# Source of Income Discrimination and Homelessness: Effects of Anti-Discrimination Laws

Luke Maddock
*Department of Economics, Colorado State University, Fort Collins, CO 80523*
*Luke.Maddock@colostate.edu*

July 13, 2025

**Abstract**

This paper analyzes how Source-of-Income Protection (SOIP) policies affect low-income renters who use Housing Choice Vouchers (HCVs) to lease housing. In many localities, landlords can legally reject tenants solely due to the use of a public subsidy. Recent legislation aims to curtail such practices by labeling voucher holders' income as "lawful source," thereby restricting landlords' ability to discriminate. To study the consequences of these reforms, I assemble a panel of Continuums of Care (CoCs) spanning 2009–2018 and track staggered policy adoption across states, counties, and municipalities. Relying on a modern difference-in-differences framework that accounts for heterogeneous treatment timing, I estimate how SOIP adoption influences outcomes for households participating in the HCV program and households in the broader locality. Results from an intent-to-treat framework indicate that housing and neighborhood characteristics improve for HCV households after an SOIP policy is passed, but rental housing costs and homelessness increase for the entire area. One interpretation is that the policy improves outcomes for some voucher tenants but triggers offsetting housing provider decisions that adversely affect other low-income renters. These findings highlight how well-intended rules against source-of-income discrimination do not necessarily curb overall housing instability.

# 1 Introduction

Across the United States, approximately 2.3 million low-income households rely on Housing Choice Vouchers (HCVs, often referred to as "Section 8" vouchers) to bridge the gap between their incomes and market rents.[1] Despite broad policy goals of voucher mobility and poverty deconcentration, the success of voucher holders in leasing quality housing often falls short of expectations. One salient barrier is the prevalence of landlord discrimination by "source of income" (SOI), which can leave voucher recipients unable to rent in higher-opportunity neighborhoods (Tighe et al., 2016).

In many jurisdictions, landlords can legally refuse prospective tenants whose primary income derives from HCVs or other public benefits, even though voucher holders are disproportionately members of historically protected groups. While some states and municipalities have adopted explicit SOI antidiscrimination provisions—for instance, forbidding landlords from advertising "No Section 8"—other jurisdictions have either excluded vouchers from legal protection or even preempted local SOI ordinances altogether. This policy patchwork has generated considerable variation in how protected voucher holders are, with potentially significant consequences for both voucher utilization and housing affordability (Tighe et al., 2016).

At the same time, a growing body of research highlights the ways in which eviction and homelessness can follow from disruptions in the rental market. Recent work documents the negative consequences of forced moves for household stability, health, and material well-being (Desmond, 2016), and the fiscal cost that homelessness imposes on local governments (Abramson, 2022). Although policies such as "Right-to-Counsel" or eviction moratoria have been implemented in some localities (Abramson, 2022), SOI protections can complement these efforts by expanding the set of housing options available to voucher recipients, potentially reducing the likelihood of displacement in the first place.

This paper investigates how the adoption of source-of-income protection (SOIP) statutes affects three key outcomes: (i) voucher utilization rates, i.e. the fraction of distributed vouchers that are eventually used to lease housing, (ii) housing affordability measures such as rents, and (iii) local homelessness rates (Abramson, 2022). The analysis bridges two strands of the literature. First, it extends the line of inquiry presented in Tighe et al. (2016), which emphasizes the uneven state and local legislation surrounding SOI protections and the challenges voucher holders face. Second, it connects to equilibrium-oriented studies (Abramson, 2022) that examine how landlord and tenant behaviors may shift under policies that prohibit discrimination or constrain eviction processes, with potential ripple effects for homelessness.

From a methodological perspective, the paper employs a modern difference-in-differences framework

[1] See https://www.cbpp.org/research/housing for national HCV statistics from the Center on Budget and Policy Priorities.

(Callaway and Sant'Anna, 2021) designed for staggered policy adoption. This setup compares CoCs that implement SOIP with those that remain untreated, relying on variation in local ordinance adoption across municipalities, counties, and states during the period of interest. By tracking outcomes before and after SOIP enactment in places with similar underlying trends, I estimate how voucher usage, rent levels, and homelessness respond to changes in the legal status of SOI-based refusals.

In exploring partial or incomplete adoption at the local level, the results highlight how gaps in coverage can weaken the intended benefits of SOI laws. An SOIP, in principle, restricts landlords' capacity to refuse HCV tenants outright, but the scope of such rules and the intensity of their enforcement may alter how effectively voucher holders can find units. The study also examines whether landlords respond to these mandates by adjusting rents or screening practices, potentially offsetting the gains in housing stability. The findings contribute to ongoing policy discussions on how best to combine income-based subsidies and legal protections to combat housing instability.

The remainder of the paper proceeds as follows. Section 2 presents institutional details on the voucher program and SOI legislation. Section 3 outlines data sources and the construction of key variables. Section 4 explains the identification strategy, including a staggered DiD approach inspired by Callaway and Sant'Anna (2021). Section 5 reports preliminary findings on voucher utilization, rent levels, and homelessness. Section 6 discusses next steps, potential robustness tests, and theoretical mechanisms that might explain why homelessness does not show a clear downward shift despite stronger voucher acceptance policies.

# 2 Institutional Background

## 2.1 Housing Choice Voucher Program

The Housing Choice Voucher (HCV) program, often called Section 8, is the largest tenant-based rental assistance initiative in the United States. Under this arrangement, low-income households contribute roughly 30 percent of their monthly adjusted income toward rent, and the local Public Housing Authority (PHA) pays the remainder up to a specific payment standard.[2] The program is intended to increase mobility by allowing recipients to lease housing in a range of neighborhoods, rather than remaining concentrated in subsidized complexes within high-poverty areas (Tighe et al., 2016).

Nonetheless, voucher usage has been hindered by several market frictions. In many localities, landlords may lawfully reject applications based on a prospective tenant's reliance on a voucher, making it difficult for participants to lease in higher-opportunity areas (Tighe et al., 2016). Some landlords also perceive

[2] For a thorough overview of the HCV program's history and its policy motivations, see Desmond (2016).

the administrative requirements of dealing with PHAs, such as inspections and paperwork, to be more burdensome than under a typical market lease, which further discourages them from accepting vouchers. Time constraints on households also matter, because recipients must locate qualifying units within a defined search window (often 60 to 120 days), after which they can lose their subsidy (**?**). In addition, many recipients ultimately reside in lower-rent neighborhoods, undermining the intended poverty-deconcentration goal (Tighe et al. 2016). Observers who favor stronger legal protections against landlord discrimination believe reducing these obstacles could raise overall voucher utilization, which may, in turn, broaden neighborhood choices for low-income households.

## 2.2 Source of Income Legislation

### 2.2.1 Variations in Coverage Across Jurisdictions

Source of income (SOI) legislation prevents landlords from refusing to rent based on a tenant's lawful income streams, such as wages, child support, or public assistance. These ordinances often amend fair housing codes to protect "lawful source of income" in the same way that categories like race or religion are protected. When a locality explicitly includes Housing Choice Vouchers within this coverage, landlords who reject voucher holders solely on that basis can be found liable for discrimination.

The extent of SOI protection varies widely. Some states (e.g. New Jersey, Maine) and cities (e.g. Washington, DC, New York) include vouchers as a fully protected income source, effectively prohibiting phrases like "No Section 8" in housing advertisements (Freeman, L., & Li, Y. , 2013; Hangen , 2022). Others provide partial coverage that excludes vouchers, allowing landlords to continue screening out applicants for relying on subsidies. In certain states such as Indiana or Texas, local attempts to enact broader SOI protections are preempted by state law, creating large "voucher deserts" where the policy offers no recourse if applicants are denied because they hold HCVs (Tighe et al., 2016).

### 2.2.2 Potential Outcomes and Challenges

Supporters of SOI legislation anticipate an increase in voucher utilization and a decline in the spatial concentration of poverty (Freeman, L., & Li, Y. , 2013). If landlords can no longer refuse voucher applicants outright, more households should be able to secure leases in a wider range of neighborhoods (Tighe et al., 2016). However, property owners who regard vouchers as riskier or more burdensome may raise rents or introduce extra fees in order to compensate, which could undermine affordability gains. Abramson (2022) highlights how legal interventions protecting tenants can lead to higher costs for landlords and, in turn, potential rent inflation.

Analyzing the net effect of SOI rules thus requires examining both direct impacts on voucher-holding households and broader shifts in the housing market. Although some might expect that restricting income-based discrimination would reduce evictions and homelessness, partial coverage, limited enforcement, or other supply constraints could dampen or even negate these expected benefits. The empirical strategy in Section 4 explains how this paper identifies and estimates the effects of staggered SOI policy adoption on voucher lease-up, local rents, and homelessness.

## 3 Data and Outcome Measures

A central component of this study is an original data set that tracks source-of-income (SOI) protection policies across the United States and aligns them with Continuums of Care (CoCs). The first element of this construction relies on the National Low Income Housing Coalition's *Tenant Protection Database*, which identifies each municipality, county, or state that has implemented an SOI ordinance. For each policy entry, the database records the date of passage, the geographic boundary of coverage, and relevant legal details specifying whether the legislation fully or partially covers recipients of Housing Choice Vouchers. Because SOI provisions apply at varying spatial scales—from small municipalities to entire states—it is necessary to determine how these boundaries intersect with CoCs, which are the units of analysis for many federal homelessness and voucher metrics.

To integrate these policy areas with CoCs, I employ shapefiles obtained from the U.S. Department of Housing and Urban Development (HUD). The shapefiles delineate the exact boundaries for every CoC. By crosswalking these CoC boundaries with each policy-passing municipality or county, it becomes possible to assess the extent to which any given CoC lies within or overlaps a jurisdiction that has enacted an SOI law. In cases where an SOI policy originates at the county or state level and the entire CoC is subsumed within that jurisdiction, the CoC is considered fully affected by the legislation. If, however, only a single municipality within a larger CoC passes its own SOI ordinance, the CoC's coverage is partial, confined to the municipal boundary.

Quantifying these partial effects requires estimates of local population and voucher-eligible households within each jurisdiction. I rely on the `tidycensus` package in R to retrieve 5-year American Community Survey (ACS) data at the tract level. After spatially joining tracts to CoCs, the data yield totals for both the full population and the population eligible for vouchers. Because multiple localities can pass SOI legislation at different times, I individually identify and aggregate the population affected by the policy within each CoC boundary. For example, if one municipality in a CoC adopts an SOI law, I sum the voucher-eligible population in that municipality. These figures are then compared to the total voucher-eligible population in the CoC to generate a continuous measure of treatment intensity:

$$\text{Treatment Intensity}_{k,t} = \frac{\text{Voucher-Eligible Pop Affected in CoC } k \text{ at time } t}{\text{Total Voucher-Eligible Pop in CoC } k \text{ at time } t}.$$

When the entire CoC population is covered, this ratio equals one, and when no policy is present, it equals zero.

In addition to measuring policy exposure, the data set includes several outcome variables. The *voucher utilization rate* represents the percentage of housing vouchers distributed by PHAs that eventually result in a household signing a lease. I identify the relevant PHAs within each CoC, track how many vouchers they distribute, and compare that number to how many are ultimately used by recipients to rent a unit. The second main outcome, *median CoC rent*, comes from tract-level 5-year ACS estimates, aggregated to the CoC level. Finally, the *homelessness rate* per 10,000 residents derives from HUD's annual Point-in-Time (PIT) counts, which record the number of sheltered and unsheltered individuals in each CoC on a single night in January. Dividing the PIT count by the CoC's total population from the ACS and scaling to a per-10,000 figure produces a standardized measure of homelessness severity.

Constructing the merged data set therefore requires multiple steps: downloading and harmonizing the Tenant Protection Database; mapping each policy to CoC boundaries; merging in ACS-based population and rent indicators through `tidycensus`; and appending PIT counts for the homeless population. The final result is a panel of CoCs with time-varying indicators of treatment intensity (the fraction of the CoC covered by SOI rules), as well as key outcomes including voucher lease-up rates, median rental prices, and homelessness rates. This combined resource underpins the empirical analysis that follows.

# 4 Identification Strategy and Estimation Procedures

A key feature of this research design is the heterogeneous timing of Source of Income Protection (SOIP) policy adoption across Continuums of Care (CoCs). Between 2009 and 2018, different municipalities, counties, or states enacted (or declined to enact) such legislation at various points in time. Because each CoC is observed annually over that horizon, the data take a staggered rollout form, in which some CoCs begin to be treated earlier than others, while a subset never receives treatment at all.

Following Callaway and Sant'Anna (2021), I implement a difference-in-differences framework adapted to multiple time periods, referred to here as "staggered DiD." The estimation proceeds by comparing outcomes of CoCs that become treated in a particular year to outcomes of CoCs that remain untreated in that same year, while accounting for the possibility that other CoCs may have already adopted the policy. Specifically, for each group $g$ and time $t$, I estimate an average treatment effect on the treated, denoted $\text{ATT}(g,t)$, which captures the impact of the policy for CoCs first treated in year $g$ at time $t$.

To ensure a well-defined counterfactual, I drop CoCs that entered 2009 with a pre-existing SOIP policy. These CoCs do not offer a valid pretreatment baseline in the context of interest, and including them could confound the staggered adoption comparisons. Consequently, in the final sample, the "never treated" CoCs function as the main comparison group. I further restrict the time window to 2009–2018. Homelessness counts from the HUD Point-in-Time (PIT) survey underwent structural changes around 2008, and the disruptions caused by COVID-19 in 2020 affect PIT measures thereafter, making consistent data beyond 2018 less reliable.

An important requirement for DiD identification is that, in the absence of treatment, the outcome trajectories of treated and untreated CoCs would have followed parallel paths. One potential threat is that policy adoption might occur more frequently in CoCs already experiencing rising homelessness or poor voucher outcomes. To investigate this, I conduct a simple balance exercise comparing never-treated CoCs to eventually treated CoCs at baseline. The sample means of the primary outcome variables—voucher utilization, rent affordability measures, and homelessness—are statistically indistinguishable across these two groups, and similar patterns emerge for demographic and economic covariates. These checks do not prove exogeneity but mitigate concerns that selection on unobservables is severe enough to invalidate the staggered DiD approach.

## 4.1 Econometric Specification and Estimation Setup

To estimate how Source of Income Protection (SOIP) policies affect voucher utilization, median rents, and homelessness, I employ the approach from Callaway and Sant'Anna (2021), which generalizes difference-in-differences (DiD) to handle multiple time periods and staggered treatment adoption. The unit of observation is a Continuum of Care (CoC) $i$, observed in year $t$, over a panel spanning 2009 to 2018. Let $G_i$ denote the first year CoC $i$ becomes treated. Denote the outcome of interest by $Y_{i,t}$, whether it is voucher utilization, median rent, or homelessness.

In a standard fixed-effects DiD, a common estimation strategy employs an ordinary least squares (OLS) equation of the form:

$$Y_{i,t} = \alpha_i + \lambda_t + \beta \mathbf{1}\{G_i \leq t\} + \varepsilon_{i,t},$$

where $\alpha_i$ and $\lambda_t$ represent CoC and year fixed effects, and $\mathbf{1}\{G_i \leq t\}$ is an indicator that CoC $i$ is treated by year $t$. However, because policies are adopted at different times and may yield heterogeneous impacts, a single $\beta$ is insufficient to capture the complexity of staggered adoption.

Instead, I use the `att_gt` function from Callaway and Sant'Anna (2021) using the `did` package in R, which obtains group-time average treatment effects on the treated (ATTs) in a doubly robust manner and accounts for different treatment cohorts.

For use, `yname` specifies the outcome variable, `tname` references the calendar year, `idname` is the unique CoC identifier, `gname` indicates the first treatment year $G_i$, and `control_group` defines CoCs that never adopt the policy as the primary comparison. The argument `est_method = "dr"` activates a doubly robust estimation procedure that uses both an outcome regression model and a propensity score model to reduce bias due to confounders or misspecification. Standard errors are clustered by state to account for correlated shocks within each state.

Repeating this procedure for each outcome of interest (voucher utilization rate, log of median CoC rent, and the rate of homelessness per 10,000 CoC population) produces a set of group-time ATT estimates, $\widehat{\text{ATT}}(g,t)$. To summarize results, I aggregate or transform these estimates using the `aggte` function (e.g., `type = "dynamic"`) to generate event-study coefficients centered on each group's adoption date. These event-study estimates show how the outcome evolves from several years before to several years after adoption, thereby revealing a time path of policy impacts. Formally, for group $g$ and year $t \geq g$,

$$\widehat{\text{ATT}}(g,t) \;=\; E\Big[Y_{i,t}(1) - Y_{i,t}(0) \;\Big|\; G_i = g\Big],$$

where $Y_{i,t}(1)$ and $Y_{i,t}(0)$ are the treated and untreated potential outcomes, respectively. The assumption of parallel trends is enforced via either outcome regression, inverse-probability weighting, or both, in a doubly robust manner. This framework also accommodates the CoC and time fixed effects implicitly by comparing within-year outcomes of treated CoCs to those of CoCs that remain untreated at that same time. Finally, by clustering standard errors at the state level, I ensure that state-specific shocks or policy environments do not spuriously inflate the precision of the estimated treatment effects.

For the current analysis, I treat the policy as strictly binary, despite the fact that partial coverage arises within some CoCs. Although Callaway, Bacon, Sant'Anna (2024) (forthcoming) and related work propose methods to handle continuous or multi-valued treatments, I have not yet implemented these approaches in the main regression results. Instead, any CoC with at least some fraction of its population covered by an SOIP law at time $t$ is labeled as "treated." Future extensions will refine this coding by incorporating the continuous treatment intensity measure introduced in Section 3, which would allow estimation of how partial vs. full coverage shifts outcomes.

.

Figure 1: Voucher Utilization Rate - Callaway & Sant'Anna (2021)

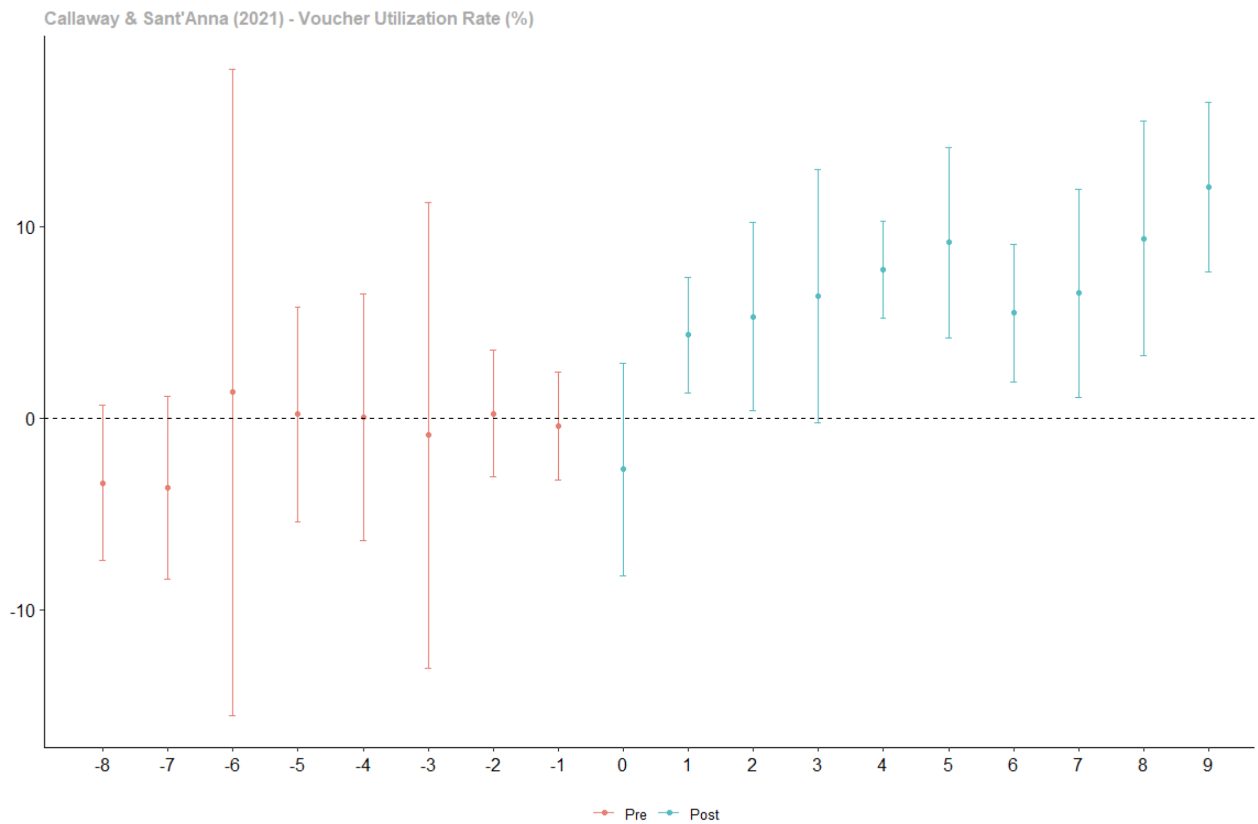


## 5 Preliminary Findings

Figure 1 first presents the event-study estimates for the voucher utilization rate, which is arguably the most direct outcome of Source of Income Protection (SOIP). The horizontal axis measures relative time, with zero denoting the first full year in which a Continuum of Care (CoC) is treated. Negative values represent pre-treatment leads, while positive values capture post-treatment lags. In the years prior to SOIP adoption, the point estimates remain close to zero and do not indicate a violation of parallel trends. After the policy takes effect, the estimates shift upward and stabilize in the range of a 5–10 percentage-point increase. Confidence intervals grow wider at later lags, but the positive shift remains robust across multiple post-treatment observations. This pattern aligns with the central policy rationale that prohibiting landlords from rejecting voucher holders *per se* permits more households to lease successfully.

Figure 2 shows the same relative-time event-study framework applied to the log of median rent. The pre-adoption estimates hover near zero, reinforcing that treated and untreated CoCs follow comparable trajectories before the policy. By about four years after adoption, the coefficients suggest rent increases in the range of 0.10–0.15 in logs, with wider confidence intervals at more distant lags. Although the magnitude

Figure 2: Log of Median Rent - Callaway & Sant'Anna (2021)

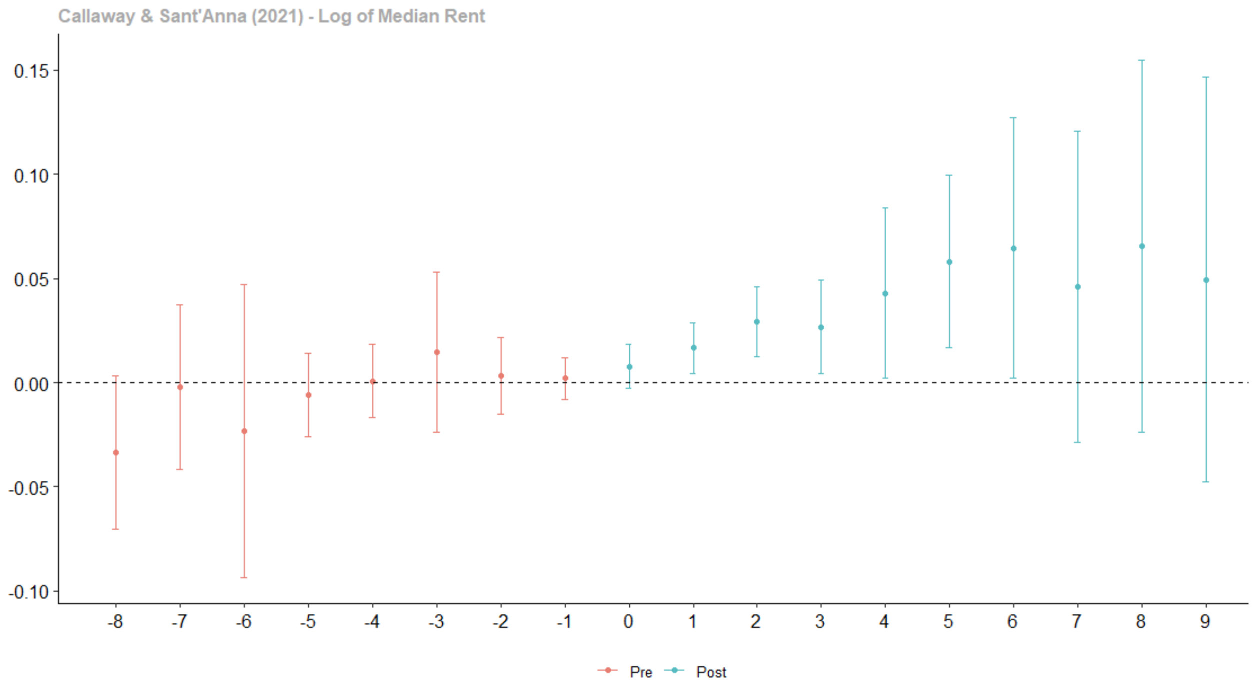


of these changes can vary, the upward trend lends credibility to the possibility that stronger SOIP coverage encourages landlords either to raise rents in anticipation of higher demand from voucher holders or to offset potential risks and administrative costs.

Figure 3 then reports results for homelessness per 10,000 individuals. As before, the horizontal axis indexes the relative event time around the adoption year. In the pre-treatment period, the estimates fluctuate around zero and show no discernible pattern, suggesting that the parallel trends assumption is not severely violated. After enactment, the estimates remain close to zero, but several lags display slightly positive effects. Summing the post-treatment effects indicates that, on average, the policy increases the homelessness rate by approximately 4.2 people per 10,000 CoC residents, statistically significant at the 10% level. While not definitively large or precise in each individual year, this aggregated figure stands contrary to a naive expectation that restricting SOI discrimination would naturally lower homelessness.

Overall, these event-study findings imply that CoCs with newly enacted SOIP experience a clear rise in voucher lease-up rates and moderate rent inflation, yet the homelessness measure shows little sign of declining. External factors—such as how thoroughly the policy is enforced, whether its coverage is only partial within a CoC, and the pre-existing housing supply constraints—could be masking or offsetting the direct benefits that higher voucher utilization might otherwise confer. Later sections expand on these questions of partial coverage, local heterogeneity, and landlord behavior, with an eye toward explaining why

Figure 3: Homelessness per 10,000 - Callaway & Sant'Anna (2021)

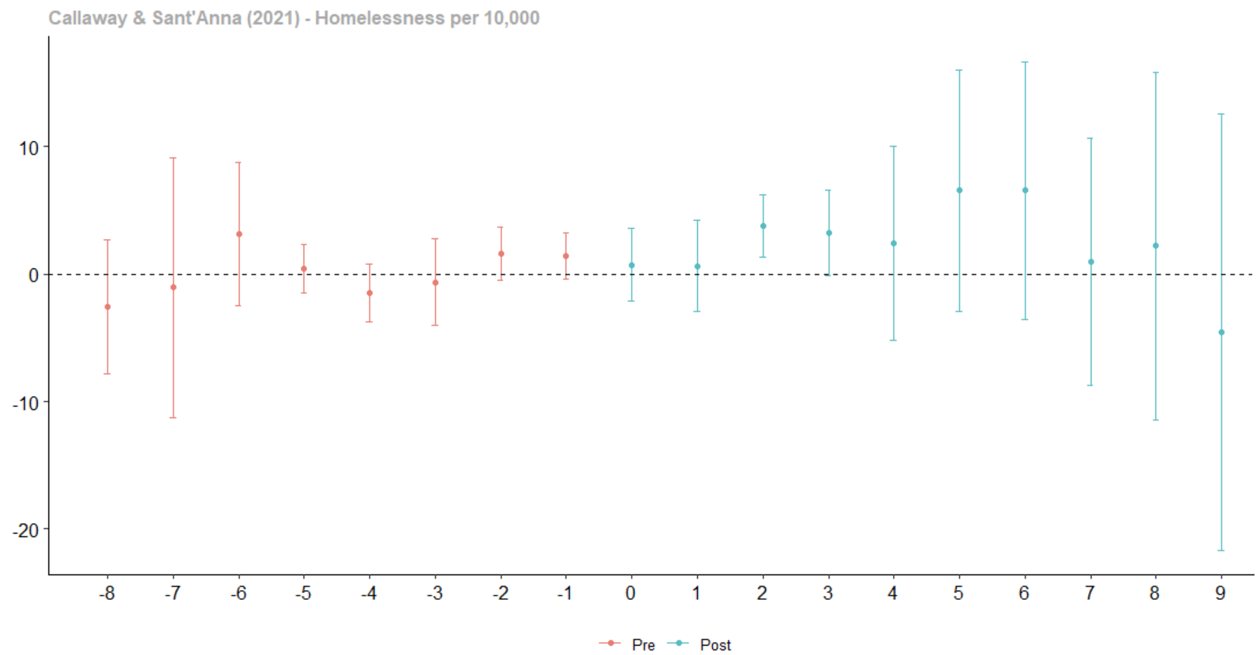


an ostensibly pro-tenant policy might not deliver its most intuitive outcome of reducing homelessness.

## 6 Next Steps and Further Analysis

Although these preliminary findings provide insight into how Source of Income Protection (SOIP) policies coincide with changes in homelessness, rent levels, and voucher utilization, important extensions remain for making the analysis both more precise and more policy-relevant. Two directions stand out: (1) extending the empirical framework to accommodate *continuous treatment intensity* in the sense of Callaway, Bacon, Sant'Anna (2024), and (2) conducting additional sensitivity analyses that clarify whether observed patterns could be artifacts of model specification, time horizons, or market conditions rather than true policy effects.

### 6.1 Callaway, Bacon, and Sant'Anna (2024) Continuous-Treatment Framework

A key limitation in the current empirical design is the binary coding of treatment. In many Continuums of Care (CoCs), legislation does not uniformly cover all inhabitants; a fraction of the CoC's population may be subject to SOI regulations while the remainder is not. Methods introduced by Callaway, Bacon, Sant'Anna (2024) tackle such scenarios by extending the difference-in-differences (DiD) logic to a setting where the "dose" of treatment (i.e., the fraction of CoC residents covered by a policy) varies between 0 and 1. Implementing this framework entails estimating the average effect of shifting from one dose level to another, rather than merely comparing the presence versus absence of treatment.

Under the approach described by Callaway, Bacon, Sant'Anna (2024), each unit (CoC-year observation) is associated with a treatment level $D_{k,t} \in [0,1]$. The first step involves nonparametrically estimating

$$E[\Delta Y_{k,t} \mid D_{k,t} = d],$$

where $\Delta Y_{k,t}$ is the change in outcomes (e.g., from a baseline period) for CoC $k$ at time $t$. Depending on the researcher's preference, this can be pursued via local polynomial regression, spline methods, or flexible series approximations that regress $\Delta Y_{k,t}$ on the continuous treatment $d$. Once the conditional mean function $m(d) = E[\Delta Y_{k,t} \mid D_{k,t} = d]$ is estimated, the next step is to construct parameters such as

$$\widehat{\mathrm{ATT}}(d \mid d) \;=\; m(d) - m(0),$$

which captures the average effect of having dose $d$ among units that actually experienced $d$. It is also possible to derive average causal responses (ACRs) by differentiating the function $m(d)$ with respect to $d$, or by examining discrete differences between adjacent dose levels. This approach thus yields a richer characterization of policy impacts across different degrees of SOI coverage, rather than forcing a binary distinction between "treated" and "untreated."

Adopting this continuous DiD estimator requires careful attention to data construction. Each CoC-year must be assigned a precise fraction of population covered by the legislation, which in turn depends on tract-level population data, boundary crosswalks, and accurate policy timing. The approach also relies on a generalized parallel trends assumption, meaning that units experiencing dose $d$ in period $t$ should exhibit outcome trends similar to those of hypothetical "comparison" units at each relevant dose level. Verifying such an assumption calls for diagnostic plots of $\Delta Y_{k,t}$ against $D_{k,t}$ in the pretreatment period and, where feasible, event-study-like comparisons across small intervals of the dose. If these checks lend credibility to a generalized parallel trends condition, then the Callaway, Bacon, Sant'Anna (2024) method can yield dose-response curves and average derivative estimates that deepen our understanding of exactly when partial coverage becomes large enough to change voucher lease-up, rents, or homelessness. Another practical challenge is ensuring a sufficient mass of observations at each relevant dose level, so that the nonparametric regressions are not overly sensitive to small-sample noise. Despite these complexities, the continuous-treatment framework promises greater realism and nuance, especially given that many CoCs are never purely "all in" or "all out" with respect to SOI policies.

## 6.2 Further Robustness Checks and Potential Mechanisms

Beyond shifting to a continuous treatment, a number of robustness checks can add credibility to the findings and clarify why rents sometimes climb while homelessness does not necessarily decline. One possibility is to introduce placebo or falsification exercises in which "never-treated" CoCs are assigned a pseudo-policy date. Re-estimating the model with these fabricated adoption dates helps distinguish actual policy effects from spurious variation in outcomes; if the placebo dates generate meaningful "treatment" effects, then pre-existing shocks may be confounding the identification.

Another strategy is to limit the sample in different ways, for instance by excluding CoCs that enact SOI rules prior to a certain year. This approach tests whether early adopters drive the overall results or whether short follow-up windows mask any delayed shifts in homelessness, rents, or voucher uptake. Such trimming of the sample can also reveal whether CoCs that adopt late look systematically different, perhaps due to political or economic conditions that influence both the legislation and the underlying housing market.

A complementary analysis would group CoCs according to pre-policy measures of market tightness or vacancy rates, so that the same staggered difference-in-differences estimates can be run in relatively slack versus highly constrained housing environments. If rent inflation is more acute in low-vacancy areas, that finding would align with basic supply-and-demand principles, reinforcing the plausibility that landlords respond to new mandates by adjusting lease terms upward. Establishing whether these outcomes are concentrated in tight markets is important, because it clarifies whether the policy's effects on homelessness might be modest or even counterproductive under specific conditions.

All of these checks serve a dual purpose. They probe the extent to which the identification is credible by examining whether treated and untreated CoCs experience idiosyncratic shocks that mimic or confound the policy timing. They also illuminate the mechanisms behind the surprising result that homelessness fails to decline significantly, despite improved access for voucher holders. By isolating subgroups or periods and observing whether effects remain robust, it is possible to discover whether particular subsets of CoCs are driving the apparent rent increases and correspondingly offsetting any homelessness gains. These sensitivity tests, when combined with deeper theoretical models of landlord behavior, offer a stronger basis for concluding whether SOI adoption genuinely leads to unintended rent hikes or whether other contextual factors—such as enforcement intensity, partial coverage, or local supply constraints—better explain why the expected drop in homelessness is not readily apparent.

## 6.3 Why Might Homelessness Rise (or Remain Unchanged)? A Sketch of the Landlord Response

Initially predicting that preventing landlords from discriminating against Housing Choice Voucher (HCV) recipients would reduce homelessness is straightforward: more voucher holders should be able to lease and maintain stable housing. However, preliminary results do not confirm a strong negative effect on homelessness. Indeed, the estimates show minimal short-run change in homeless rates, which in some cases hints at modest increases. This subsection provides a simple modeling perspective, partly inspired by Abramson (2022), on how landlord actions could undermine the direct benefits of an SOI policy.

### 6.3.1 Voucher Utilization and Landlord Incentives

Let $\eta_o = V_u/V_{Tot}$ represent the baseline voucher utilization (i.e., fraction of distributed vouchers that eventually secure housing). Suppose that once the SOI law passes, landlords can no longer reject applicants simply because they hold an HCV. The new utilization rate, $\eta_{SOI}$, can be written as:

$$\eta_{SOI} \;=\; \eta_o \;+\; \delta(\gamma,\pi)\Big(1-\eta_o\Big),$$

where $\delta(\gamma,\pi)$ is a parameter in $[0,1]$ capturing the expected strength of the law's enforcement. It depends on $\gamma$, the probability of being caught violating SOI rules, and $\pi$, the associated penalty. Higher $\gamma$ or $\pi$ means a stronger deterrent, thus $\delta(\gamma,\pi)$ is larger and the policy effectively increases $\eta_{SOI}$ more.

### 6.3.2 Homelessness and Rent Adjustments

Let $H$ denote the homelessness rate, defined in a simplified form as

$$H \;=\; (1-\eta)\,P\Big(y+v<r\Big) \;+\; P\Big(y<r\Big),$$

where $y$ is household income, $v$ is the voucher subsidy, $r$ is equilibrium rent, and $\eta$ is the share of voucher holders (proportionate to the population needing a voucher) who successfully avoid homelessness by securing a lease. The first term captures households who would otherwise need a voucher to afford rent but fail to utilize it (hence remain or become homeless). The second term captures those ineligible for vouchers but whose income still falls below the cost of housing.

At first glance, we might expect $\eta$ to increase under stricter SOI protections, thus reducing $H$. However, a crucial channel involves the possibility that landlords raise $r$ in response to perceived extra risk or regulatory burden. If $r$ increases significantly, some non-voucher households are priced out (the second term goes up), and *some* voucher holders might also fail to lease up when the higher rent exceeds the payment standard ($r > y + v$), or if the voucher requires tenant-supplied top-ups they cannot afford. Hence,

$$\frac{\partial H}{\partial r} > 0$$

for the income segments near threshold $r = y + v$, and for the segment near $r = y$. Landlords' rent-setting behavior can arise from a profit-maximization perspective if taking on voucher tenants implies additional administrative overhead or a perceived higher cost of default. Even if explicit discrimination is not permitted, the landlord can adjust *prices* to offset the new risk environment.

#### 6.3.3 Implications for Equilibrium Analysis

The short-run net effect on homelessness can be ambiguous. A sufficiently large increase in $\eta$ through stronger enforcement might lower $H$, but this effect can be partially neutralized if landlords collectively respond by raising $r$. The model suggests that policy-enforced acceptance of vouchers is *not* a silver bullet against homelessness in markets where supply constraints or landlord risk perceptions are strong. Over longer horizons, new supply or stronger enforcement could shift these dynamics, yet the magnitude of the price response remains an empirical question.

## References


Abramson, B. (2022). The welfare effects of eviction and homelessness policies. *Working Paper*.

Callaway, B. and Sant'Anna, P. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. (2024). Difference-In-Differences with a Continuous Treatment. *Working Paper*.

Desmond, M. (2016). *Evicted: Poverty and Profit in the American City*. Crown Publishers, New York.

Freeman, L., & Li, Y. (2013). Do Source of Income Anti-discrimination Laws Facilitate Access to Less Disadvantaged Neighborhoods? *Housing Studies*, 29(1):88–107.

Hangen, F. (2022) The Choice to Discriminate: How Source of Income Discrimination Constrains Opportunity for Housing Choice Voucher Holders. *Urban Affairs Review*, 59(5).

Jones, M. M., & Byrne, T. (2020). New faces of homelessness: Implications for research and policy. *American Journal of Public Health*, *110*(12), 1777–1778.

O'Flaherty, B. (2019). Homelessness research: A guide for economists (and friends). *Journal of Housing Economics*, *44*, 1–25.

Schwemm, R. (2016). Source-of-Income Discrimination and the Fair Housing Act. *Case Western Reserve Law Review*, 70(1):573–600.

Tighe, J. R., Hatch, M. E., and Mead, J. (2016). Source of income discrimination and fair housing policy. *Journal of Planning Literature*, 32(1):3–15.