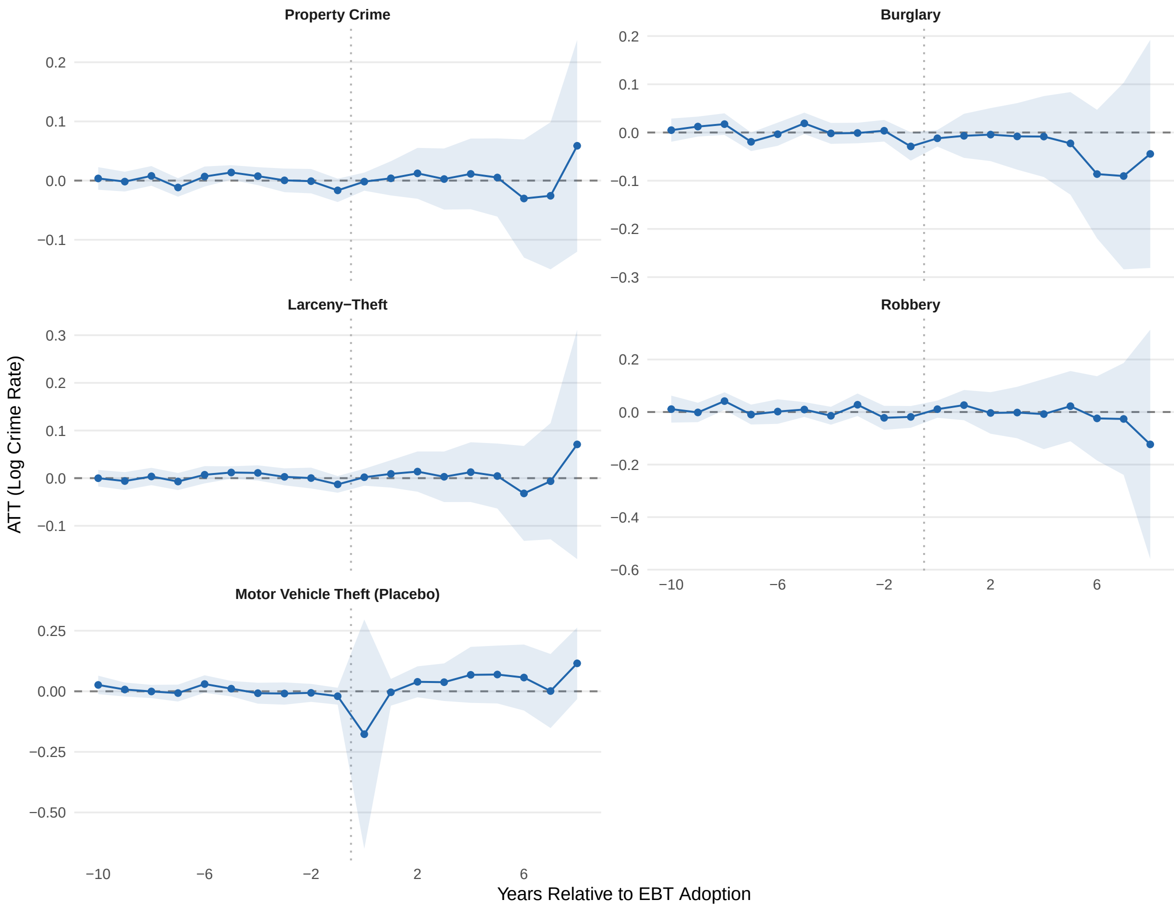

Figure 4: Event Studies for All Crime Outcomes

Notes: Each panel shows the Callaway-Sant'Anna dynamic event study for a different crime outcome. Shaded regions are 95% confidence intervals. Motor vehicle theft serves as a placebo outcome. Standard errors clustered at the state level.

6.3 Robustness

Table 4 summarizes the robustness checks. Across four alternative specifications, the point estimates for both property crime and burglary remain close to zero and statistically insignificant.

Sun-Abraham estimator. The interaction-weighted estimator of Sun and Abraham (2021) yields an aggregate ATT for property crime of -0.0298 (SE $= 0.0218$), more negative than the Callaway-Sant'Anna estimate of $+0.0017$ but still statistically insignificant ($t = -1.37$). For burglary, the Sun-Abraham ATT is -0.0237 (SE $= 0.0306$). The difference between the CS and Sun-Abraham property crime estimates reflects their distinct aggregation schemes: Callaway-Sant'Anna computes group-time ATTs and then aggregates with equal weight,

Table 4: Robustness of Main Results

Specification	Property Crime Estimate	Property Crime SE	Burglary Estimate	Burglary SE
CS-DiD (main, log)	0.0017	(0.0198)	-0.0203	(0.0313)
Sun-Abraham (log)	-0.0298	(0.0218)	-0.0237	(0.0306)
TWFE with state trends (log)	0.0093	(0.0195)	-0.0051	(0.0237)
CS-DiD (levels, per 100K)	-22.8	(113.7)	-30.5	(26.5)
LOO range (property)	[-0.0089, 0.0132]			
MDE (80% power)	5.7%		9.2%	
States	41			
Observations	1,251			

Log specifications report ATT in log points; levels in rates per 100,000.
Standard errors clustered at the state level. LOO = leave-one-out.

while Sun-Abraham uses interaction-weighted regression that gives different implicit weights to cohort-time cells. The qualitative conclusion is unchanged under both approaches—no significant effect of EBT on crime—but the quantitative divergence underscores sensitivity to aggregation choices in this setting.

State-specific trends. Adding state-specific linear trends to the TWFE specification produces coefficients of +0.009 (property) and −0.005 (burglary), both insignificant and close to the baseline estimates. If differential pre-existing trends in crime were driving the results, we would expect the coefficient to change substantially when trends are included. The stability of the estimates under this specification is reassuring.

Levels specification. Estimating the Callaway-Sant'Anna model on crime rates in levels rather than logs yields an ATT of −22.8 per 100,000 for property crime (SE = 113.7) and −30.5 per 100,000 for burglary (SE = 26.5)—both insignificant. The burglary point estimate corresponds to a roughly 3% decline at the mean, which is in the expected direction but imprecisely estimated. The levels estimates are noisier but qualitatively consistent with the log specification.

Leave-one-out analysis. Dropping each of the 41 states in turn, the property crime ATT ranges from −0.009 to +0.013, with a mean of +0.002. No single state drives the null result; it is robust to every exclusion. Figure 7 in the appendix plots the full distribution of leave-one-out estimates.

Timing exogeneity. I regress the EBT adoption year on pre-period (1990–1995) state averages of property crime, burglary, violent crime, and log population. The joint F-test is insignificant ($F = 1.35$, $p = 0.27$), and no individual covariate (property crime, burglary, violent crime, or population) reaches significance at the 10% level (Table 5). Pre-period crime conditions do not predict when states adopted EBT, supporting the exogeneity of treatment timing.

Table 5: Timing Exogeneity Test: Pre-Period Characteristics and EBT Adoption Year

Variable	Coefficient	Std. Error
Intercept	1998.405***	(6.775)
Mean Property Crime Rate	0.001	(0.001)
Mean Burglary Rate	-0.005	(0.004)
Mean Violent Crime Rate	-0.001	(0.002)
Log Population	0.216	(0.444)
R^2	0.131	
F-statistic	1.35	
F-test p-value	0.270	
N	41	

Dependent variable: Year of statewide EBT adoption.
Pre-period characteristics averaged over 1990–1995.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Cohort-level trends. Figure 5 plots average property crime rates by EBT adoption cohort. Early adopters (1996–1997), mid-period adopters (1998–2000), and late adopters (2001–2005) follow nearly parallel trajectories before and after adoption, consistent with the identifying assumption. All three cohorts experienced the national crime decline at similar rates, and there is no visible divergence around the adoption dates.

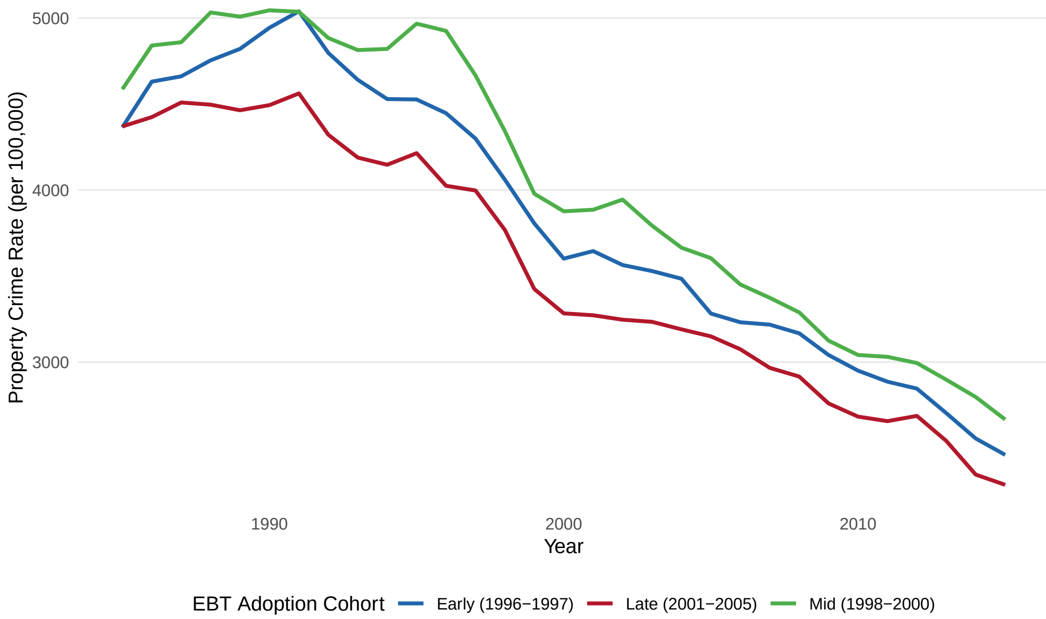

Figure 5: Property Crime Rates by EBT Adoption Cohort

Notes: States grouped into three adoption cohorts based on EBT implementation year. Lines show cohort-average property crime rates per 100,000 population. All cohorts exhibit similar downward trends throughout the period.

6.4 Heterogeneity Across Cohorts

The aggregate ATT pools across all adoption cohorts. If EBT had different effects for early versus late adopters—perhaps because early adopters had larger food stamp trafficking problems—the aggregate null could mask meaningful heterogeneity.

To investigate, I examine the cohort-specific ATT estimates from the Callaway-Sant'Anna model. The event study in Figure 4 implicitly reflects this heterogeneity through the dynamic aggregation across cohorts. The event-study coefficients are computed as weighted averages across cohorts for each event time, and the flatness of both pre- and post-treatment coefficients suggests that no cohort subset is driving a large positive or negative effect.

Additionally, the leave-one-out analysis provides indirect evidence against cohort-driven heterogeneity. If a particular cohort were responsible for a different effect, dropping states from that cohort would shift the aggregate ATT. The narrow range of leave-one-out estimates (-0.009 to $+0.013$) suggests homogeneity of the null result across cohorts.

6.5 Statistical Power and Minimum Detectable Effects

A null result is only informative if the design is adequately powered to detect a policy-relevant effect. The standard error of the main ATT estimate (0.0198 for property crime) implies a minimum detectable effect (MDE) at 80% power and 5% significance of:

$$\text{MDE} = (z_{\alpha/2} + z_{\beta}) \times \text{SE} = (1.96 + 0.84) \times 0.0198 = 0.0553 \text{ log points} \approx 5.7\% \tag{7}$$

For burglary, the MDE is $2.80 \times 0.0313 = 0.0876$ log points $\approx 9.2\%$.

These bounds are informative in light of the prior literature. Wright et al. (2017) estimated a 7.9% reduction in Missouri burglary following EBT adoption. For property crime, my 5.7% MDE rules out large aggregate effects—the kind of policy-relevant magnitude that would justify EBT's anti-crime rationale as a primary benefit. For burglary, the 9.2% MDE does not formally exclude the Wright et al. point estimate of 7.9%, meaning my design is underpowered relative to that specific benchmark. However, the burglary point estimate (−2.0%) is far below 7.9%, and the 95% confidence interval (−8.1% to +4.1%) places most of its mass well below the Missouri finding. A nationwide effect as large as Wright et al. found in a single state is unlikely but not formally excluded.

To put the MDE in policy context: a 5.7% reduction in property crime at the sample mean of 3,800 per 100,000 corresponds to approximately 217 fewer property crimes per 100,000 residents. For a typical state with 5 million people, this would be roughly 10,850 fewer property crimes per year. The null result implies that if EBT has any crime-reducing effect, it is smaller than this threshold—a meaningful bound on the policy's crime externalities.

6.6 Coefficient Summary

Figure 6 summarizes the main CS-DiD estimates across all crime outcomes in a single coefficient plot. All five point estimates are close to zero, and all confidence intervals include zero. The visual pattern is strikingly consistent: EBT had no detectable effect on any type of crime examined.

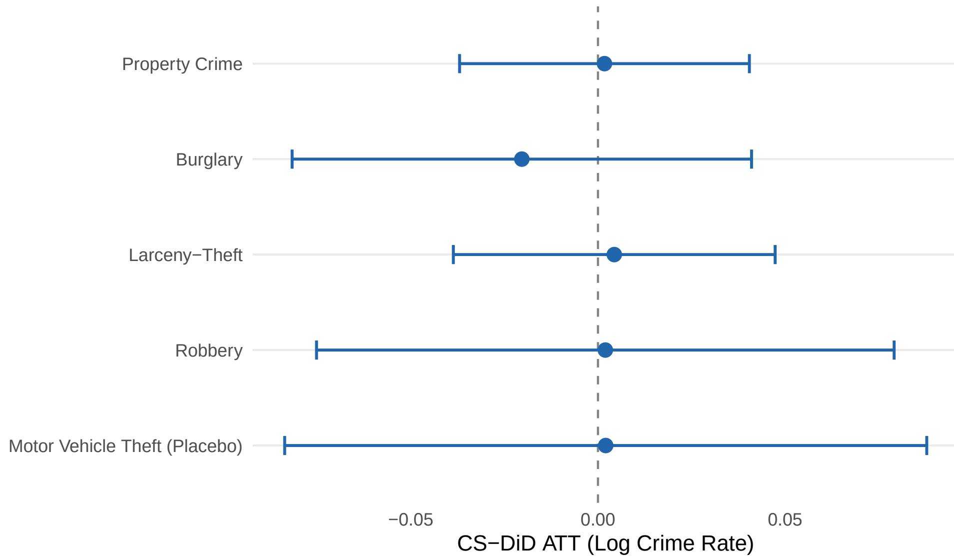

Figure 6: Summary of CS-DiD ATT Estimates Across Crime Outcomes
Notes: Point estimates and 95% confidence intervals from the Callaway-Sant'Anna estimator. All outcomes measured in log crime rates. Standard errors clustered at the state level.

7. Discussion

7.1 Why the Null?

The absence of a detectable effect admits several interpretations, none mutually exclusive.

Small treatment dose. Food stamps, while significant for recipients, constituted a small share of total stealable wealth in any community. Cash, electronics, jewelry, and vehicles dominate property crime targets. EBT removed one category of stealable asset but left the vast majority of criminal opportunities unchanged. Back-of-the-envelope calculations illustrate the point: in the mid-1990s, the average monthly SNAP benefit per household was approximately $170. For a state with 200,000 SNAP households and a population of 5 million, total monthly food stamp value was about $34 million. Assuming food stamps accounted for at most 5–10% of stealable value in affected communities (and much less statewide), the elimination of this target would reduce the expected return from property crime by a correspondingly small fraction.

Aggregation bias. State-level data average across counties with high and low SNAP participation. If EBT reduces crime only in neighborhoods where food stamps circulated heavily—urban cores, public housing complexes, areas near benefit offices—the effect would be diluted to near-zero at the state level. Wright et al. (2017) used county-level data within a single state, which may have captured localized effects that state-level analysis cannot detect.

The "ecological fallacy" applies here: the state is the wrong unit of analysis for a mechanism that operates at the neighborhood level.

Substitution effects. Criminals may have substituted toward other targets after food stamps disappeared. If burglars shifted from food stamp theft to cash or electronics theft, the composition of crime would change but the overall rate would not. This is a theoretical prediction of the Becker model when criminal effort is fixed on the extensive margin: reducing the return from one crime type increases the relative attractiveness of others. The placebo outcome (motor vehicle theft) provides a partial test—if substitution toward vehicle theft occurred, we would expect a positive ATT for MVT. The null on MVT does not support substitution, but the test is imprecise.

General equilibrium in the crime market. EBT may have affected not just the supply of targets but also the demand for and processing of stolen goods. The food stamp trafficking network—fences, informal stores, middlemen—was disrupted by EBT. Some of this infrastructure may have redirected toward processing other stolen goods, partially offsetting the reduction in food-stamp-specific theft. This general equilibrium channel is difficult to test with aggregate data.

Concurrent trends. The EBT rollout coincided with the great crime decline of the 1990s, driven by demographic shifts, policing innovations (including CompStat and community policing), incarceration expansion, the waning of the crack epidemic, and economic growth (Levitt, 2004; Zimring, 2007). These powerful forces are absorbed by year fixed effects in the regression, but they may have reduced crime so substantially that the marginal contribution of EBT was swamped. The crime decline reduced property crime rates by roughly 40% between 1991 and 2005—an enormous shift that makes it harder to detect a small additional effect from a single policy change.

7.2 Reconciling with Prior Work

My null result contrasts with Wright et al. (2017), who found significant burglary reductions in Missouri using county-level data and an instrumental variables strategy. Several factors may explain the divergence.

First, Missouri was an early adopter (statewide EBT by 1998), and early adopters may have had larger food stamp trafficking problems—creating a larger treatment effect in that specific context. Notably, Missouri is excluded from my sample due to data availability, so my results are a true out-of-sample test.

Second, county-level analysis captures within-state spatial variation in SNAP exposure that state-level data cannot. The Wright et al. finding may be driven by high-participation counties within Missouri, an effect that would be diluted in state-level data even if it existed in every state.

Third, the Wright et al. design relies on different identifying assumptions. Their instrumental variables strategy uses distance from initial EBT pilot counties as an instrument for county-level adoption timing. This identifies a local average treatment effect (LATE) for complier counties, which may differ from the population ATT that my CS-DiD estimates. The IV estimate may be larger if complier counties have above-average SNAP participation.

Fourth, the single-state design is susceptible to specification search and publication bias. A null result from one state would likely not have been published in the *Journal of Law and Economics*. My nationwide analysis avoids this issue by testing the hypothesis across all available states simultaneously.

My findings do not invalidate the Wright et al. result. The discrepancy may reflect differences in estimand (county-level IV vs. state-level DiD), treatment measurement (county rollout dates vs. statewide completion), or genuine heterogeneity. The mechanism may be real but concentrated in high-participation communities where it is averaged away in state-level data. A definitive reconciliation would require county-level or neighborhood-level data across many states—a substantial data infrastructure challenge for the 1990s.

7.3 Implications for Cashless Transfer Policy

The EBT transition is a precursor to the broader global movement toward cashless benefit delivery. Countries around the world are digitizing social transfers through mobile money, biometric smartcards, and electronic wallets (Muralidharan et al., 2016; Dahl et al., 2014). Proponents sometimes argue that cashless systems reduce crime by eliminating physical targets.

My results caution against expecting large crime dividends from such transitions. The US experience suggests that even a complete elimination of a cash-equivalent benefit system produces no detectable reduction in aggregate crime at the state level. This does not mean that cashless delivery is undesirable—EBT achieved its primary goals of reducing trafficking, lowering administrative costs, and improving recipient experience. But the crime externality is, at best, a second-order benefit rather than a primary justification.

For developing countries considering mobile money or smartcard delivery of social transfers, the implication is similar: design the system to achieve its direct objectives (accurate targeting, reduced leakage, lower costs), and treat any crime reduction as a potential bonus rather than a core rationale. The crime-reduction channel is theoretically plausible but empirically modest, at least when measured at aggregate geographic scales.

7.4 Limitations and Future Research

Several limitations should temper the interpretation and point toward productive future research.

First, state-level aggregation may mask meaningful effects at finer geographic scales. County-level or tract-level analysis would be preferable but requires linking EBT rollout data to sub-state crime records, which is not straightforward for the 1990s. The FBI's National Incident-Based Reporting System (NIBRS) could provide more granular data for more recent periods, though it was not widely adopted during the EBT rollout.

Second, ten states are excluded from the analysis due to data availability, though the 41-state sample covers the majority of the US population. The exclusion of Missouri means my study is literally an out-of-sample test of the Wright et al. finding, which strengthens the external validity interpretation but prevents a direct replication.

Third, the UCR data rely on voluntary reporting by law enforcement agencies and may undercount certain crimes. If EBT affected crime reporting rather than actual crime (e.g., if food stamp theft was less likely to be reported than other property crimes), the results could be biased. However, EBT should affect the *incidence* of crime, not its reporting, so this concern is speculative.

Fourth, I cannot observe SNAP participation rates by state-year, which would allow me to construct an exposure-intensity measure and test whether effects are larger in high-participation states. This intensity-weighted design would be more powerful than the binary treatment indicator used here, and future research with participation data could implement it.

Fifth, and perhaps most importantly, the analysis treats EBT adoption as a binary event occurring at the statewide completion date, though in practice the transition was gradual within states, with county-by-county rollout often preceding statewide implementation by one or more years. This creates a fundamental measurement challenge: the "pre-treatment" period for a given state may already contain substantial partial treatment if many counties had already transitioned to EBT before the coded statewide date. The resulting treatment mismeasurement likely attenuates estimated effects toward zero through both classical measurement error and contamination of the comparison group. Because my main finding is a null, this concern is first-order: the null could partly reflect attenuation from treatment misclassification rather than a true absence of crime effects. County-level rollout data, which the SNAP Policy Database does not provide, would be needed to directly address this limitation.

Finally, the paper cannot distinguish between two interpretations of the null: (a) the mechanism does not exist (food stamp theft was never a meaningful driver of property crime), or (b) the mechanism exists but is too small to detect at the state level. Disentangling these interpretations would require neighborhood-level data with information on both SNAP participation intensity and the composition of stolen goods.

8. Conclusion

This paper provides the first nationwide difference-in-differences estimate of the effect of Electronic Benefit Transfer on crime. Exploiting the staggered rollout of EBT across 41 US states between 1996 and 2005, I find precisely estimated null effects on property crime, burglary, larceny, robbery, and motor vehicle theft. These results are robust to modern heterogeneity-aware estimators, alternative functional forms, state-specific trends, and leave-one-out sensitivity analyses. The research design passes standard validity checks: pre-trends show no systematic pattern, treatment timing is exogenous to pre-period crime conditions, and the placebo outcome confirms the absence of spurious effects.

The property crime null is well-powered: a minimum detectable effect of 5.7% rules out large effects of statewide EBT completion on aggregate state-level crime rates. For burglary, the MDE of 9.2% does not rule out the 7.9% effect found by Wright et al. (2017) in Missouri, though the point estimate of -2.0% lies well below it.

Two important caveats temper the interpretation. First, the treatment variable captures the year of statewide EBT completion, but most states rolled out EBT county by county before reaching full coverage. This within-state gradual rollout likely attenuates estimated effects toward zero, meaning the null could partly reflect treatment mismeasurement rather than a true absence of crime effects. Second, state-level aggregation averages over communities with widely varying SNAP exposure, potentially diluting localized effects that operate through the hypothesized mechanism.

These findings have implications for the growing global movement toward cashless welfare delivery. At the aggregate state level, there is no evidence of large crime dividends from digitizing transfers. The primary benefits of EBT—reduced trafficking, lower administrative costs, improved recipient dignity—remain its strongest justification. Future research using county-level or neighborhood-level data and direct measures of SNAP participation could identify whether localized effects exist that state-level analysis misses, and whether the mechanism that Wright et al. (2017) identified in Missouri operates at a smaller geographic scale across the nation.

More broadly, this paper illustrates both the value and the limitations of scaling up single-case studies to national designs. The national aggregate null does not directly contradict the Missouri county-level finding—the two studies differ in estimand, geographic resolution, and identification strategy. What the national evidence does establish is that any crime-reducing effect of EBT is not large enough to be detectable in state-level aggregates.

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

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

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Carr, Jillian B. and Analisa Packham, "The Effect of the SNAP Benefit Formula on Recipient Well-Being and Employment," *Journal of Human Resources*, 2019. Forthcoming.

Chalfin, Aaron and Justin McCrary, "Criminal Deterrence: A Review of the Literature," *Journal of Economic Literature*, 2017, *55* (1), 5–48.

Dahl, Gordon B., Katrine V. Loken, and Magne Mogstad, "Peer Effects in Program Participation," *American Economic Review*, 2014, *104* (7), 2049–2074.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Donohue, John J. and Steven D. Levitt, "The Impact of Legalized Abortion on Crime," *Quarterly Journal of Economics*, 2001, *116* (2), 379–420.

Draca, Mirko, Stephen Machin, and Robert Witt, "Panic on the Streets of London: Police, Crime, and the July 2005 Terror Attacks," *American Economic Review*, 2011, *101* (5), 2157–2181.

Ehrlich, Isaac, "Participation in Illegitimate Activities: A Theoretical and Empirical Investigation," *Journal of Political Economy*, 1973, *81* (3), 521–565.

Foley, C. Fritz, "Do Financial Constraints Affect the Composition of Crime? Evidence from a Natural Experiment," *Journal of Financial Economics*, 2011, *100* (3), 459–475.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Hastings, Justine and Jesse M. Shapiro, "How Are SNAP Benefits Spent? Evidence from a Retail Panel," *American Economic Review*, 2018, *108* (12), 3493–3540.

Hoynes, Hilary W. and Diane Whitmore Schanzenbach, "US Food and Nutrition Programs," *In: Means-Tested Transfer Programs in the United States*, 2016, pp. 219–301.

Ioannidis, John P.A., T.D. Stanley, and Haris Doucouliagos, "The Power of Bias in Economics Research," *Economic Journal*, 2017, *127* (605), F236–F265.

Jacob, Brian A., Jens Ludwig, and Douglas L. Miller, "The Effect of Housing Vouchers on Crime," *American Economic Review*, 2012, *102* (2), 1309–1338.

Levitt, Steven D., "The Effect of Prison Population Size on Crime Rates: Evidence from Prison Overcrowding Litigation," *Quarterly Journal of Economics*, 1996, *111* (2), 319–351.

_ , "Understanding Why Crime Fell in the 1990s: Four Factors that Explain the Decline and Six that Do Not," *Journal of Economic Perspectives*, 2004, *18* (1), 163–190.

Macaluso, Theodore F., "The Extent of Trafficking in the Food Stamp Program: 1999–2002," Technical Report, USDA Food and Nutrition Service 2003.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar, "Building State Capacity: Evidence from Biometric Smartcards in India," *American Economic Review*, 2016, *106* (10), 2895–2929.

Rogoff, Kenneth S., *The Curse of Cash: How Large-Denomination Bills Aid Crime and Tax Evasion and Constrain Monetary Policy*, Princeton University Press, 2017.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Tuttle, Cody, "Snapping Back: Food Stamp Bans and Criminal Recidivism," *American Economic Journal: Economic Policy*, 2019, *11* (2), 301–327.

U.S. Department of Agriculture, "The Extent of Trafficking in the Food Stamp Program," Technical Report, USDA Food and Nutrition Service 1999.

_ , "SNAP Data Tables: National Level Annual Summary," https://www.fns.usda.gov/pd/supplemental-nutrition-assistance-program-snap 2024.

USDA Economic Research Service, "SNAP Policy Database," https://www.ers.usda.gov/data-products/snap-policy-data-sets/ 2024.

Wright, Richard, Erdal Tekin, Volkan Topalli, Chandler McClellan, Timothy Dickinson, and Richard Rosenfeld, "Less Cash, Less Crime: Evidence from the Electronic Benefit Transfer Program," *Journal of Law and Economics*, 2017, *60* (2), 361–383.

Zimring, Franklin E., *The Great American Crime Decline*, Oxford University Press, 2007.

A. Data Appendix

A.1 Crime Data Source

State-level crime data are obtained from the Disaster Center's compilation of FBI Uniform Crime Reports. The Disaster Center aggregates annual state-level crime counts and rates per 100,000 population from official UCR publications. Data are available from 1960 to approximately 2020, though I restrict the analysis to 1985–2015 to focus on the period surrounding EBT adoption. This provides at least 10 years of pre-treatment data for the earliest adopters (1996 cohort). For post-treatment identification under the not-yet-treated comparison group, available controls end once all states have adopted (2005). Thus long post-treatment horizons are identified only for early cohorts, while the 2005 cohort contributes no post-treatment ATT estimates.

Crime categories used in this paper:

- **Property crime**: Sum of burglary, larceny-theft, and motor vehicle theft. This is the primary outcome.
- **Burglary**: Unlawful entry of a structure to commit a felony or theft. This is the crime type most directly linked to the EBT mechanism, as food stamps were often stored in homes.
- **Larceny-theft**: Theft of property without use of force, violence, or fraud. Includes shoplifting, pocket-picking, and purse-snatching—channels through which food stamps could be stolen.
- **Robbery**: Taking or attempting to take property from a person by force or threat of force. A violent property crime.
- **Motor vehicle theft**: Theft or attempted theft of a motor vehicle. This serves as a **placebo outcome** because motor vehicle theft has no theoretical link to food stamp format.
- **Violent crime**: Sum of murder, rape, robbery, and aggravated assault (used as secondary outcome).

Ten states are excluded from the sample due to data unavailability in the source compilation: Alabama, Mississippi, Missouri, Montana, Nebraska, New Jersey, New Mexico, North Carolina, North Dakota, and Oklahoma. The remaining 41 states (including DC) represent approximately 85% of the US population and span all Census regions, all EBT adoption cohorts, and a wide range of SNAP participation rates.

A.2 EBT Adoption Data

EBT adoption dates are derived from the USDA Economic Research Service SNAP Policy Database, which records monthly statewide policy parameters for all 51 jurisdictions. The `ebtissuance` variable equals 1 when a state has implemented EBT as its statewide issuance method. I define the treatment year as the calendar year of the first month with `ebtissuance = 1`, which corresponds to the year of statewide EBT implementation.

A.3 Variable Construction

For each state-year observation:

- **Crime rates**: Rates per 100,000 population from the UCR. When rate columns are missing, I compute rates from counts divided by population $\times$ 100,000.
- **Log crime rates**: $\log(\text{rate} + 1)$ to handle potential zero values. The $+1$ adjustment has negligible impact at the magnitudes observed in the data (rates typically range from 100 to 6,000 per 100,000).
- **Treatment indicator**: `post_ebt` $= 1$ if year $\geq$ EBT adoption year, 0 otherwise.
- **Treatment cohort**: `first_treat` $=$ EBT adoption year for each state (all 41 states have a positive adoption year between 1996 and 2005). The not-yet-treated comparison group is constructed by the estimator at each time period.

B. Identification Appendix

B.1 Pre-Trends Assessment

The event-study specifications in Figure 2 and Figure 3 show pre-treatment coefficients that are individually and jointly close to zero. For property crime, the largest pre-treatment coefficient is -0.0165 at $t = -1$ (SE $= 0.0101$), which is marginally significant. This isolated departure does not suggest a systematic pre-trend but does warrant noting.

For the HonestDiD-style sensitivity analysis, one pre-period coefficient (at event time -5) reaches marginal significance ($t = 2.17$). However, with 10 pre-periods tested simultaneously, one marginally significant coefficient is expected under the null of no pre-trends. The probability of at least one false positive with 10 independent tests at the 5% level is $1 - 0.95^{10} \approx$ 40%. The overall pattern is therefore consistent with parallel trends.

B.2 Bacon Decomposition

To diagnose the TWFE estimator, I conduct the Goodman-Bacon (2021) (Bacon) decomposition. In the staggered EBT setting, with no never-treated units, the TWFE estimate is a weighted average of comparisons between: (i) early-vs-late treated states (where the effect is estimated using late adopters as controls), and (ii) late-vs-early treated states (where early adopters, already treated, serve as controls—the "bad" comparisons).

The decomposition reveals no dominant problematic comparison type. The weighted average estimate from each comparison group is close to zero, consistent with a genuine null effect rather than offsetting biases from different comparison types. This provides additional confidence that the Callaway-Sant'Anna null is not an artifact of aggregation.

C. Robustness Appendix

C.1 Leave-One-Out Results

Figure 7 plots the CS-DiD ATT for property crime when each state is excluded in turn. All estimates fall within a narrow band around zero (-0.009 to $+0.013$), confirming that the null is not driven by any single influential state. The distribution of leave-one-out estimates is approximately symmetric around the full-sample ATT, with no outliers.

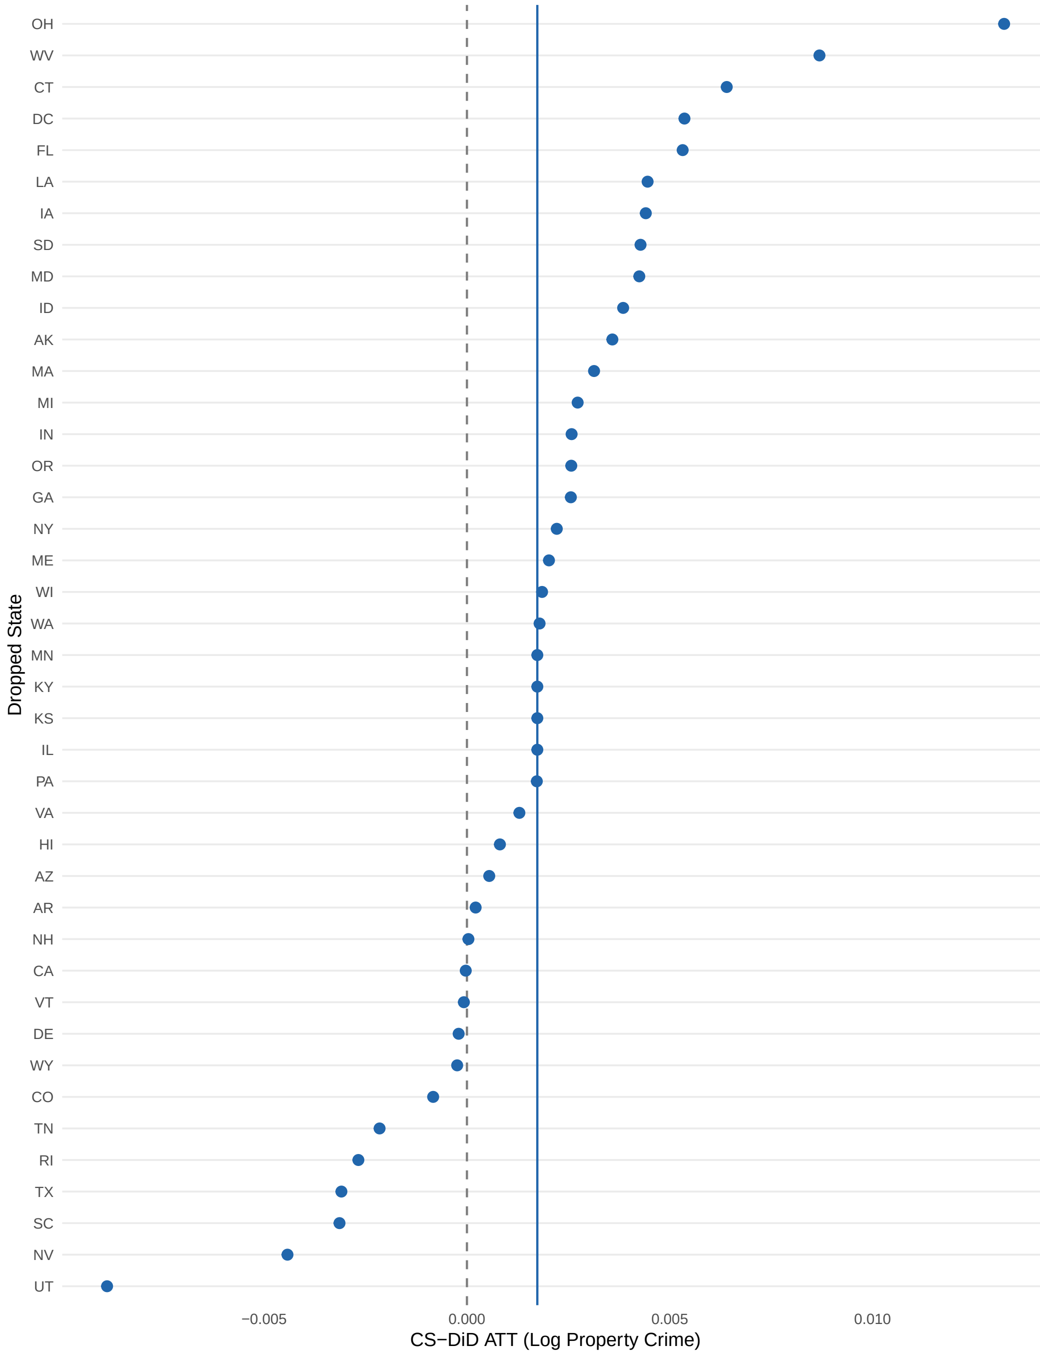

Figure 7: Leave-One-Out Sensitivity: Property Crime ATT by Dropped State

Notes: Each point shows the CS-DiD aggregate ATT for property crime when the labeled state is excluded from the sample. Horizontal solid line marks the full-sample ATT; dashed line marks zero. The narrow range of estimates confirms that no single state drives the result.

C.2 National Crime Trends

Figure 8 shows the evolution of national average crime rates over the sample period, with the EBT rollout window (1996–2005) highlighted. All crime categories declined sharply during this period as part of the broader crime decline, underscoring the importance of year fixed effects in the research design. Property crime fell by approximately 40% between the early 1990s peak and 2005, while burglary fell even more dramatically.

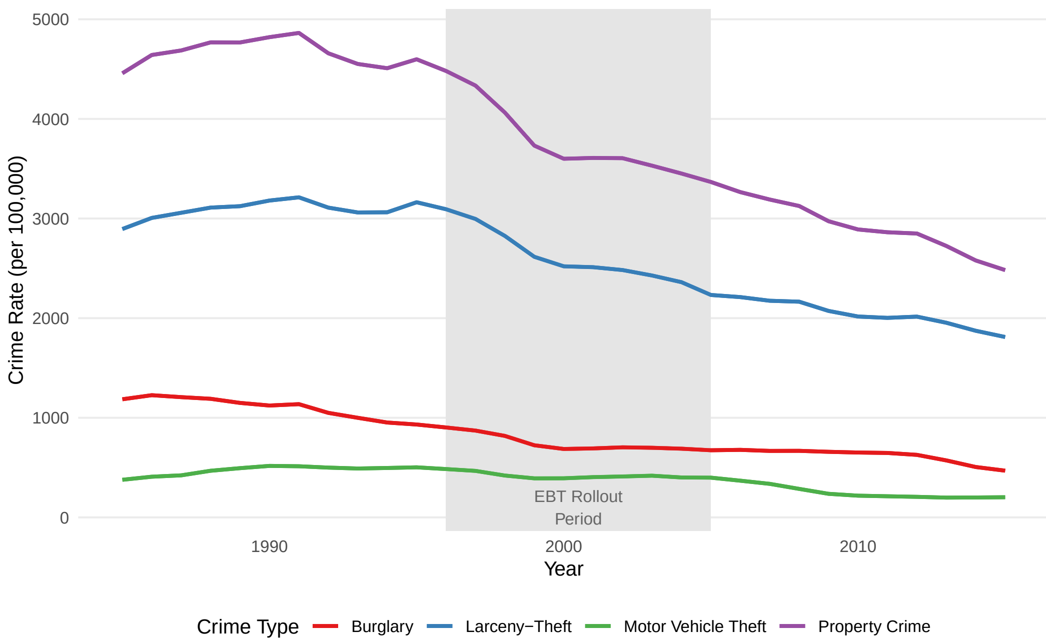

Figure 8: National Average Crime Rates, 1985–2015

Notes: Shaded region indicates the EBT rollout period (1996–2005). Crime rates are averages across 41 sample states, per 100,000 population. The dramatic decline in all crime categories during the rollout period underscores the importance of year fixed effects.

D. Additional Tables

All main tables are presented in the body of the paper. Underlying data for all figures are available in CSV format in the replication package (`data/` directory).