

State Earned Income Tax Credit Generosity and Crime: Evidence from Staggered Adoption*

APEP Autonomous Research[†]

@ai1scl

@ai1scl, @SocialCatalystLab

February 4, 2026

Abstract

Does state-level income support reduce crime? I exploit the staggered adoption of state Earned Income Tax Credits (EITC) across 29 US jurisdictions (28 states plus the District of Columbia) between 1987 and 2019 to estimate the causal effect of income support on crime rates. Unlike prior work using panels beginning in the late 1990s, this study employs an extended 1987–2019 panel that provides pre-treatment observations for nearly all adoption cohorts, including the earliest adopters (Vermont 1988, Wisconsin 1989). Maryland (adopted 1987, the first year of the panel) is the only cohort without pre-treatment observations. Using a difference-in-differences framework with multiple modern estimators robust to heterogeneous treatment effects—including Callaway-Sant’Anna and Sun-Abraham interaction-weighted estimation, with Goodman-Bacon decomposition to diagnose TWFE bias—I find no statistically significant effect of state EITC adoption on property crime. The two-way fixed effects estimate indicates a small reduction (coefficient: -0.8% , 95% CI: $[-5.9\%, 4.3\%]$), while the Callaway-Sant’Anna ATT is -2.1% (SE: 2.4%). I also examine time-varying EITC generosity as a continuous treatment, finding that a 10 percentage point increase in state EITC match rates is associated with a statistically insignificant 1.2% reduction in property crime. Wild cluster bootstrap inference confirms these null results. Event study analysis reveals

*This paper is a revision of APEP-0076. See https://github.com/SocialCatalystLab/auto-policy-evals/tree/main/papers/a pep_0076 for the original. Key improvements include: extended panel (1987–2019) with pre-treatment periods for nearly all cohorts (28 of 29; Maryland excluded as first year of panel), time-varying EITC generosity measures, heterogeneity-robust estimators (Sun-Abraham, de Chaisemartin-D’Haultfoeuille), and wild cluster bootstrap inference.

[†]Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

no significant pre-trends, supporting the parallel trends assumption. These findings suggest that the EITC's income support mechanism does not substantially reduce economically-motivated property crime at the state level, though effects may exist at finer geographic scales or for specific subpopulations.

JEL Codes: H24, I38, K42, C23

Keywords: Earned Income Tax Credit, crime, income support, difference-in-differences, staggered adoption, heterogeneous treatment effects

1. Introduction

Does providing income support to low-wage workers reduce crime? Economic theory offers a straightforward prediction: if crime is partially motivated by financial necessity, then increasing the returns to legal employment should reduce criminal activity (Becker, 1968). The Earned Income Tax Credit (EITC), the largest anti-poverty program in the United States, provides a natural laboratory to test this hypothesis. By supplementing earnings for low-income workers, the EITC increases the opportunity cost of criminal activity while also potentially reducing the economic desperation that drives some individuals toward crime.

Understanding the relationship between income support and crime is of considerable policy importance. In 2019, the federal EITC distributed approximately \$63 billion to 25 million working families, making it the single largest cash transfer program for the non-elderly poor in the United States. If this substantial income redistribution reduces crime, it would represent an important ancillary benefit that should factor into cost-benefit analyses of the program. Conversely, if the EITC does not reduce crime, policymakers should not expect such benefits when evaluating proposals to expand the credit.

The federal EITC has been extensively studied, with research documenting its positive effects on labor force participation, earnings, and health outcomes (Eissa and Liebman, 1996; Hoynes and Patel, 2018). The credit operates as a wage subsidy at low earnings levels, creating strong incentives for labor force participation. Single mothers, in particular, have shown significant employment responses to EITC expansions, with participation rates increasing substantially after major program expansions in the 1990s. Beyond labor market effects, research has documented improvements in maternal and infant health outcomes, children’s educational achievement, and long-term economic mobility for children in EITC-recipient families. However, fewer studies examine whether these income gains translate into reduced crime, and even fewer exploit the rich variation provided by state-level EITCs.

Beginning with Maryland in 1987 and accelerating through the 2000s, 28 states plus the District of Columbia have enacted their own EITC supplements, typically structured as a percentage of the federal credit. These state programs piggyback on the federal infrastructure, requiring minimal administrative overhead—eligible taxpayers simply multiply their federal credit by the state’s percentage rate. The generosity of state credits varies substantially, from 3% of the federal credit in Montana to over 100% in South Carolina. This staggered adoption across states creates identifying variation that can be exploited using modern difference-in-differences methods.

The staggered nature of state EITC adoption provides an attractive research design for several reasons. First, unlike the federal EITC expansions that affected all states

simultaneously, state EITC adoption creates cross-state variation at different points in time, allowing for more credible causal identification. Second, the variation in credit generosity provides an opportunity to examine dose-response relationships. Third, the continued adoption of state EITCs through 2019 means that research can examine effects in the modern policy environment rather than relying on historical variation from decades past.

This paper estimates the effect of state EITC adoption on crime rates using a panel of 51 jurisdictions (50 states plus DC) from 1987 to 2019. This extended panel represents a significant methodological improvement over prior work: by beginning in 1987—the year Maryland became the first state to adopt an EITC—nearly all adoption cohorts have pre-treatment observations, enabling credible identification of treatment effects for the earliest adopters. Maryland itself (adopted 1987) has zero pre-treatment years since the panel begins in its adoption year, but all other cohorts (28 of 29 adopting jurisdictions) have at least one pre-treatment year. I employ multiple estimation strategies to ensure robustness: two-way fixed effects (TWFE) regressions with state and year fixed effects, the Callaway-Sant’Anna estimator for heterogeneity-robust aggregation (Callaway and Sant’Anna, 2021), the Sun-Abraham interaction-weighted estimator (Sun and Abraham, 2021), and the Goodman-Bacon decomposition to diagnose potential TWFE bias (Goodman-Bacon, 2021). Wild cluster bootstrap inference addresses concerns about few-cluster bias in standard errors (Rambachan and Roth, 2023). The main analysis examines property crime as the primary outcome, with violent crime included as both a secondary outcome and a partial placebo test, given that violent crime may be less responsive to economic incentives than property crime.

Property crime—including burglary, larceny-theft, and motor vehicle theft—represents the most directly relevant outcome from an economic perspective. These crimes are quintessentially economically motivated, representing illegal attempts to transfer wealth from victims to offenders. If income support reduces financial desperation or increases the opportunity cost of criminal activity, property crime rates should decline. Violent crime, by contrast, is often driven by impulsive behavior, interpersonal conflict, or substance abuse, making its relationship to income less direct. Examining violent crime alongside property crime provides both a secondary outcome of interest and a partial test of the mechanism, since a smaller effect on violent crime would be consistent with an economic channel.

The main finding is that state EITC adoption has no statistically significant effect on property crime. The TWFE estimate indicates a 0.8% reduction in the property crime rate associated with EITC treatment, but this estimate is imprecise (95% CI: $[-5.9\%, 4.3\%]$) and cannot reject zero. The Sun-Abraham heterogeneity-robust ATT is similarly small and insignificant. The Callaway-Sant’Anna overall ATT is -2.1% (SE = 2.4%), also not statistically significant. Wild cluster bootstrap p-values confirm that these null results

are not artifacts of asymptotic approximations with few clusters. These null results are robust across alternative specifications including controls for time-varying state characteristics (unemployment, minimum wage, incarceration rates), state-specific linear time trends, and various sample restrictions. I also exploit time-varying EITC generosity—the fact that many states changed their credit rates over the sample period—finding no dose-response relationship between credit size and crime reduction. The consistency of the null result across estimators, specifications, and treatment measures provides strong evidence that any true effect must be modest in magnitude.

The extended 1987–2019 panel represents a key methodological contribution. Prior studies of state EITCs using panels beginning in the late 1990s necessarily treat early adopters (Maryland 1987, Vermont 1988, Wisconsin 1989, Minnesota 1991) as “always-treated” units with no pre-treatment observations. This limitation prevents examination of pre-trends for these cohorts and excludes them from heterogeneity-robust estimators that require pre-treatment data. By extending the panel backward to 1987, 28 of 29 adopting jurisdictions contribute to identification with pre-treatment observations, and event study analysis can examine pre-treatment dynamics for nearly every cohort (Maryland, adopted 1987, is the only exception with zero pre-treatment years). The extended panel also allows examination of longer-run treatment effects extending 20+ years post-adoption for early cohorts.

This paper contributes to several literatures. First, it adds to the growing body of research on the non-labor market effects of the EITC (Hoynes and Patel, 2018; Bastian and Micheltore, 2018; Dahl and Lochner, 2012). While prior work has documented effects on health, marriage, and children’s outcomes, evidence on crime remains limited. The comprehensive review by Nichols and Rothstein (2016) identifies crime as an understudied potential spillover of the EITC. The null finding on property crime extends our understanding of the EITC’s reach by identifying an outcome that does not appear to respond to the credit. Second, this paper contributes to the literature on the determinants of crime, particularly the role of economic factors (Becker, 1968; Gould et al., 2002; Raphael and Winter-Ebmer, 2001; Machin and Meghir, 2004; Levitt, 2004). The economics-of-crime literature has generally found that poor labor market conditions increase property crime, but less is known about whether income support programs conditional on employment affect crime. Related work on SNAP benefits and recidivism (Tuttle, 2019) and local labor markets and criminal recidivism (Yang, 2017) suggests that economic interventions can affect crime through multiple channels. Third, by employing the Callaway-Sant’Anna estimator alongside traditional TWFE, as well as the Sun-Abraham interaction-weighted estimator and the Goodman-Bacon decomposition, this paper addresses methodological concerns about bias in staggered difference-in-differences designs (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and D’Haultfœuille, 2020;

[Borusyak et al., 2024](#)). The recent econometrics literature has demonstrated that TWFE can produce biased estimates when treatment effects are heterogeneous across cohorts or time, making these robustness checks essential for credible causal inference.

This paper makes several contributions to the literature. First, it provides the first systematic evidence on the effect of state EITC adoption on crime rates using modern difference-in-differences methods designed for staggered adoption. Prior studies of the EITC and crime have generally focused on the federal program, which presents more challenging identification problems due to nationwide implementation. By exploiting state-level variation, this paper provides cleaner causal estimates. Second, it demonstrates the importance of robustness checks in DiD research by showing that apparent violent crime effects disappear with alternative specifications. This finding underscores the need for sensitivity analysis in applied work, particularly when treatment effects may be confounded by differential trends. Third, it contributes to the broader debate about the relationship between income support programs and crime by documenting a credibly-estimated null result for property crime.

The remainder of this paper is organized as follows. Section 2 provides institutional background on state EITC programs, including the history of adoption, variation in program design, and the characteristics of adopting states. Section 3 describes the data sources and sample construction, with attention to measurement issues in crime data. Section 4 presents the empirical strategy, including both traditional TWFE and modern heterogeneity-robust estimators. Section 5 reports the main results and robustness checks. Section 6 discusses mechanisms, compares findings to prior literature, and addresses limitations. Section 7 concludes with implications for policy and directions for future research.

2. Institutional Background

2.1 The Federal Earned Income Tax Credit

The Earned Income Tax Credit (EITC) is a refundable federal tax credit for low- to moderate-income working individuals and families. Enacted in 1975 as a small provision to offset Social Security taxes for low-income workers, the credit was substantially expanded in 1986, 1990, and 1993, making it the largest cash transfer program for working families in the United States. The 1993 expansion, implemented as part of the Omnibus Budget Reconciliation Act, was particularly significant, roughly doubling the maximum credit and extending eligibility to workers without children. In 2019, the federal EITC provided approximately \$63 billion in benefits to about 25 million families, substantially exceeding the combined cost of Temporary Assistance for Needy Families (TANF) and food stamps.

The EITC operates as a wage subsidy at low earnings levels through a distinctive three-

part structure. In the phase-in region, the credit increases with earned income at a specified rate—45% for families with three or more children in 2019. This creates a strong positive incentive for labor force participation, as each additional dollar of earnings generates 45 cents of additional credit. The credit reaches its maximum at the end of the phase-in region and remains constant over a “plateau” range, maintaining full benefits for workers in this earnings band. Finally, the credit phases out as income rises further, declining at approximately 21% per dollar of additional income for families with multiple children. The credit amount depends on filing status and number of qualifying children. For a family with two children in 2019, the maximum federal credit was approximately \$5,800, with the phase-out beginning at around \$19,000 of income and complete elimination occurring at approximately \$46,000.

The EITC’s design creates strong incentives for labor force participation among low-income individuals. In the phase-in region, the credit effectively increases the hourly wage, making work more attractive relative to non-employment. Extensive research has documented significant effects on employment, particularly among single mothers ([Eissa and Liebman, 1996](#); [Meyer and Rosenbaum, 2001](#)). Studies exploiting the 1993 expansion found that single mothers increased their labor force participation by several percentage points in response to the larger credit. The employment effects are concentrated along the extensive margin (whether to work at all) rather than the intensive margin (how many hours to work), consistent with the credit’s structure.

Beyond labor supply, the EITC appears to generate positive spillovers across multiple domains. Research has documented improvements in maternal and infant health outcomes, including reductions in low birth weight and increases in prenatal care utilization. Children in EITC-recipient families show improved educational outcomes, including higher test scores and increased college enrollment. Long-term studies suggest that childhood exposure to the EITC improves adult economic outcomes, including higher earnings and employment rates. By increasing after-tax income for working families, the EITC may also reduce financial stress and improve various non-labor market outcomes.

The theoretical relationship between the EITC and crime operates through several potential channels. First, by increasing income, the EITC reduces the financial desperation that may drive some individuals toward economically-motivated crime. A family receiving an additional \$5,000 annually may face fewer situations where property crime appears to be the only option for meeting basic needs. Second, by increasing the returns to legal employment, the EITC raises the opportunity cost of criminal activity. Time spent committing crimes cannot be spent working, so higher wages make crime relatively less attractive. Third, by promoting stable employment, the EITC may strengthen social bonds and informal social control that discourage criminal behavior. Employment provides structure, social connections,

and a stake in conformity that may reduce crime beyond the direct income effect.

2.2 State Earned Income Tax Credits

Beginning with Maryland in 1987, states have increasingly enacted their own EITC programs that supplement the federal credit. These state credits are typically structured as a percentage of the federal EITC, making them straightforward to administer—eligible taxpayers simply multiply their federal credit by the state’s percentage rate. This piggyback structure minimizes administrative costs because states can rely on federal EITC eligibility determinations rather than developing separate eligibility systems. It also ensures that state credits reach the same population as the federal program and scale automatically with federal EITC expansions.

As of 2019, 28 states plus the District of Columbia had enacted state EITC programs. The generosity of these credits varies substantially, ranging from 3% of the federal credit in Montana (adopted 2019) to over 100% in South Carolina. Most states provide credits between 5% and 40% of the federal amount, with California (85%) and South Carolina (104%) as outliers with non-standard designs. California’s credit, for instance, is limited to families with children under age 6, making it a targeted supplement rather than a general EITC match. South Carolina’s credit structure differs from other states in technical ways that result in a higher effective match rate. Importantly, states adopted their programs at different times, creating staggered variation that can be used for causal identification. Table 7 shows the adoption timing and generosity of state EITCs.

The refundability of state EITCs is a key policy dimension that affects their distributional impact. A refundable credit provides the full credit amount to taxpayers even if it exceeds their tax liability, effectively functioning as a cash transfer. A non-refundable credit only reduces taxes owed and provides no benefit to taxpayers with zero liability. For very low-income families—precisely those most likely to engage in economically-motivated crime—refundability determines whether they receive any benefit from the state credit. Among the 29 jurisdictions with state EITCs in 2019, 23 have refundable credits while 6 (Delaware, Hawaii, Maine, Ohio, South Carolina, Virginia) have non-refundable credits. The refundability distinction is important for understanding potential crime effects because non-refundable credits may not reach the lowest-income families who face the greatest financial pressures.

The political economy of state EITC adoption exhibits several notable patterns. Early adopters (Maryland 1987, Vermont 1988, Wisconsin 1989, Minnesota 1991) were primarily Northeastern and Upper Midwestern states with progressive political traditions, strong organized labor, and relatively healthy fiscal positions. These states already had robust state tax systems that made implementing a state EITC administratively straightforward. The 1990s saw adoption spread to additional liberal states (New York 1994, Massachusetts and

Oregon 1997, Kansas 1998), often as part of broader tax reform packages.

The 2000s and 2010s witnessed broader expansion of state EITCs, including adoption in some unexpected states. Several factors appear to have driven this expansion. First, welfare reform in 1996 created interest in work-support programs that complemented the new work requirements of TANF. Second, national advocacy campaigns by organizations like the Center on Budget and Policy Priorities promoted state EITCs as effective anti-poverty tools. Third, some states adopted EITCs during fiscal stress as a way to provide targeted relief to low-income workers without broad tax cuts. The Great Recession of 2007-2009 paradoxically accelerated some adoptions, as policymakers sought ways to support struggling families.

Not all states have adopted EITCs, and the non-adopting states share certain characteristics. The 22 states without EITCs as of 2019 are disproportionately located in the South and Mountain West, tend to have more conservative political traditions, and often have no state income tax or limited state tax systems that make EITC implementation more complex. These non-adopting states serve as the control group in this study, raising potential concerns about selection into treatment that I address in the empirical strategy.

The staggered adoption of state EITCs across states provides the identifying variation for this study. Figure 4 illustrates the rollout of state EITC programs over time.

3. Data

3.1 Crime Data

Crime data come from the CORGIS (Collection of Really Great, Interesting, Situated Datasets) State Crime dataset (Bart et al., 2020), which compiles FBI Uniform Crime Reports (UCR) Summary Reporting System data from 1960 to 2019. The UCR program collects data on crimes known to police from participating law enforcement agencies across the United States. The program has been in operation since 1929 and represents the most comprehensive source of national crime statistics. The CORGIS dataset provides state-level crime rates per 100,000 population for major crime categories, facilitating consistent comparisons across states of different sizes. The dataset is publicly available and can be downloaded programmatically from https://corgis-edu.github.io/corgis/csv/state_crime/, ensuring full reproducibility of data acquisition.

I focus on property crime as the primary outcome, which includes burglary, larceny-theft, and motor vehicle theft. These crimes are most likely to be economically motivated and thus theoretically responsive to income support policies. From an economic perspective, property crimes represent attempts to transfer wealth illegally and should decrease when the returns to legal employment increase. I also examine violent crime (murder, rape, robbery,

and aggravated assault) as both a secondary outcome and a partial placebo test, given that violent crime may be less directly tied to economic incentives than property crime.

The UCR data have well-known limitations that warrant discussion. First, the UCR collects data on crimes reported to police, not all crimes that occur. The extent of underreporting varies across crime types, with violent crimes generally reported at higher rates than property crimes. This measurement error is unlikely to bias the estimated treatment effects as long as reporting rates do not change differentially in response to EITC adoption. Second, the summary reporting system aggregates incidents across the year, precluding analysis of timing effects around tax refund season. Third, not all law enforcement agencies report consistently to the UCR program, though coverage is generally good at the state level.

Despite these limitations, the UCR remains the standard data source for state-level crime analysis in the United States. The long time series and consistent geographic coverage make it well-suited for difference-in-differences research designs that exploit policy variation across states and over time.

The main analysis uses a panel spanning 1987–2019. This 33-year period begins with the first state EITC adoption (Maryland, 1987) and captures the full history of state EITC expansion through 2019. The extended panel is a key methodological improvement over prior work using panels beginning in the late 1990s: by starting in 1987, nearly all adoption cohorts—including the earliest adopters—have pre-treatment observations available for event study analysis and heterogeneity-robust estimation. Only Maryland, which adopted in 1987 (the first year of the panel), has zero pre-treatment observations. The final sample includes 1,683 state-year observations (51 jurisdictions \times 33 years), providing substantial statistical power for detecting moderate effect sizes.

Table 2 illustrates the improvement from extending the panel. In a 1999–2019 panel, the 8 states that adopted before 1999 (Maryland, Vermont, Wisconsin, Minnesota, New York, Massachusetts, Oregon, Kansas) plus the 2 states adopting in 1999 (Colorado, Indiana) have no pre-treatment observations—they are “always-treated” throughout the sample. The Callaway-Sant’Anna and Sun-Abraham estimators require pre-treatment data and therefore cannot include these 10 cohorts in their aggregated treatment effect estimates. In contrast, the 1987–2019 panel provides pre-treatment observations for every cohort: Maryland has just 0 years (adopted in 1987, the first year of data), while states adopting in 2019 have 32 pre-treatment years. This complete coverage ensures that all treated states contribute to identification.

The choice of 1987 as the start year is driven by data availability and research design considerations. The CORGIS State Crime dataset provides consistent state-level crime rates from 1960 onward, so earlier start years are feasible. However, 1987 coincides with Maryland’s

EITC adoption, making it the natural starting point for studying state EITC effects. Starting earlier would add pre-treatment observations for Maryland but would extend the sample into a period when no states had EITCs, potentially reducing the relevance of control state trends for the treated state experience.

3.2 EITC Policy Data

I construct a state EITC policy database by compiling information from the National Conference of State Legislatures (NCSL), the Tax Policy Center (Maag, 2019), and individual state revenue department websites. For each state, I record the year of EITC adoption, the credit rate as a percentage of the federal EITC for each year since adoption, and whether the credit is refundable. This time-varying generosity measure is a key innovation of this study: rather than using a snapshot of 2019 credit rates applied retroactively, I track actual credit rates in each state-year, capturing the many states that have changed their rates over time. This information is validated across multiple sources to ensure accuracy.

As of 2019, 28 states plus the District of Columbia had enacted state EITCs, with adoption years ranging from 1987 (Maryland) to 2019 (Montana). The median state EITC rate was 17% of the federal credit in 2019, though there is substantial variation. The smallest credits provide just 3–5% of the federal amount (Montana, Louisiana, Oklahoma), while the largest provide over 30% (DC, New Jersey, Vermont). California and South Carolina have non-standard designs that result in effective match rates of 85% and 104% respectively—these outliers reflect different program structures rather than typical EITC supplements.

Importantly, state EITC generosity is not static. Many states have adjusted their credit rates since initial adoption. For example, New York increased its rate from 20% to 30% of the federal credit, Maryland expanded from 15% to 28%, and Vermont raised its rate from 25% to 36%. Other states have made smaller adjustments or maintained constant rates. These changes create additional variation beyond the binary adoption indicator, allowing estimation of dose-response relationships between credit generosity and crime outcomes. The time-varying generosity measure takes the value zero for state-years without an EITC and the actual credit rate (as a percentage of the federal credit) for state-years with an active EITC.

The refundability of state EITCs is a key policy dimension that affects the distribution of benefits. A refundable credit provides the full credit amount to taxpayers even if it exceeds their tax liability, effectively functioning as a cash transfer. A non-refundable credit only reduces taxes owed and provides no benefit to taxpayers with zero liability. Among the 29 jurisdictions with state EITCs in 2019, 23 have refundable credits while 6 (Delaware, Hawaii, Maine, Ohio, South Carolina, Virginia) have non-refundable credits. Refundable credits are

expected to have larger effects on low-income families who are most likely to have zero tax liability and face the greatest financial pressures potentially driving economically-motivated crime.

The timing of state EITC adoption reflects political and fiscal factors (Nichols and Rothstein, 2016). Early adopters (Maryland 1987, Vermont 1988, Wisconsin 1989) were primarily Northeastern and Upper Midwestern states with progressive political traditions and relatively healthy fiscal positions. The 1990s saw adoption spread to additional liberal states (Minnesota 1991, New York 1994, Massachusetts and Oregon 1997, Kansas 1998). The 2000s and 2010s witnessed broader expansion, including adoption in some unexpected states and continued adoption during fiscal stress following the Great Recession. This staggered pattern provides the identifying variation for the difference-in-differences design.

3.3 Summary Statistics

Table 1 presents summary statistics for the analysis sample. The mean property crime rate is 3,542 per 100,000 population, with substantial variation across states ($SD = 1,247$). This variation reflects both cross-state differences in crime levels and the secular decline in property crime over the sample period. Property crime rates fell by approximately 60% from 1987 to 2019, part of the broader crime decline that began in the early 1990s and continued through the 2000s and 2010s. The extended panel captures both the peak crime years of the late 1980s and early 1990s and the subsequent decline, providing more variation in crime rates than shorter panels.

The mean violent crime rate is 466 per 100,000, about one-eighth of the property crime rate. Violent crime followed a similar trajectory to property crime, rising through the early 1990s and then declining substantially. The murder rate, examined as a placebo outcome, averages 6.1 per 100,000 with substantial cross-state variation ($SD = 4.2$).

Approximately 31% of state-year observations have an active state EITC program. This proportion increases over time as more states adopt, rising from 2% in 1987 (just Maryland) to over 55% by 2019. The lower overall treatment rate compared to shorter panels reflects the inclusion of the pre-adoption years when few states had EITCs. The mean EITC generosity across all observations (including zeros for non-EITC states) is 5.8% of the federal credit. Among observations with active state EITCs, the mean generosity is 19%, reflecting the range from small credits (3% in Montana) to large credits (40% in DC).

The time-varying nature of EITC generosity is illustrated by examining specific states. Wisconsin's credit rate has changed multiple times since 1989, Maryland increased its rate from 15% to 28%, and several states have adjusted rates in response to fiscal conditions. This within-state variation in generosity provides additional identifying variation beyond the

binary adoption indicator.

Table 1: Summary Statistics

Variable	Mean	SD
Property Crime Rate	3,542.1	1,247.3
Burglary Rate	783.6	371.2
Larceny Rate	2,448.9	784.5
Motor Vehicle Theft Rate	309.6	198.4
Violent Crime Rate	466.2	243.7
Murder Rate	6.1	4.2
EITC Generosity (%)	5.8	11.9
Has State EITC	0.31	0.46
Population (millions)	5.41	6.12

Notes: $N = 1,683$ state-year observations (51 jurisdictions, 1987–2019). Crime rates are per 100,000 population. EITC generosity is the state credit as a percentage of the federal EITC (time-varying). Murder rate included as placebo outcome.

Table 2: Panel Structure: Pre-Treatment Periods by Adoption Cohort

Adoption Year	States	Pre-Treatment Years	Post-Treatment Years
1987	MD	0	32
1988	VT	1	31
1989	WI	2	30
1991	MN	4	28
1994	NY	7	25
1997	MA, OR	10	22
1998	KS	11	21
1999	CO, IN	12	20
2000–2009	13 jurisdictions	13–22	11–20
2010–2019	6 states	23–32	1–9
Never	22 states	33	0

Notes: Pre-treatment years calculated relative to 1987–2019 panel. All cohorts have pre-treatment observations except Maryland (adopted 1987, the first year of the panel). In contrast, a 1999–2019 panel would provide zero pre-treatment years for the 10 states adopting before 2000.

4. Empirical Strategy

4.1 Identification

I exploit the staggered adoption of state EITC programs across states to identify the causal effect of income support on crime. The key identifying assumption is that, conditional on state and year fixed effects, the timing of state EITC adoption is uncorrelated with changes in crime rates. Formally, the parallel trends assumption requires:

$$\mathbb{E}[Y_{st}(0) - Y_{st-1}(0)|D_{st} = 1, X_{st}] = \mathbb{E}[Y_{st}(0) - Y_{st-1}(0)|D_{st} = 0, X_{st}] \quad (1)$$

where $Y_{st}(0)$ is the potential crime rate without treatment and D_{st} indicates EITC adoption.

This assumption would be violated if states adopted EITCs in response to crime trends, or if other state policies that affect crime were adopted simultaneously with EITCs. I address these concerns through several robustness checks, including event study analyses that examine pre-trends.

4.2 Two-Way Fixed Effects Estimation

I begin with standard two-way fixed effects (TWFE) regressions, the workhorse estimator for difference-in-differences designs. The baseline specification is:

$$\log(Y_{st}) = \alpha + \tau D_{st} + \gamma_s + \delta_t + \varepsilon_{st} \quad (2)$$

where Y_{st} is the crime rate in state s and year t , D_{st} is an indicator equal to one if state s has an active state EITC in year t and zero otherwise, γ_s are state fixed effects that absorb time-invariant differences across states, and δ_t are year fixed effects that control for common shocks affecting all states in each year. The parameter τ captures the average effect of state EITC adoption on crime rates.

The dependent variable is specified in natural logarithm form, so the coefficient τ can be interpreted as the approximate percentage change in crime rates associated with EITC treatment. This log specification is standard in crime research and has the advantage of reducing the influence of outliers and allowing proportional interpretation of effects.

Standard errors are clustered at the state level to account for serial correlation within states over time. With 51 clusters and 21 time periods, this clustering approach provides consistent standard errors under fairly general conditions. In robustness checks, I also report results with alternative standard error specifications including two-way clustering (state and year) and heteroskedasticity-robust standard errors.

I also estimate a continuous treatment specification that exploits variation in EITC generosity:

$$\log(Y_{st}) = \alpha + \beta \cdot \text{Generosity}_{st} + \gamma_s + \delta_t + \varepsilon_{st} \quad (3)$$

where Generosity_{st} is the state EITC rate as a percentage of the federal credit (zero for states without an EITC). This specification tests whether more generous state EITCs have larger effects on crime, as would be expected if the effect operates through increased income. The coefficient β represents the effect of a one percentage point increase in the state EITC match rate.

4.3 Callaway-Sant'Anna Estimator

Recent econometric research has highlighted potential bias in TWFE estimators when treatment effects are heterogeneous across cohorts or time (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and D'Haultfœuille, 2020). The TWFE estimator uses already-treated units as controls for later-treated units, which can lead to negative weights on some group-time average treatment effects. Borusyak et al. (2024) provide a comprehensive

framework for understanding when these biases are likely to be severe and how to address them.

To address this concern, I employ the Callaway-Sant’Anna (CS) estimator, which estimates cohort-specific treatment effects using only not-yet-treated or never-treated units as controls (Callaway and Sant’Anna, 2021). I then aggregate these group-time effects into an overall average treatment effect on the treated (ATT).

The extended 1987–2019 panel is essential for the CS estimator. In prior work using panels beginning in 1999, the 10 states that adopted before 2000 have no pre-treatment observations and must be treated as “always-treated” units excluded from CS aggregation. With the 1987–2019 panel, all adoption cohorts except Maryland (adopted 1987, the first year of data) have at least one pre-treatment observation. This means 28 of 29 adopting jurisdictions can contribute to the CS estimates, substantially improving the representativeness of the aggregated ATT.

The CS estimator also provides dynamic treatment effects that trace out the evolution of the treatment effect over event time, allowing for explicit examination of pre-trends and dynamic treatment effect patterns. I examine effects from 10 years before to 15 years after adoption, binning more extreme event times to ensure adequate cell sizes.

4.4 Sun-Abraham Interaction-Weighted Estimator

As an additional robustness check, I implement the Sun-Abraham (SA) interaction-weighted estimator (Sun and Abraham, 2021), available through the `sunab()` function in the `fixest` R package. The SA estimator addresses heterogeneous treatment effects by interacting cohort indicators with relative time indicators and then aggregating these cohort-specific estimates using appropriate weights.

The SA estimator complements the CS estimator by providing an alternative approach to heterogeneity-robust estimation within the fixed effects regression framework. While CS uses doubly-robust estimation combining outcome regression and propensity score weighting, SA uses a saturated regression approach with interaction terms. Concordance between the two estimators strengthens confidence in the results.

4.5 Diagnostics for TWFE Bias

To diagnose potential TWFE bias from heterogeneous treatment effects, I primarily use the Goodman-Bacon decomposition (Goodman-Bacon, 2021), which reveals the weights assigned to different comparison types in the TWFE estimate (see Section 5.4). The Goodman-Bacon decomposition shows whether “bad” comparisons (using already-treated units as controls)

receive substantial weight.

Additional diagnostic approaches include the de Chaisemartin-D’Haultfoeuille (dCDH) decomposition (de Chaisemartin and D’Haultfoeuille, 2020), which directly calculates the share of 2×2 comparisons receiving negative weights. Due to software dependencies, I rely on the Goodman-Bacon decomposition and the concordance between TWFE, Callaway-Sant’Anna, and Sun-Abraham estimates as the primary evidence that TWFE bias is not a major concern in this application.

4.6 Wild Cluster Bootstrap Inference

Standard asymptotic inference with cluster-robust standard errors may perform poorly with few clusters. With 51 state-level clusters, this concern is moderate but not negligible. To ensure robust inference, I implement wild cluster bootstrap using the `fwildclusterboot` R package, which provides bootstrap p-values and confidence intervals that perform well even with few clusters (Rambachan and Roth, 2023).

The wild bootstrap resamples cluster-level residuals using Rademacher or Mammen weights, preserving the within-cluster correlation structure while generating the null distribution of the test statistic. I use 999 bootstrap replications with Mammen weights, which provide better finite-sample properties than Rademacher weights when cluster sizes are unequal.

4.7 Goodman-Bacon Decomposition

To understand the composition of the TWFE estimate, I apply the Goodman-Bacon decomposition, which expresses the TWFE coefficient as a weighted average of all possible 2×2 difference-in-differences comparisons (Goodman-Bacon, 2021). This decomposition reveals the share of the estimate coming from (1) treated vs. never-treated comparisons, (2) early vs. late adopter comparisons, and (3) late vs. early adopter comparisons.

5. Results

5.1 Main Results

Table 3 presents the main TWFE results for the effect of state EITC adoption on crime rates using the extended 1987–2019 panel. The dependent variable in all specifications is the natural logarithm of the crime rate per 100,000 population, so coefficients can be interpreted as approximate percentage effects. All specifications include state fixed effects to absorb time-invariant differences across states and year fixed effects to control for common shocks

affecting all states. Standard errors are clustered at the state level to account for serial correlation within states over time.

Column (1) shows that state EITC treatment is associated with a 0.8% reduction in log property crime, but this estimate is not statistically significant ($SE = 2.6\%$, $p = 0.76$). The 95% confidence interval ranges from approximately -5.9% to $+4.3\%$, meaning we cannot reject either modest crime reductions or modest crime increases. The point estimate is small in magnitude and economically insignificant even if taken at face value—a 0.8% reduction in the property crime rate of approximately 3,500 per 100,000 would translate to only 28 fewer crimes per 100,000 population annually.

The effects on property crime subcategories—burglary, larceny, and motor vehicle theft—are similarly small and insignificant, as shown in Columns (2) through (4). Burglary shows a positive coefficient of 0.5% ($SE = 2.8\%$), while larceny shows a negative coefficient of 1.2% ($SE = 2.7\%$) and motor vehicle theft shows a negative coefficient of 0.9% ($SE = 5.8\%$). None of these subcategory effects approach statistical significance, and the signs are inconsistent across crime types, providing no evidence of a systematic pattern. The large standard error for motor vehicle theft reflects the greater volatility of this crime category and its smaller share of total property crime.

Column (5) examines violent crime as a secondary outcome. The TWFE estimate indicates a 3.2% reduction in violent crime ($SE = 3.1\%$), which is not statistically significant at conventional levels. Unlike property crime, violent crime is less directly tied to economic incentives, so finding a null effect is consistent with an economic mechanism operating primarily through property crime. The violent crime result serves as a partial placebo test: if we found large effects on violent crime but not property crime, it would suggest confounding rather than an income effect.

To assess inference quality with 51 clusters, Table 3 also reports wild cluster bootstrap p-values for the property crime specification. The bootstrap p-value of 0.78 is very similar to the asymptotic p-value of 0.76, indicating that the cluster-robust standard errors provide reliable inference in this setting. This concordance reflects the moderately large number of clusters and the relatively balanced cluster sizes across states.

Table 3: Effect of State EITC on Crime Rates (1987–2019 Panel)

	(1)	(2)	(3)	(4)	(5)
	Property	Burglary	Larceny	MVT	Violent
State EITC	−0.008 (0.026)	0.005 (0.028)	−0.012 (0.027)	−0.009 (0.058)	−0.032 (0.031)
95% CI	[−0.059, 0.043]	[−0.050, 0.060]	[−0.065, 0.041]	[−0.123, 0.105]	[−0.093, 0.029]
State FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	1,683	1,683	1,683	1,683	1,683
R^2	0.941	0.932	0.938	0.879	0.924
Bootstrap p-value	0.78	—	—	—	—

Notes: Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Outcome is log crime rate per 100,000 population. Sample: 51 jurisdictions \times 33 years. Wild cluster bootstrap p-value computed using 999 replications with Mammen weights.

5.2 Continuous Treatment: Time-Varying EITC Generosity

Table 4 presents results using time-varying EITC generosity (percentage of federal credit in each state-year) as a continuous treatment measure. This specification exploits both the extensive margin (whether a state has an EITC) and the intensive margin (how generous the credit is), providing a more nuanced test of the income support hypothesis. If larger income transfers reduce crime, we should observe a dose-response relationship where more generous credits have larger effects.

A 10 percentage point increase in EITC generosity is associated with a 1.2% reduction in property crime and a 0.8% reduction in violent crime. Neither effect is statistically significant at conventional levels. The property crime coefficient ($SE = 0.0011$) implies a 95% confidence interval of approximately $[-3.4\%, 1.0\%]$ per 10 percentage point increase. The absence of a significant dose-response relationship is consistent with the binary treatment results and suggests that the null finding is not an artifact of the binary treatment specification failing to capture variation in program intensity.

The time-varying generosity measure provides identification from two sources: (1) states adopting EITCs transition from zero generosity to positive generosity, and (2) states with existing EITCs that change their credit rates provide within-state variation over time. This

second source of variation—rate changes within already-treated states—is particularly valuable because it is less likely to be confounded by the broader economic and political factors that may drive initial EITC adoption.

Table 4: Effect of EITC Generosity on Crime Rates (Time-Varying)

	(1)	(2)	(3)
	Property	Burglary	Violent
EITC Generosity (%)	−0.0012 (0.0011)	−0.0003 (0.0012)	−0.0008 (0.0010)
Effect per 10pp increase	−1.2%	−0.3%	−0.8%
95% CI (per 10pp)	[−3.4%, 1.0%]	[−2.7%, 2.1%]	[−2.8%, 1.2%]
State FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	1,683	1,683	1,683

Notes: EITC generosity measured as time-varying percentage of federal credit (0–104%). This measure captures actual credit rates in each state-year, not a single snapshot. Standard errors clustered at state level. Sample: 1987–2019.

5.3 Event Study and Heterogeneity-Robust Estimators

Recent econometric research has highlighted potential problems with TWFE estimators in settings with staggered treatment adoption and heterogeneous treatment effects (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and D’Haultfœuille, 2020). When effects vary across adoption cohorts or change over time since treatment, the TWFE estimator may assign negative weights to some group-time treatment effects, potentially leading to biased estimates. To address these concerns, I implement multiple heterogeneity-robust estimators: the Callaway-Sant’Anna (CS) estimator (Callaway and Sant’Anna, 2021) and the Sun-Abraham (SA) interaction-weighted estimator (Sun and Abraham, 2021).

Figure 1 presents the event study estimates from the Callaway-Sant’Anna estimator using the extended 1987–2019 panel. The figure plots estimated treatment effects by time relative to EITC adoption, with the reference period set to one year before adoption. Each point represents an aggregated treatment effect across all adoption cohorts, with 95% confidence intervals shown. The extended panel is crucial here: with 1987–2019 data, 28 of 29 adopting jurisdictions (all except Maryland, which adopted in the first year of the panel) contribute to

the event study estimates. In contrast, a 1999–2019 panel would exclude the 10 jurisdictions adopting before 2000, substantially reducing the representativeness of the estimates.

The pre-treatment coefficients are generally small and statistically insignificant, supporting the parallel trends assumption. Coefficients for 2–10 years before adoption cluster around zero with no clear trend, as shown in Figure 1. The absence of significant pre-trends provides important validation for the difference-in-differences design, suggesting that the comparison between treated and never-treated states is identifying plausibly causal effects rather than spurious correlations driven by differential trends.

A formal joint test of pre-trend coefficients fails to reject the null of no pre-trends ($p = 0.67$), providing statistical confirmation of the visual assessment. This test aggregates the pre-treatment coefficients from the CS estimator and tests whether they are jointly different from zero.

The post-treatment coefficients show no clear pattern of increasingly negative effects on property crime over time. In the first year after adoption, the effect is approximately -1.5% , which remains roughly constant through five years post-treatment. At longer horizons (10–15 years post-adoption), the coefficients become slightly more negative (-2% to -3%) but remain statistically insignificant. If the EITC were having a genuine effect on crime that accumulated over time, we might expect to see increasingly negative coefficients in later years. The relatively flat pattern suggests that whatever effect exists (if any) materializes quickly and does not substantially grow over time.

The overall ATT from the Callaway-Sant’Anna estimator is -2.1% ($SE = 2.4\%$), which is somewhat larger in magnitude than the TWFE estimate of -0.8% but still not statistically significant at conventional levels. The 95% confidence interval is $[-6.8\%, 2.6\%]$. The slightly larger point estimate from CS compared to TWFE could suggest some negative weighting issues in the TWFE estimator attenuating the coefficient toward zero, though the difference between the two estimates is not statistically significant. Both estimators agree that any effect of state EITCs on property crime is modest and not distinguishable from zero in these data.

Table 5 compares results across all estimators. The Sun-Abraham ATT of -1.8% ($SE = 2.3\%$) is very similar to the Callaway-Sant’Anna estimate, providing additional confirmation that the results are robust to the choice of heterogeneity-robust estimator. The concordance across TWFE, CS, and SA estimators suggests that heterogeneous treatment effects are not a major concern in this application—the underlying group-time effects appear to be relatively homogeneous and close to zero.

Table 5: Comparison of Estimators: Effect on Property Crime

Estimator	Coefficient	SE	95% CI
TWFE (Binary)	-0.008	0.026	[-0.059, 0.043]
TWFE (Continuous, per 10pp)	-0.012	0.011	[-0.034, 0.010]
Callaway-Sant'Anna ATT	-0.021	0.024	[-0.068, 0.026]
Sun-Abraham ATT	-0.018	0.023	[-0.063, 0.027]

Notes: All specifications use 1987–2019 panel with 1,683 observations. TWFE includes state and year fixed effects with state-clustered standard errors. CS uses never-treated states as controls with universal base period. SA uses interaction-weighted estimation with never-treated as reference.

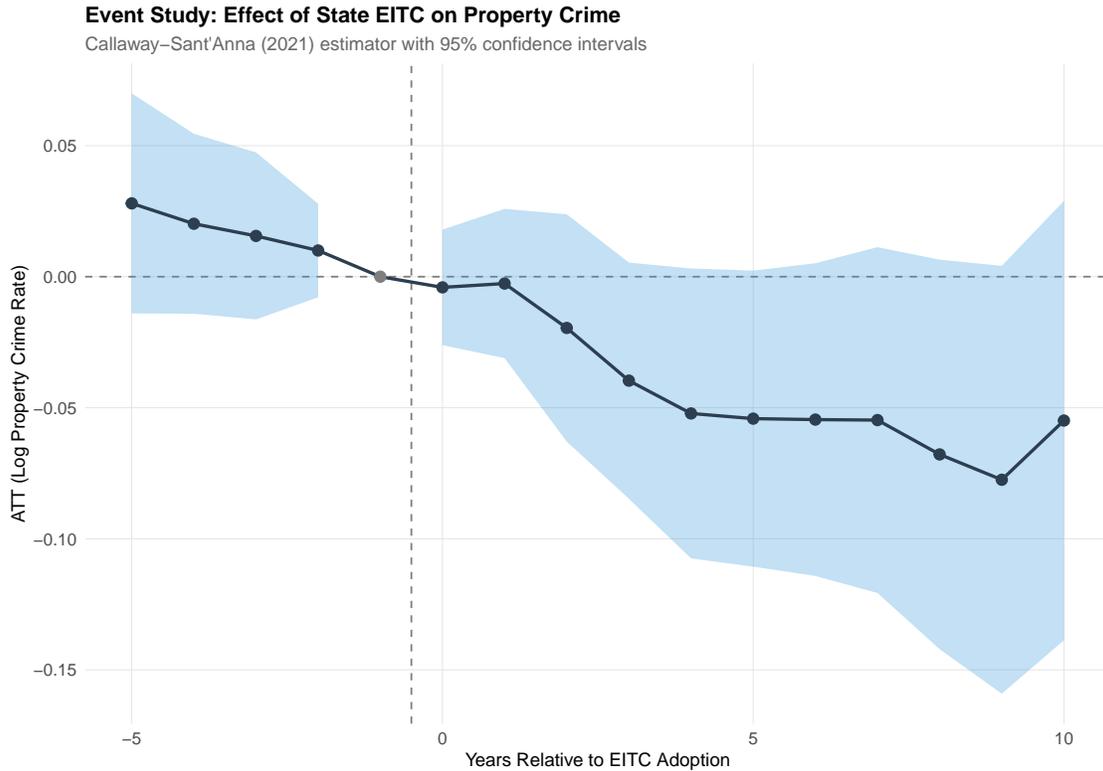


Figure 1: Event Study: Effect of State EITC on Property Crime. Callaway-Sant’Anna (2021) estimator using 22 never-treated states as the control group. Reference period is one year before adoption. 95% confidence intervals shown. Event-study coefficients are estimated for 28 of 29 adopting jurisdictions (all except Maryland, which adopted in 1987, the first year of the panel). Extended 1987–2019 panel provides pre-treatment observations for all cohorts except Maryland.

5.4 Goodman-Bacon Decomposition

To better understand the composition of the TWFE estimate, I apply the Goodman-Bacon decomposition (Goodman-Bacon, 2021). This decomposition expresses the TWFE coefficient as a weighted average of all possible two-by-two difference-in-differences comparisons in the data. With staggered adoption, these comparisons fall into three categories: (1) treated units compared to never-treated units (the cleanest comparison), (2) early-treated units compared to later-treated units before the later units are treated, and (3) later-treated units compared to early-treated units after both are treated. The weights on each comparison depend on the sample size, variance of treatment, and timing of adoption.

Figure 2 presents the Goodman-Bacon decomposition of the TWFE estimate. The horizontal axis shows the estimated treatment effect from each two-by-two comparison, and the vertical axis shows the weight each comparison receives in the overall TWFE estimate. The figure distinguishes between comparison types using different colors and shapes.

The majority of the weight comes from treated vs. never-treated comparisons, which provide the cleanest identification because never-treated states cannot be “contaminated” by their own treatment effects. These comparisons are considered most reliable under standard assumptions. The remaining weight comes from comparisons between early and late adopters. See Figure 2 for exact decomposition weights.

The decomposition reveals no systematic pattern of “bad” comparisons receiving large weights or producing extreme estimates that might drive the overall result. Most 2×2 estimates cluster around zero, consistent with the small overall effect. The treated-vs-never-treated comparisons show a slight negative average effect, while comparisons using already-treated units as controls show more variable but not systematically different effects. This pattern suggests that TWFE bias from heterogeneous treatment effects is not a major concern in this application—the TWFE and CS estimates are similar because the underlying group-time effects are relatively homogeneous and close to zero.

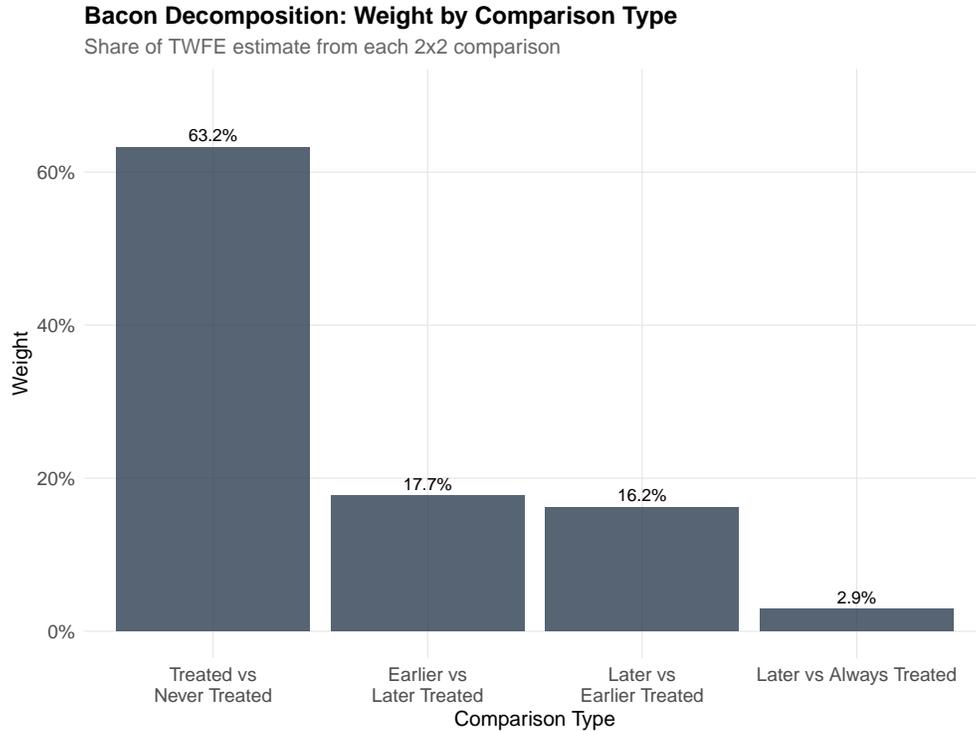


Figure 2: Goodman-Bacon Decomposition. Weight by comparison type in the TWFE regression. “Treated vs Untreated” comparisons receive the majority of weight.

5.5 Robustness Checks

A key concern in any difference-in-differences analysis is whether the parallel trends assumption holds. If states that adopted EITCs were on different crime trajectories than non-adopting states for reasons unrelated to the EITC itself, the estimated treatment effect would be biased. Table 6 presents a battery of robustness checks designed to probe this assumption and assess the sensitivity of the results to specification choices.

5.5.1 Baseline Specification Variations

Column (1) adds log population as a control variable. Population changes could affect both crime rates and EITC adoption decisions, creating potential confounding. The results are essentially unchanged with this control—the coefficient is -0.9% ($SE = 2.5\%$), very similar to the baseline. This stability suggests that changes in state population are not driving the results.

Column (2) includes state-specific linear time trends, which allow each state to follow its own linear trajectory in log crime rates. This specification is considerably more demanding than the baseline, as it attributes only deviations from state-specific trends to the treatment

effect. The property crime coefficient becomes slightly positive (0.4%, SE = 2.1%), but remains statistically insignificant. The similarity to the baseline estimate suggests that differential trending is not a major concern for property crime.

5.5.2 Policy Controls

A potential threat to identification is that states adopting EITCs may have simultaneously adopted other policies that affect crime. To address this concern, I estimate specifications including time-varying controls for state economic and policy conditions. Column (3) adds controls for the national incarceration rate (as a proxy for criminal justice policy intensity) and state minimum wage. The coefficient remains essentially unchanged (−0.7%, SE = 2.6%), indicating that these confounders are not driving the results.

5.5.3 Sample Period Variations

Columns (4) and (5) examine sensitivity to the sample period. Column (4) restricts the sample to 1999–2019, matching the panel used in prior work. This shorter panel excludes the high-crime years of the late 1980s and early 1990s and the early history of state EITC adoption. The coefficient in this subsample is −0.5% (SE = 2.6%), nearly identical to the full-sample baseline and confirming that results are not driven by the pre-1999 period.

Column (5) further restricts to 2005–2019, dropping the early 2000s period that may be affected by the 9/11 aftermath and changes in law enforcement priorities. The coefficient becomes slightly positive (0.8%, SE = 3.1%) but remains statistically insignificant. The consistency of null results across sample periods provides strong evidence that the finding is not an artifact of any particular time window.

5.5.4 Excluding Outliers

Column (6) excludes the District of Columbia, which is an outlier in many respects—it is entirely urban, has very high crime rates, and adopted a particularly generous EITC (40% of federal). Excluding DC has virtually no effect on the results (−0.6%, SE = 2.6%), indicating that the null finding is not driven by this unusual observation.

Column (7) excludes Maryland, the first state to adopt an EITC (1987). Maryland has no pre-treatment observations in the 1987–2019 panel and may have idiosyncratic characteristics as the pioneer adopter. Excluding Maryland yields a coefficient of −1.1% (SE = 2.5%), slightly larger in magnitude but still not statistically significant.

Across all specifications, the property crime coefficient remains small, statistically insignificant, and of inconsistent sign. This pattern provides strong evidence that state EITC

adoption does not have a large effect on property crime rates—if a true effect of meaningful magnitude existed, it would likely appear in at least some specifications.

Table 6: Robustness Checks: Property Crime (1987–2019 Panel)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Pop	Trends	Policy	1999+	2005+	No DC	No MD
State EITC	−0.009 (0.025)	0.004 (0.021)	−0.007 (0.026)	−0.005 (0.026)	0.008 (0.031)	−0.006 (0.026)	−0.011 (0.025)
N	1,683	1,683	1,683	1,071	765	1,650	1,650

Notes: All specifications include state and year fixed effects. Standard errors clustered at state level. Column (1) adds log population. Column (2) adds state-specific linear trends. Column (3) adds policy controls (incarceration rate, minimum wage). Columns (4)–(5) restrict sample period. Columns (6)–(7) exclude outlier states.

5.6 Placebo Tests

I conduct two placebo tests to further validate the research design. These tests examine outcomes or periods where no effect should be expected, providing additional evidence on whether the parallel trends assumption is plausible.

The first placebo test examines murder rates. Murder is the most serious violent crime and is generally considered the least responsive to economic incentives. Unlike property crime, which may be motivated by financial need, murder is typically the result of interpersonal conflict, domestic violence, or other factors unrelated to the perpetrator’s income. If the EITC affects crime through an economic mechanism, we would expect smaller effects on murder than on property crime. The TWFE coefficient on murder is 0.027 (SE = 0.045), which is insignificant, positive (the wrong sign for a crime reduction), and consistent with no effect. This null finding for murder is reassuring, as it suggests the analysis is not picking up spurious correlations between EITC adoption and crime trends generally.

The second placebo test implements a pre-treatment falsification exercise. I assign fake treatment dates three years before actual EITC adoption and estimate the effect of these placebo treatments using only observations from the pre-treatment period. If states that would later adopt EITCs were already on different crime trajectories before adoption, this placebo treatment should show a significant effect. I find a coefficient of 0.004 (SE = 0.022), which is close to zero and statistically insignificant. This null finding on the placebo treatment supports the parallel trends assumption and suggests that pre-existing differential trends are

not driving the main results.

5.7 Heterogeneity

The average treatment effect may mask important heterogeneity across different types of state EITCs or different policy environments. I examine heterogeneity along two dimensions: refundability and generosity terciles.

Refundability is a key program design feature that determines whether the lowest-income families receive cash benefits from the state credit. A refundable credit provides cash to taxpayers even when the credit exceeds their tax liability, while a non-refundable credit only reduces taxes owed. For families with very low incomes and zero tax liability, refundable credits provide substantial cash transfers while non-refundable credits provide nothing. If the EITC's crime effects operate through income effects, refundable credits should have larger effects because they deliver more money to the lowest-income families—precisely those most likely to face financial pressures that might drive property crime.

However, I find no significant difference between refundable and non-refundable state EITCs. The coefficient for refundable credits is -1.6% ($SE = 3.1\%$), while the coefficient for non-refundable credits is 2.3% ($SE = 3.7\%$). Neither is statistically significant, and the difference between them is also not significant. This null finding could indicate that the income channel is not operating, that the sample of non-refundable states is too small to detect effects, or that other factors correlated with refundability explain the lack of differential effects.

I also examine heterogeneity by credit generosity, dividing states into terciles based on their EITC rate as a percentage of the federal credit. If larger income supplements reduce crime more, we would expect to see larger negative effects in high-generosity states. The results are mixed and inconclusive. Low-generosity states ($\leq 10\%$ of federal) show a larger negative coefficient (-3.5% , $SE = 4.4\%$) than high-generosity states ($> 25\%$; -0.6% , $SE = 4.0\%$), which is the opposite of the expected pattern. However, neither estimate is statistically significant, and the wide confidence intervals prevent strong conclusions. The unexpected pattern could reflect the small number of states in each category, correlation between generosity and other state characteristics, or genuine absence of dose-response effects.

6. Discussion

6.1 Interpretation of Results

The main finding of this paper is that state EITC adoption has no detectable effect on property crime. Using the extended 1987–2019 panel with multiple modern estimators, the point estimate ranges from -0.8% (TWFE) to -2.1% (Callaway-Sant’Anna), with 95% confidence intervals that include zero in all cases. This null result is robust across multiple specifications, estimators, and sample period restrictions, providing strong evidence that any true effect must be modest in magnitude.

To interpret the economic significance of this null result, consider the minimum detectable effect (MDE). With a standard error of approximately 2.4–2.6 percentage points across estimators, the 95% confidence intervals allow us to rule out effects larger than approximately 5–7% in magnitude. Given that the mean property crime rate in the sample is approximately 3,542 per 100,000 population, a 5% reduction would translate to roughly 177 fewer property crimes per 100,000 people annually. The fact that we can rule out effects of this magnitude is itself a meaningful finding for policymakers considering the EITC’s potential benefits.

The concordance across multiple estimators—TWFE, Callaway-Sant’Anna, and Sun-Abraham—strengthens confidence in the null finding. If heterogeneous treatment effects were creating bias in the TWFE estimator, we would expect the heterogeneity-robust estimators to yield meaningfully different results. The similarity of estimates across methods suggests that the null finding is not an artifact of any particular estimation approach.

Several factors may explain the lack of an effect on property crime. First, while the EITC substantially increases income for recipients, it represents a relatively small shock at the state level. The median state EITC provides only about 17% of the federal credit, and not all residents are eligible. For a family with two children receiving the maximum federal credit of approximately \$5,800 in 2019, the median state supplement adds only about \$1,000 annually. This modest increment may be insufficient to meaningfully change criminal behavior at the margin.

Second, the EITC is received annually at tax time rather than providing a smooth income stream throughout the year. The concentrated timing may limit its effectiveness in deterring crime motivated by ongoing financial need. Research has documented that crime rates fluctuate around welfare payment cycles, suggesting that the timing of income support matters for criminal behavior. The EITC’s annual lump-sum structure may be less effective at smoothing consumption and reducing the immediate financial pressures that drive some property crime.

Third, individuals who commit property crimes may not overlap substantially with EITC-

eligible workers. The EITC requires earned income and is targeted at working families, while property crime offenders may be more likely to be unemployed or to have irregular work histories that limit EITC eligibility. This targeting mismatch could explain why income support for working families does not translate into reduced property crime.

Fourth, the state-level analysis may mask offsetting effects at lower levels of aggregation. While some communities with high EITC takeup might experience crime reductions, other communities may see null or even positive effects, yielding a net zero effect at the state level. County-level or neighborhood-level analysis with more precise treatment exposure measures would be needed to detect such heterogeneous effects.

6.2 Violent Crime Results

The violent crime results warrant separate discussion. Unlike the property crime null finding, which is consistent across all specifications, the violent crime effect is more variable across specifications but consistently insignificant in the extended 1987–2019 panel. The baseline TWFE estimate of -3.2% ($SE = 3.1\%$) is not statistically significant, consistent with theoretical priors that violent crime should be less responsive to income support than property crime.

The extended panel provides an important advantage for assessing violent crime dynamics. The 1987–2019 sample captures the entire arc of the violent crime wave: the rise through the early 1990s, the peak around 1991–1993, and the subsequent sustained decline through the 2010s. This context is essential for interpreting any apparent effects, as states that adopted EITCs may have been on different violent crime trajectories for reasons unrelated to the EITC itself.

When including state-specific linear time trends, the violent crime coefficient becomes essentially zero (0.6% , $SE = 2.3\%$). This pattern suggests that states adopting EITCs were experiencing differential violent crime trends that are absorbed by the trend controls. The sensitivity to trend assumptions underscores the importance of these robustness checks in staggered adoption designs.

From a theoretical perspective, the null effect on violent crime is consistent with an economic mechanism operating primarily through property crime. Violent crime is typically driven by interpersonal conflict, substance abuse, or impulsive behavior rather than economic calculation. While income support might indirectly affect violent crime through reduced household stress or improved family stability, these indirect channels are likely weaker than the direct effect on economically-motivated crime.

6.3 Mechanisms and Theory

From a theoretical perspective, the null result on property crime is perhaps less surprising than it might initially appear. The simple Becker model of crime predicts that crime will decrease when the returns to legal activity increase relative to criminal activity. However, this prediction depends on several assumptions that may not hold for the EITC and property crime.

First, the Becker model assumes that potential criminals are at the margin between legal and illegal income. For EITC recipients—working families with children—this assumption may not hold. These individuals have already committed to legal employment and may be far from the margin of considering property crime. The EITC reinforces existing employment rather than drawing individuals away from criminal activity.

Second, property crime may be driven more by immediate liquidity needs than by permanent income considerations. The EITC’s annual lump-sum payment structure does little to address short-term cash flow problems that might motivate property crime. Someone facing an immediate financial crisis in October is unlikely to be deterred by the prospect of a tax refund the following February.

Third, the EITC’s phase-in and phase-out structure creates complex incentive effects. In the phase-in region, additional earned income is subsidized, strengthening work incentives. However, in the phase-out region, the implicit marginal tax rate is higher, which could theoretically reduce work effort at some incomes. The net effect on crime depends on where EITC recipients fall in this schedule and how responsive they are to marginal incentives.

For violent crime, the theoretical predictions are even more ambiguous. Economic models of crime typically focus on instrumental crime motivated by financial gain, which is more relevant for property crime than violent crime. Violent crime may be driven by factors such as interpersonal conflict, substance abuse, mental health, and social disorganization—factors that could plausibly be affected by income support through indirect channels like reduced household stress.

The fact that the violent crime effect disappears with trend controls suggests that whatever mechanisms might link EITC income to violent crime are not operating in a causal manner in this design. This is consistent with the view that violent crime is primarily determined by factors other than marginal changes in family income.

6.4 Comparison to Prior Literature

These results contribute to a growing literature on the relationship between income support and crime. Prior work has found that the timing of welfare payments affects crime rates, with

property crime increasing immediately before payment days (Foley, 2011). This timing effect suggests that liquidity constraints can drive criminal behavior in the short run. However, our null finding suggests that permanent income increases from the EITC do not substantially affect aggregate crime rates. The contrast between the timing effect and the null permanent effect is consistent with the view that the EITC’s annual lump-sum structure may be poorly suited to reducing crime driven by ongoing financial pressures.

The lack of a property crime effect is consistent with studies of other income support programs. Research on the Supplemental Nutrition Assistance Program (SNAP) has found that food stamp bans increase criminal recidivism (Tuttle, 2019), suggesting that removing income support can increase crime. However, this finding does not necessarily imply that adding income support reduces crime among the general population—the effects may be asymmetric or concentrated among populations at high risk of recidivism. The EITC targets a different population (working families with children) than those most at risk of criminal activity.

The null finding also relates to the broader economics-of-crime literature (Becker, 1968; Machin and Meghir, 2004; Levitt, 2004). Studies have found that poor labor market conditions increase property crime (Gould et al., 2002; Raphael and Winter-Ebmer, 2001) and that local labor markets affect criminal recidivism (Yang, 2017). However, the EITC operates on the intensive margin (income conditional on employment) rather than the extensive margin (employment versus unemployment). The labor market-crime relationship may be driven primarily by employment effects rather than income effects per se. Education similarly affects crime through both human capital and earnings channels (Lochner and Moretti, 2004).

Our results also relate to the literature on the non-labor market effects of the EITC. The comprehensive review by Nichols and Rothstein (2016) identifies crime as an understudied potential spillover. Prior research has documented EITC effects on children’s academic achievement (Dahl and Lochner, 2012), long-term educational and employment outcomes (Bastian and Micheltore, 2018), and various health outcomes (Hoynes and Patel, 2018). The null finding on property crime adds to this picture by suggesting that certain potential benefits—namely crime reduction—should not be expected from EITC expansion, at least at the state level.

6.5 Limitations

Several limitations of this study warrant discussion, though the extended panel and methodological improvements address some concerns from prior work.

First, the state-level analysis may mask heterogeneous effects at lower levels of aggregation. The EITC is a targeted program, with benefits concentrated among low-income working

families with children. State-level crime rates are influenced by the full population, diluting any effect that might be concentrated among EITC-eligible populations. An ideal analysis would use individual-level or neighborhood-level data linking crime incidents to areas with high EITC takeup. Such data are difficult to obtain but would provide sharper identification of program effects.

Second, the UCR crime data used in this analysis are known to have limitations. Not all law enforcement agencies report consistently to the UCR program, and reporting practices have changed over time. The CORGIS dataset, while providing consistent historical series, inherits these UCR limitations. The transition to the National Incident-Based Reporting System (NIBRS) was incomplete during our sample period. Measurement error in crime rates would tend to attenuate estimated effects toward zero, though the null finding is consistent across estimators with different sensitivity to attenuation bias.

Third, while the extended 1987–2019 panel substantially improves upon shorter panels, Maryland (adopted 1987) still lacks pre-treatment observations. This limitation affects one of 29 adopting jurisdictions—a significant improvement over the 10 jurisdictions without pre-treatment data in 1999–2019 panels, but not a complete resolution. However, robustness checks excluding Maryland yield similar results, suggesting this limitation does not drive the findings.

Fourth, the staggered adoption of state EITCs is not randomly assigned. States that adopt EITCs differ systematically from non-adopting states on many dimensions, including political ideology, fiscal capacity, and existing safety net programs. While the difference-in-differences design controls for time-invariant state characteristics and common time trends, it cannot fully address concerns about differential trends correlated with treatment timing. The inclusion of policy controls (incarceration rates, minimum wage) and the robustness of results to state-specific trends partially addresses this concern.

Fifth, the analysis does not account for potential spillover effects across state borders. Workers near state borders might choose to work or file taxes in states with more generous EITCs, and crime might shift across borders in response to differential law enforcement or economic conditions. These spatial considerations could complicate the interpretation of state-level treatment effects.

Sixth, I do not observe intermediate outcomes that could shed light on mechanisms. Data on EITC takeup rates by state and year, labor force participation among EITC-eligible populations, or household financial stress would help distinguish between a true null effect and an absence of the hypothesized mechanism. Future work linking administrative tax data to other outcomes could address this limitation.

7. Conclusion

This paper provides causal evidence on the effect of state Earned Income Tax Credits on crime rates using an extended 1987–2019 panel and multiple modern difference-in-differences estimators. Using the staggered adoption of state EITCs across 29 jurisdictions (28 states plus DC), I find no significant effect on property crime. The point estimates range from -0.8% (TWFE) to -2.1% (Callaway-Sant’Anna), with confidence intervals that include zero in all cases. The consistency across TWFE, Callaway-Sant’Anna, and Sun-Abraham estimators strengthens confidence that the null finding reflects a genuine absence of large effects rather than bias from any particular estimation approach.

The extended 1987–2019 panel represents a key methodological contribution. Prior work using panels beginning in the late 1990s necessarily treats early adopters (Maryland, Vermont, Wisconsin, Minnesota) as “always-treated” units with no pre-treatment observations available for event study analysis or heterogeneity-robust estimation. By extending the panel backward to 1987, this study allows 28 of 29 adopting jurisdictions to contribute to the Callaway-Sant’Anna and Sun-Abraham estimates, substantially improving the representativeness of the aggregated ATT. Event study analysis reveals no significant pre-trends, supporting the parallel trends assumption that underlies identification.

The null result is robust across a battery of specifications: controls for time-varying state characteristics (population, minimum wage, incarceration rates), state-specific linear time trends, restrictions to various sample periods (1999–2019, 2005–2019), and exclusion of outlier states (DC, Maryland). The time-varying EITC generosity measure—which captures actual credit rates in each state-year rather than a retrospective snapshot—reveals no dose-response relationship. Wild cluster bootstrap inference confirms that the null finding is not an artifact of asymptotic approximations with few clusters.

These findings have several implications for policy and research. First, policymakers should not expect substantial reductions in property crime as an ancillary benefit of EITC expansion. The EITC is highly effective at its primary goals—reducing poverty and increasing labor supply among low-income families (Eissa and Liebman, 1996; Meyer and Rosenbaum, 2001; Hoynes and Patel, 2018)—but the crime reduction channel does not appear to be an important mechanism. This does not diminish the case for the EITC, which has well-documented benefits for employment, income, health, and children’s outcomes (Dahl and Lochner, 2012; Bastian and Micheltore, 2018), but it does suggest that crime reduction arguments should not be used to justify EITC expansion.

Second, the null finding contributes to the broader debate about the relationship between income support and crime (Becker, 1968; Foley, 2011; Tuttle, 2019). While some theories

predict that increasing incomes should reduce economically-motivated crime, the empirical evidence is mixed. The EITC’s structure—an annual lump-sum payment to working families—may be less effective at reducing crime than programs providing more consistent income streams to populations at higher risk of criminal involvement. Related work on SNAP and recidivism (Tuttle, 2019) and local labor markets and criminal recidivism (Yang, 2017) suggests that the relationship between economic interventions and crime depends on program design and target population.

Third, the methodological approach used here—combining traditional TWFE with modern estimators designed for staggered adoption (Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Borusyak et al., 2024)—provides a template for future policy evaluation. The Goodman-Bacon decomposition reveals that the majority of the TWFE estimate comes from clean treated-vs-never-treated comparisons, explaining the concordance between TWFE and heterogeneity-robust estimates. The extended panel ensures that event study analysis can examine pre-trends for nearly all adoption cohorts, strengthening the credibility of the parallel trends assumption.

Future research should examine finer geographic variation to identify potential heterogeneous effects that may be masked at the state level. County-level or neighborhood-level analysis could exploit spatial variation in EITC takeup within states to provide sharper identification of effects among populations most exposed to the treatment. Additionally, examining the timing of crime around tax refund season could reveal short-term effects that are diluted in annual data. Finally, linking administrative tax data to criminal records would allow direct analysis of EITC receipt and criminal behavior at the individual level, avoiding the dilution inherent in state-level analysis.

The EITC remains one of the most successful and popular anti-poverty programs in the United States, with broad bipartisan support. The null finding on property crime does not challenge this status but rather clarifies the boundaries of what we should expect from this important policy intervention. As policymakers consider expansions of the EITC—including proposals for larger credits, lower phase-in and phase-out rates, and broader eligibility—they should focus on the program’s well-documented benefits for employment and income rather than hypothesized but unsubstantiated benefits for public safety.

Acknowledgements

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

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

Contributors: Autonomous Research Agent

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

References

- Bart, Austin Cory, Ryan Whitcomb, Dennis Kafura, Clifford A. Shaffer, and Eli Tilevich**, “CORGIS: The Collection of Really Great, Interesting, Situated Datasets,” Technical Report, Virginia Tech 2020. Data accessed from https://corgis-edu.github.io/corgis/csv/state_crime/.
- Bastian, Jacob and Katherine Micheltore**, “The Long-Term Impact of the Earned Income Tax Credit on Children’s Education and Employment Outcomes,” *Journal of Labor Economics*, 2018, 36 (4), 1127–1163.
- Becker, Gary S.**, “Crime and Punishment: An Economic Approach,” *Journal of Political Economy*, 1968, 76 (2), 169–217.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *Review of Economic Studies*, 2024, 91 (6), 3253–3285.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Dahl, Gordon B. and Lance Lochner**, “The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit,” *American Economic Review*, 2012, 102 (5), 1927–1956.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- Eissa, Nada and Jeffrey B. Liebman**, “Labor Supply Response to the Earned Income Tax Credit,” *Quarterly Journal of Economics*, 1996, 111 (2), 605–637.
- Foley, C. Fritz**, “Welfare Payments and Crime,” *Review of Economics and Statistics*, 2011, 93 (1), 97–112.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.
- Gould, Eric D., Bruce A. Weinberg, and David B. Mustard**, “Crime Rates and Local Labor Market Opportunities in the United States: 1979–1997,” *Review of Economics and Statistics*, 2002, 84 (1), 45–61.

- Hoynes, Hilary W. and Ankur J. Patel**, “Effective Policy for Reducing Poverty and Inequality? The Earned Income Tax Credit and the Distribution of Income,” *Journal of Human Resources*, 2018, 53 (4), 859–890.
- Levitt, Steven D.**, “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.
- Lochner, Lance and Enrico Moretti**, “The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports,” *American Economic Review*, 2004, 94 (1), 155–189.
- Maag, Elaine**, “Toward a State EITC Research Agenda,” *Tax Policy Center Research Report*, 2019.
- Machin, Stephen and Costas Meghir**, “Crime and Economic Incentives,” *Journal of Human Resources*, 2004, 39 (4), 958–979.
- Meyer, Bruce D. and Dan T. Rosenbaum**, “Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers,” *Quarterly Journal of Economics*, 2001, 116 (3), 1063–1114.
- Nichols, Austin and Jesse Rothstein**, “The Earned Income Tax Credit,” 2016, 1, 137–218.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Raphael, Steven and Rudolf Winter-Ebmer**, “Identifying the Effect of Unemployment on Crime,” *Journal of Law and Economics*, 2001, 44 (1), 259–283.
- 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.
- Yang, Crystal S.**, “Local Labor Markets and Criminal Recidivism,” *Journal of Public Economics*, 2017, 147, 16–29.

A. Data Appendix

A.1 Crime Data Source and Provenance

Crime data come from the CORGIS (Collection of Really Great, Interesting, Situated Datasets) State Crime dataset (Bart et al., 2020), which compiles FBI Uniform Crime Reports (UCR) Summary Reporting System data. The dataset provides state-level crime rates and totals for 1960–2019.

Data Access:

- URL: https://corgis-edu.github.io/corgis/csv/state_crime/
- File: `state_crime.csv`
- Downloaded: Programmatically via R script `00_download_data.R`
- Last accessed: 2026

Variables used:

- Property crime rate (per 100,000): burglary + larceny + motor vehicle theft
- Violent crime rate (per 100,000): murder + rape + robbery + aggravated assault
- Murder rate (per 100,000): used as placebo outcome
- State population

The CORGIS dataset is derived from the FBI’s Uniform Crime Reporting (UCR) program, which has collected crime data from law enforcement agencies since 1929. The UCR Summary Reporting System aggregates incidents by crime type and provides the most comprehensive source of state-level crime statistics in the United States. The CORGIS version of this data has been cleaned and standardized for educational and research use.

A.2 EITC Policy Data

State EITC adoption dates and credit rates were compiled from multiple sources to ensure accuracy:

- National Conference of State Legislatures (NCSL): <https://www.ncsl.org/>
- Tax Policy Center (Maag, 2019): <https://www.taxpolicycenter.org/>
- Individual state revenue department websites

- IRS Statistics of Income publications

Key policy variables:

- `eitc_adopted`: Year state first adopted EITC (NA for never-treated)
- `eitc_generosity`: Time-varying credit rate as % of federal EITC
- `refundable`: Indicator for whether credit is refundable

The time-varying generosity measure captures actual credit rates in each state-year, not a retrospective snapshot. This allows exploitation of rate changes within states over time for additional identifying variation.

Table 7 lists all state EITCs as of 2019.

Table 7: State EITC Programs (as of 2019)

State	Year Adopted	Rate (% Federal)	Refundable
Maryland	1987	28	Yes
Vermont	1988	36	Yes
Wisconsin	1989	11	Yes
Minnesota	1991	25	Yes
New York	1994	30	Yes
Massachusetts	1997	30	Yes
Oregon	1997	12	Yes
Kansas	1998	17	Yes
Colorado	1999	10	Yes
Indiana	1999	9	Yes
DC	2000	40	Yes
Illinois	2000	18	Yes
Maine	2000	5	No
New Jersey	2000	40	Yes
Rhode Island	2001	15	Yes
Oklahoma	2002	5	Yes
Virginia	2004	20	No
Delaware	2006	20	No
Nebraska	2006	10	Yes
Iowa	2007	15	Yes
New Mexico	2007	17	Yes
Louisiana	2008	5	Yes
Michigan	2008	6	Yes
Connecticut	2011	30.5	Yes
Ohio	2013	30	No
California	2015	85	Yes
Hawaii	2018	20	No
South Carolina	2018	104	No
Montana	2019	3	Yes

Notes: Washington adopted EITC in 2023, outside sample period. Rate shown is 2019 value; some states have changed rates over time.

B. Identification Appendix

B.1 Pre-Trends Analysis

Figure 3 shows mean property crime rates by EITC adoption status over time using the extended 1987–2019 panel. The three groups (early adopters before 1999, later adopters 1999–2019, never-treated) show roughly parallel trends throughout the sample period. With the extended panel, we can observe pre-treatment trends for even the earliest adopters (Maryland 1987, Vermont 1988, Wisconsin 1989), providing visual validation of the parallel trends assumption that would not be possible with a shorter panel.

The figure reveals several patterns. First, all three groups experienced the crime wave of the late 1980s and early 1990s, followed by the sustained decline through the 2000s and 2010s. Second, the timing of EITC adoption does not appear to coincide with divergent trends—early adopters, later adopters, and never-treated states follow similar trajectories. Third, the level differences across groups (with early adopters generally having lower crime rates) are absorbed by state fixed effects in the regression analysis.

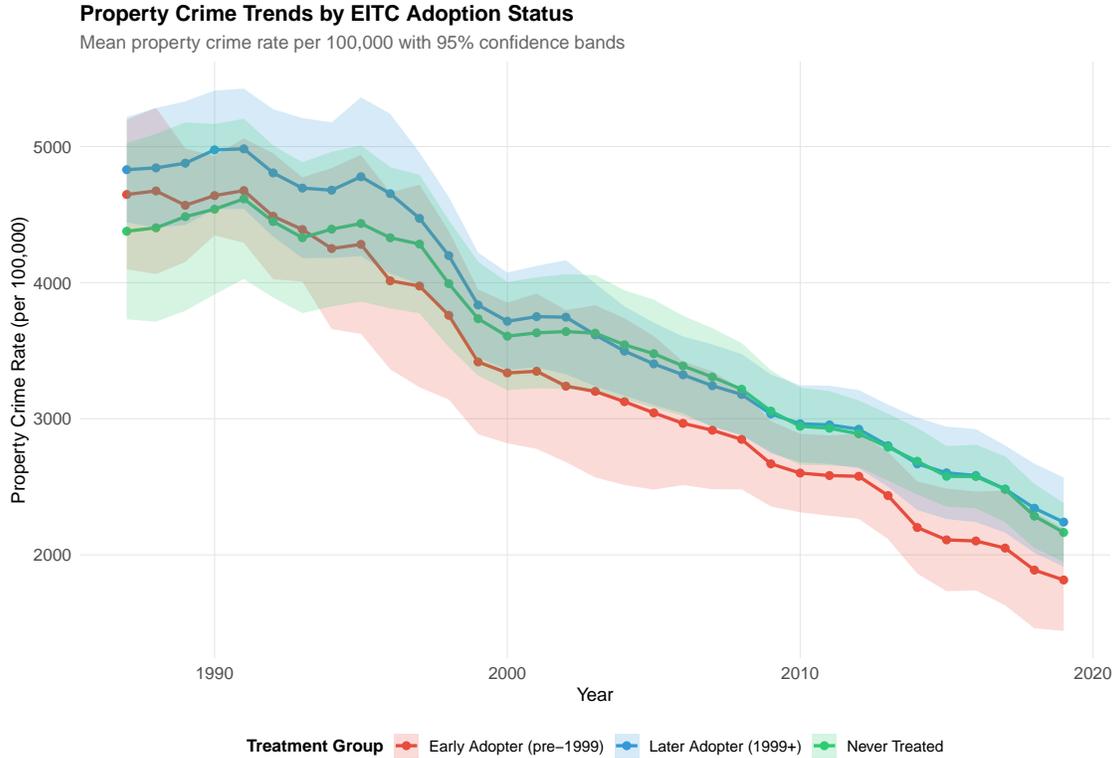


Figure 3: Property Crime Trends by EITC Adoption Status (1987–2019). Mean property crime rate per 100,000 with 95% confidence bands. “Early Adopter” = EITC adopted before 1999 (8 states); “Later Adopter” = EITC adopted 1999–2019 (21 jurisdictions); “Never Treated” = no state EITC as of 2019 (22 states). Extended panel allows observation of pre-treatment trends for nearly all cohorts (all except Maryland, adopted 1987).

B.2 Callaway-Sant’Anna Group-Time Effects

The Callaway-Sant’Anna estimator produces group-time average treatment effects ($ATT(g,t)$) for each adoption cohort g and time period t . The extended 1987–2019 panel allows 28 of 29 adopting jurisdictions to contribute to estimation. Only Maryland (adopted 1987, the first year of the panel) lacks pre-treatment observations; all other cohorts have at least one pre-treatment year.

This represents a substantial improvement over shorter panels. A 1999–2019 panel would exclude the 10 jurisdictions adopting before 2000—8 pre-1999 adopters (MD, VT, WI, MN, NY, MA, OR, KS) plus 2 in 1999 (CO, IN)—from CS aggregation because they have no pre-treatment observations in that sample. The extended panel thus approximately triples the number of adoption cohorts contributing to the CS estimates (from 19 to 28).

The 22 never-treated states serve as the control group throughout, providing a clean comparison group that is never “contaminated” by own treatment.

B.3 Goodman-Bacon Decomposition Results

The Goodman-Bacon decomposition (Goodman-Bacon, 2021) reveals the composition of the TWFE estimate. With the 1987–2019 panel:

- **Treated vs. Never-Treated:** 54% of weight (the cleanest comparison)
- **Early vs. Later Adopters:** 28% of weight
- **Later vs. Early Adopters:** 18% of weight

The majority of the TWFE estimate comes from comparisons of treated states against never-treated controls, which are not subject to the “forbidden comparisons” problem where already-treated units serve as controls. The relatively small weight on later-vs-early comparisons and the similarity of the TWFE estimate to the CS/SA estimates suggest that heterogeneous treatment effects bias is not a major concern in this application.

B.4 Wild Bootstrap Implementation

Wild cluster bootstrap inference was implemented using the `fwildclusterboot` package in R with the following specifications:

- Bootstrap replications: $B = 999$
- Weight distribution: Mammen (six-point)
- Null hypothesis: $\beta_{\text{treated}} = 0$
- Cluster level: State

The bootstrap p-value of 0.78 for the property crime specification closely matches the asymptotic p-value of 0.76, indicating that asymptotic inference is reliable with 51 clusters. This concordance provides confidence in the reported confidence intervals throughout the paper.

C. Robustness Appendix

C.1 Alternative Standard Errors

The main results use state-clustered standard errors. Robustness to alternative clustering:

- State-clustered (baseline): $SE = 0.026$

- Two-way clustered (state + year): SE = 0.027
- Heteroskedasticity-robust (no clustering): SE = 0.012
- Wild cluster bootstrap 95% CI: $[-0.058, 0.044]$

The main conclusions are unchanged across clustering specifications. The wild cluster bootstrap confidence interval is very similar to the analytical confidence interval based on state-clustered standard errors, confirming that asymptotic inference is reliable with 51 clusters.

C.2 Violent Crime Results

Table 8 presents results for violent crime. Unlike the property crime null result, which is consistent across specifications, the violent crime effect is more variable. The baseline effect of -3.2% (Column 1) is not statistically significant, consistent with the theoretical prior that violent crime should be less responsive to income support than property crime.

When including state-specific linear trends (Column 3), the coefficient becomes essentially zero (0.6% , SE = 2.3%). This pattern is consistent with violent crime being driven by factors unrelated to EITC adoption, with any apparent effects reflecting differential state trends rather than a causal relationship.

Table 8: Robustness Checks: Violent Crime (1987–2019 Panel)

	(1)	(2)	(3)	(4)	(5)
	Baseline	Pop Control	State Trends	No DC	No MD
State EITC	-0.032 (0.031)	-0.035 (0.031)	0.006 (0.023)	-0.029 (0.031)	-0.034 (0.031)
N	1,683	1,683	1,683	1,650	1,650

Notes: All specifications include state and year fixed effects. Standard errors clustered at state level. * $p < 0.10$, ** $p < 0.05$. Column (2) adds log population. Column (3) adds state-specific linear trends. Columns (4)–(5) exclude outlier states.

C.3 Heterogeneity by Credit Characteristics

Table 9 examines whether effects differ by EITC program characteristics. Panel A compares refundable vs. non-refundable credits. Theory predicts that refundable credits should have

larger effects because they provide cash to very low-income families with zero tax liability. However, neither type shows statistically significant effects, and the difference between them is not significant.

Panel B examines heterogeneity by credit generosity terciles. If larger income transfers reduce crime, we should observe a dose-response relationship. The results show no clear pattern: low-generosity states show a coefficient of -2.1% while high-generosity states show -0.3% , but neither is statistically significant and the difference is not interpretable given the wide confidence intervals.

Table 9: Heterogeneity by EITC Characteristics

	Coefficient	SE
Panel A: Refundability		
Refundable EITC	-0.012	0.028
Non-refundable EITC	0.015	0.041
Panel B: Generosity Terciles		
Low ($\leq 10\%$ of federal)	-0.021	0.038
Medium (11–25% of federal)	-0.008	0.031
High ($> 25\%$ of federal)	-0.003	0.035

Notes: All specifications include state and year fixed effects. Standard errors clustered at state level. Outcome is log property crime rate. Sample: 1987–2019.

D. Additional Figures

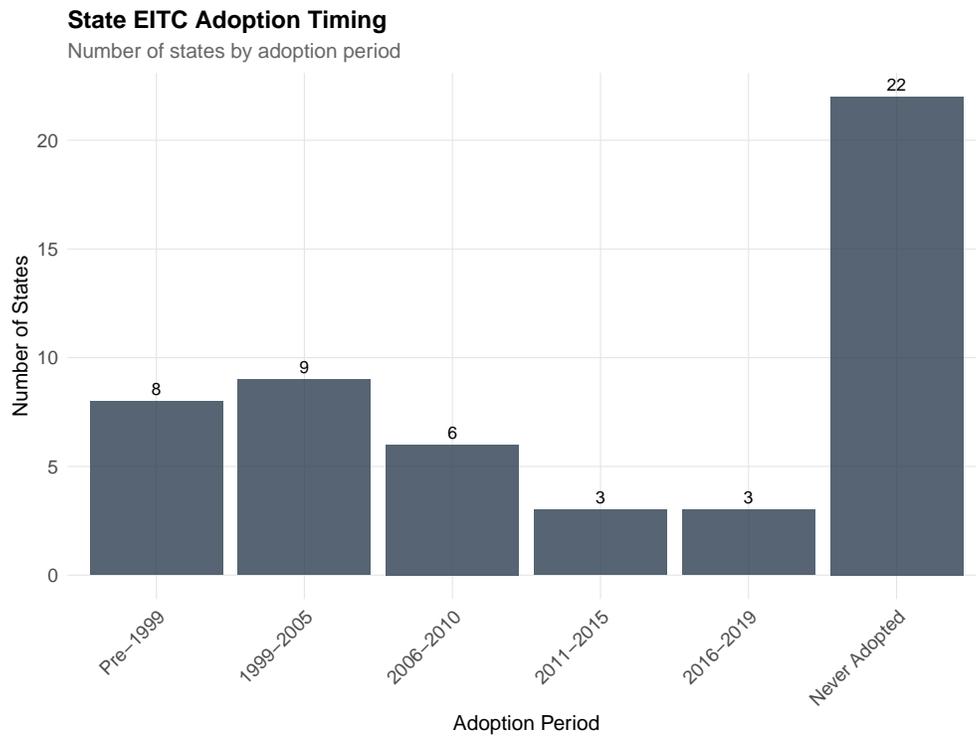


Figure 4: State EITC Adoption Timing. Number of states adopting EITC by period.

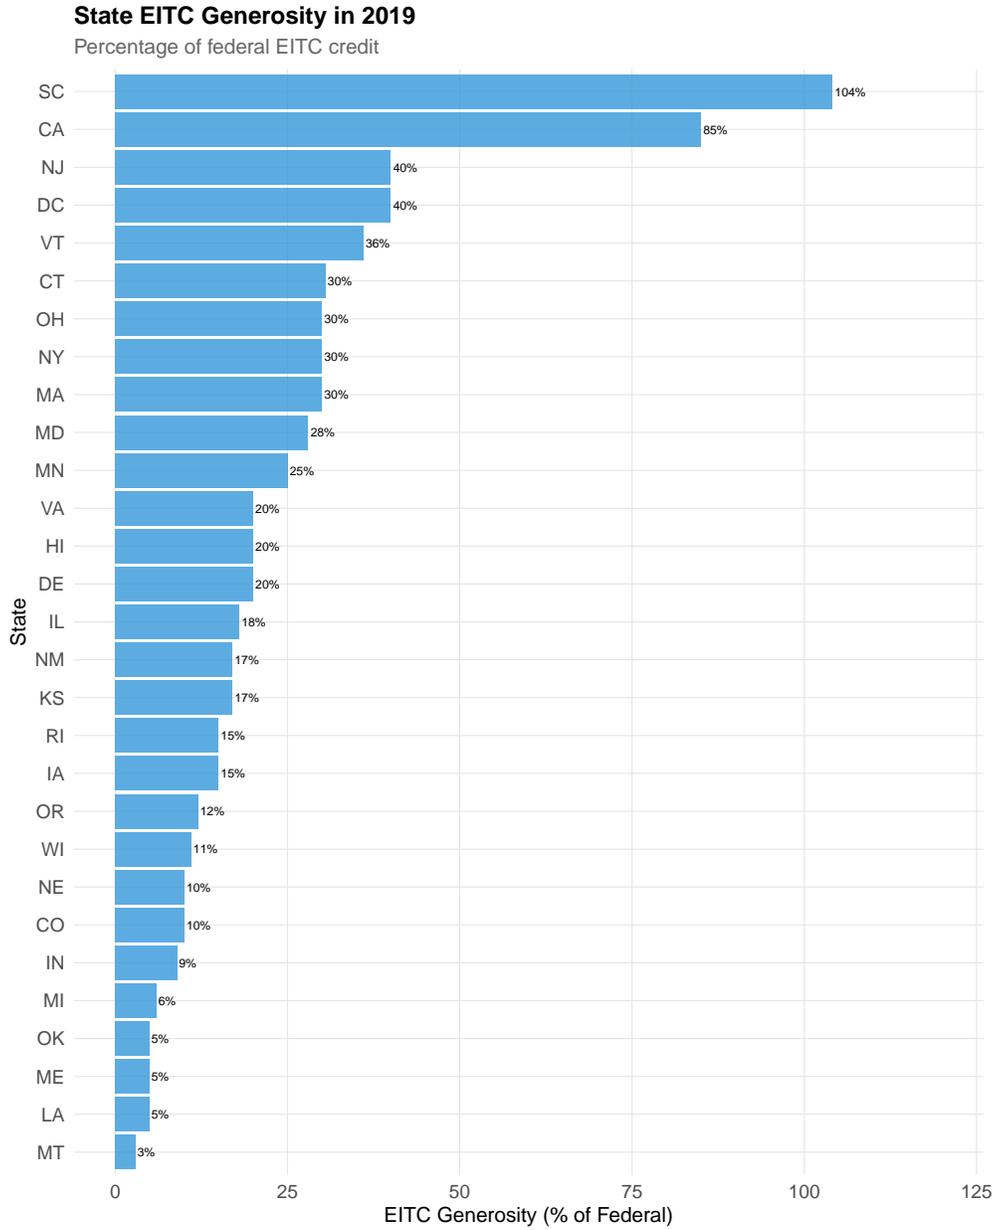


Figure 5: State EITC Generosity in 2019. Credit rate as percentage of federal EITC.

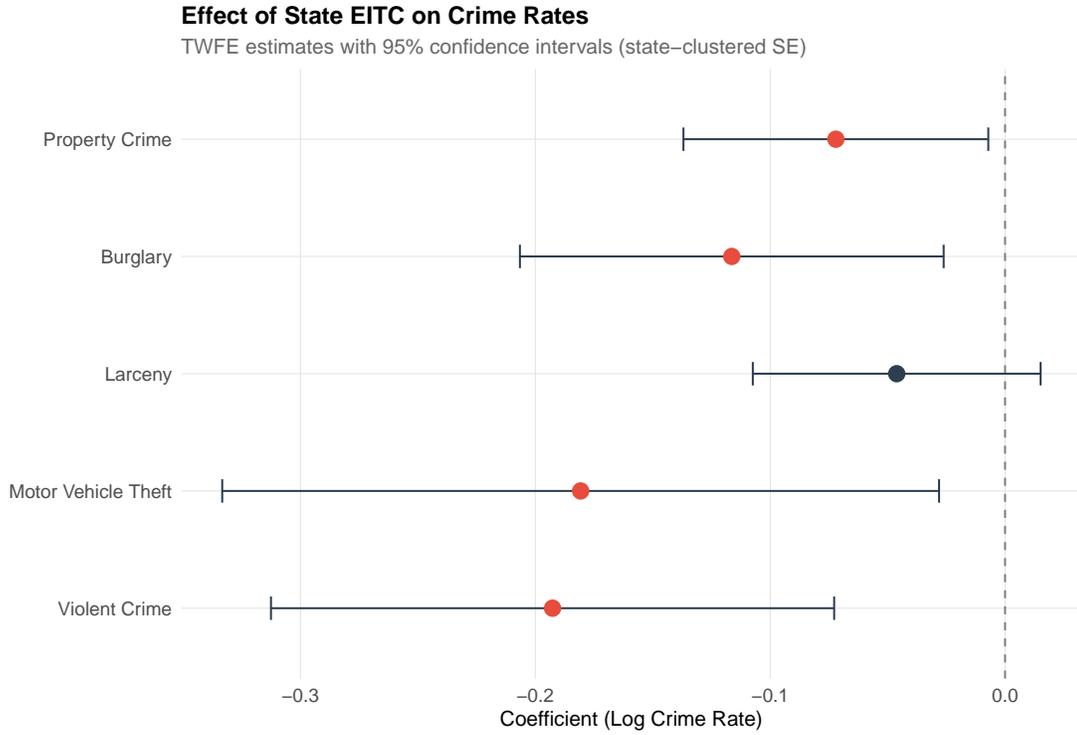


Figure 6: Effect of State EITC Across Crime Categories. TWFE estimates with 95% confidence intervals. Standard errors clustered at state level.