Government assistance protects low-income families from eviction and rent nonpayment*

Ian Lundberg[ab], Sarah L. Gold[b], Louis Donnelly[bc], Jeanne Brooks-Gunn[d], and Sara S. McLanahan[abc]

[a]Department of Sociology, Princeton University
[b]Office of Population Research, Princeton University
[c]Center for Research on Child Wellbeing, Princeton University
[d]Teachers College and College of Physicians and Surgeons, Columbia University

Last updated: 6 January 2019. 7,557 words.

Contents

*Draft prepared for presentation at the 2019 Annual Meeting of the Population Association of America. Address correspondence to Ian Lundberg, Office of Population Research, Princeton University, Princeton, NJ 08540, ilundberg@princeton.edu. We thank Catherine Doren, the Stewart Lab, the Inequality Working Group, and a housing roundtable sponsored by the W.T. Grant Foundation at Princeton University for helpful comments on earlier drafts. Research reported in this publication was supported by the Robert Wood Johnson Foundation and by The Eunice Kennedy Shriver National Institute of Child Health & Human Development of the National Institutes of Health under Award Number P2CHD047879. Funding for the Fragile Families Study was provided through Award Numbers R01HD36916, R01HD39135, and R01HD40421 and by a consortium of private foundations. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Abstract

Unaffordable housing is among the most pressing problems facing low-income American families today, many of whom are evicted from their homes or live in fear of eviction due to missed rent payments. Housing assistance programs to address this problem include public housing and other assistance in which a government agency offsets the cost of private market housing. This paper assesses whether receipt of either category of assistance reduces the probability that a family will (a) not pay the full amount of rent or mortgage or (b) be evicted from their home in the subsequent six years. Because no randomized trial has assessed these two outcomes, we use observational data and formalize the conditions under which a causal interpretation is warranted. Families receiving assistance experience less hardship conditional on other variables, and we argue that the statistical evidence points toward a causal conclusion that assistance protects against rent nonpayment and eviction.

1 Introduction

Housing insecurity – both eviction and the threat of eviction due to missed rental payments – is a common experience among low-income families in the United States and has damaging ramifications. Nearly half of all renter households and more than 75 % of low-income renter households are housing cost burdened, spending more than 30 % of their incomes on housing (Joint Center for Housing Studies, 2018). The burden of housing costs translates for many into missed rent payments and eviction: in 2016 alone, 2.3 % of renter-occupied U.S. households were evicted (Desmond et al., 2018). The problem became especially severe during the housing bubble that led to the Great Recession, with the prevalence of eviction peaking at 3.1 % in 2006 (Desmond et al., 2018). Families with children face an especially high risk (Desmond et al., 2013), particularly those with low incomes: more than one in four children born into deep poverty in large U.S. cities in 1998—2000 were evicted by age 15 (Author redacted for peer review, 2019). The consequences of eviction are far-reaching and include further material hardship (Desmond and Kimbro, 2015), residential instability (Desmond, 2012; Desmond et al., 2015), and worse maternal and child health (Desmond and Kimbro, 2015). Future landlords may look unfavorably on a history of eviction, making it hard for those with an eviction to find stable housing in the future (Desmond, 2012). Housing assistance policies – which include public housing and government-subsidized private-market housing (e.g. vouchers) – may reduce eviction and help families to avoid the missed or partial rent payments that can lead to eviction. Though implementation varies across jurisdictions, these policies typically target assistance at families earning below 50 % of the area median income: precisely the low-income group at the greatest risk of eviction.

Despite the importance of housing policies, little empirical evidence has assessed their effects on eviction. This paper evaluates whether receipt of these forms of assistance, compared with no help, reduces (1) the probability of nonpayment of rent or mortgage, a common precursor to eviction, and (2) the probability of eviction, for a sample of low-income American families (the Fragile Families and Child Wellbeing Study, hereafter Fragile Families Study). We provide evidence that receipt of housing assistance when a child is approximately 9 years old is associated with lower risk of rent nonpayment and eviction between age 9 and 15, conditional on other variables. These associations are substantively large: we estimate that residence public housing is associated with a 41 % reduction in the risk of eviction, for instance. We state the assumptions under which this statistical evidence points toward a causal effect and assess the extent to which these assumptions would need to be violated to undermine our claims. While receipt of assistance does not eliminate the risk of hardship, we conclude that expanding assistance programs would likely reduce the prevalence of rent nonpayment and eviction for low-income American families.

1.1 Assistance may protect against eviction

Housing assistance — public housing and other government assistance for rent — may reduce housing hardship for several reasons: by improving families' financial wellbeing, by providing legal protections against eviction, and by connecting families to additional social services. While this paper does not attempt to distinguish these pathways, they provide theoretical reasons to suspect that housing assistance may reduce eviction.

In principle, housing assistance programs are designed to limit families' rental payments to approximately a third of their household income. The reduction in housing costs increases the resources families can allocate to other expenses, reducing the possibility that a large expense could prevent families from making a rent payment. In turn, families that pay their rent in full are less likely to be evicted (e.g., Desmond 2012). This mechanism may be of particular importance to low-income renters, about three-quarters of whom spend more than 30 % of their income on rent (Joint Center for Housing Studies, 2018), putting them one medical bill or one flat tire away from nonpayment of rent.

Housing assistance may also provide legal protections against eviction, though these protections are stronger among families in public housing than those with vouchers due to differences in program design. Public housing is owned and operated by public housing authorities and this housing stock is permanently affordable. By contrast, families who use vouchers to rent from landlords in the private market are not promised that the unit will remain perpetually affordable.

Families living in public housing can only be evicted if there is "good cause" which includes not paying rent, substantial violations of the rental agreement, or repeated minor violations of the rental agreement (Code of Federal Regulations, 1976). Additionally, public housing authorities are required to go through a judicial process to evict, which may make informal evictions less common in this context (Code of Federal Regulations, 1976). Though voucher recipients are also protected by "good cause" reasoning, landlords can choose to terminate their acceptance of a voucher at the end of a lease or after the initial lease period for personal or business reasons, rather than renewing (Code of Federal Regulations, 1994). This leaves families at risk of an unwanted move. In terms of legal protections, one might suspect that public housing would be most effective at preventing eviction.

The third mechanism through which housing assistance may protect families from eviction is by connecting them to other social services. Low-income families receiving housing assistance may also receive other supports such as child care subsidies, food assistance, and job training (e.g., Park et al. 2014). By reducing financial hardship, these programs may indirectly protect against missed rent payments and eviction.

On the other hand, it is possible that housing assistance may *not* protect families from eviction. By virtue of receiving assistance, families become more directly connected to the social welfare state and may become subject to increased surveillance. For example, if a family member is convicted of a crime, it is possible that this increased surveillance could result in an eviction that might not have occurred if the family rented without housing assistance. In spite of this possibility, we believe the opposing claim is more plausible: assistance protects families from eviction. This leads us to hypothesize that housing assistance provides a buffer against eviction.

1.2 Assistance is assigned partly by need and partly by luck

A thorough understanding of how families select into housing assistance is critical to any argument about the effect of housing assistance. Because each housing authority operates semi-independently, a universal statement about how housing assistance programs are allocated is not possible. Nonetheless, two key aspects are central to the assignment process: family eligibility and limited availability.

Whether a family qualifies for housing assistance is primarily determined by three eligibility criteria: annual gross income relative to family size, family composition, and immigration status. Families are eligible if their incomes fall below income limits that vary geographically based on the area median income, with the priority given to families below 50 % (very low income) and 80 % (low income) of the median income (U.S. Department of Housing and Urban Development, 2018). Those whose family composition includes someone who is elderly, disabled, or a child typically receive higher priority. Public housing authorities may also choose to prioritize other populations including homeless families or victims of domestic violence. Finally, the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA, welfare reform) limits federal assistance to citizens and those with legal immigration status; housing authorities may require applicants to provide evidence of this for all household members, though the specific requirements vary across housing authorities (McCarty and Siskin, 2012). Because state and local assistance are not regulated by PRWORA, these forms of assistance may be available to those who are undocumented.

Beyond eligibility criteria, the primary determinant of housing assistance is limited availability. Housing authorities typically receive far more applications for assistance than their resources can support, leading to long waiting lists that sometimes close when availability is severely limited (U.S. Department of Housing and Urban Development, 2018). In 2012, the last time national data on housing authority waiting lists was collected, only 4 % of public housing agencies reported that assistance was available without a waiting list, 90 % of agencies had open waiting lists, and 6 % of waiting lists were closed (Public and Affordable Housing Research Corporation, 2016). There is even greater unmet demand for vouchers: only 1 % of housing authorities reported voucher availability with no wait and almost half of waiting lists were closed (Public and Affordable Housing Research Corporation, 2016). In 2015, only 25 % of households whose incomes met the general cutoff for eligibility (below 50 % of area median income) received some form of government housing assistance (Joint Center for Housing Studies, 2018, p. 5). Because of limited availability, housing assistance is assigned partly by eligibility and partly by luck. Many housing authorities open the wait list in particular periods, accept applications, and then use a lottery to determine which applicants are permitted to join the wait list (Moore, 2016). The role of luck in determining the assignment of assistance is critical to identifying the causal effect of housing assistance on housing hardship.

1.3 Design-based evidence: Studies exploiting random assignment

Several prior studies exploit lotteries to provide experimental evidence that directly assesses the effect of housing assistance on a variety of child and family outcomes. This has been especially true for the effect of vouchers designed for use in the private market. In a housing lottery, families who enter the lottery are randomly chosen to receive housing assistance. Among those who enter the lottery, assignment of housing assistance is thus unrelated to families' potential outcomes in the absence of assistance, alleviating concerns about non-ignorable selection into housing assistance and identifying the causal effect. Causal inference is essential to policy in this setting: policies to expand housing assistance will only reduce housing hardship if the association between the two is causal. For this reason, studies exploiting random assignment are often considered the gold standard of evidence.

Despite the advantage of controlling the assignment of treatment, lottery-based studies face three critical limitations. First, follow-up surveys do not always record the housing hardship indicators of greatest interest to policymakers, such as eviction in the years immediately following the intervention. Second, the credibility of results is often sensitive to assumptions that attrition from follow-up can be ignored. Third, the comparison is often between families receiving different types of assistance (e.g. public housing vs. vouchers), rather than between families receiving a given type of assistance and those receiving no assistance at all. For these reasons, we treat experimental evaluations of housing assistance as informative but not as categorically superior to observational studies. Experimental and observational assessments each have strengths and weaknesses, and evidence from both sources should be considered when evaluating policies.

Much of the research making use of randomized designs comes from the Moving to Opportunity (MTO) experiment. The experiment began with 4,604 low-income families residing in public housing in Baltimore, Boston, Chicago, Los Angeles, and New York. Each family was randomly assigned to one of three treatment conditions: one group received a housing voucher that could be used only in low-poverty neighborhoods, a second group received a voucher that could be used anywhere, and a third group received no change to their assistance and continued to reside in public housing (Sanbonmatsu et al., 2011). After four to seven years, adults with vouchers reported improved safety and neighborhood satisfaction and better health while effects were mixed for children, with reductions in violent behavior for all but increased risky behaviors for boys and improved mental health for girls (Sanbonmatsu et al., 2011). After 10 to 15 years, housing and neighborhood conditions and social networks were better among those assigned to a voucher than those in the control group (Sanbonmatsu et al., 2011). Although initial results after 4–7 years offered no evidence that vouchers promoted academic achievement (Sanbonmatsu et al., 2006), later research using even longer-term MTO data found positive effects on college attendance and adult earnings for children who moved to low poverty neighborhoods before age 13 (Chetty et al., 2016).

The MTO study specifically examined eviction and rent nonpayment in a long-term follow-up survey conducted 10–15 years after treatment assignment. The study collected two measures of housing hardship: whether in the past twelve months participants had ever (1) been more than 15 days late on rent payments or (2) been threatened with eviction due to nonpayment. Those who had been assigned to either voucher condition were less likely to report either hardship in this 12-month period which occurred 10–15 years after the voucher was assigned, though the difference was statistically significant only for effects on nonpayment of rent (Sanbonmatsu et al., 2011). Three important considerations, however, limit the informativeness of this estimate. First, we might be most interested in the effect of housing assistance on housing hardship in the years *immediately* after housing assistance is assigned. The MTO results are informative instead about effects that occur more than a decade after housing assistance is assigned. Second, the benefits of random assignment are hindered by the possibility of non-ignorable attrition. The adult follow-up attempted to reach 4,142 adults (limited by funding constraints) and actually gathered data for 3,273 adults. Claims from the MTO adult follow-up are therefore causally identified only if the attrition of 869 families is independent of potential eviction in the absence of assistance. Third, the MTO study evaluates the effect of housing vouchers relative to public housing. The estimand of greatest interest, however, is the effect of each of these forms of assistance relative to no government assistance at all. Because the MTO sample was drawn among those residing in public housing, it cannot speak to this question. These limitations motivate the present analysis; while one would certainly prefer a randomized trial for the precise effect of interest with no attrition, even the MTO study does not provide such a trial.

Research examining other lotteries has found mixed effects of housing assistance on a variety of outcomes. A Chicago lottery in 1997 received 82,607 applications for vouchers, of whom a subset of 35,000 were randomly selected to receive assistance and 18,110 ultimately did receive assistance.[1] By merging the lottery list with administrative records, Jacob et al. (2014) found that receipt of a voucher had almost no effect on education, crime, or health outcomes, compared with residing in the private market with no voucher. In a similar study, the Department of Housing and Urban Development commissioned an evaluation of the effect of vouchers provided through the Welfare to Work program (Mills et al., 2006). Between 450 and 1,400 vouchers were assigned in each of six sites ranging from Atlanta, Georgia to Spokane, Washington. In each site, the study selected a sample of families eligible for Temporary Aid to Needy Families (TANF) who were eligible for but not receiving tenant-based assistance, with various other requirements (see Mills et al. 2006).[2] Although results are mixed for a variety of outcomes, vouchers did reduce the risk of homelessness (Mills et al., 2006), an outcome possibly related to nonpayment of rent and

[1] Jacob et al. (2014) note that vouchers were offered gradually as resources became available, and some of the 35,000 selected families never received a voucher because the agency ran out of resources. They do not make clear whether the 18,110 to ultimately receive vouchers (their treatment group) are a random subset of the 35,000 initially chosen for the treatment group.

[2] Mills et al. (2006) selected a sample of TANF-eligible households because these households would likely need assistance and could be easily identified through a sampling frame of current and former TANF recipients. The eligibility criteria for TANF vary by state; many but not all families eligible for housing assistance are also eligible for TANF.

eviction. As with MTO, the validity of these claims is subject to assumptions about the ignorability of attrition from follow-up.

Experiments have yielded rich information about the effects of housing assistance on children and families. Housing lotteries still fall short of gold-standard evidence, however, because of the possibility of non-ignorable attrition, because none record the outcomes of interest in our study (nonpayment of rent and eviction in the period immediately following receipt of housing assistance), and because lotteries do not always focus on a comparison group that receives no assistance at all. The only data available to answer this question at present is observational. For this reason, we turn to observational data to study the effect of housing assistance on housing hardship.

1.4 Prior observational evidence

One published study has examined the relationship between housing assistance and subsequent hardship in observational data. Using data on 417 families in the Detroit area in 2009–2011, Kim et al. (2017) find that receipt of housing assistance is associated with reduced risk of an aggregate indicator of housing insecurity which captures whether families experienced any of a range of outcomes: moving for cost reasons, foreclosure, eviction, homelessness, moving in with others to share expenses, or missing rent payments. The authors motivate the use of a comprehensive indicator because each specific hardship is rare and the overall sample size is small. Effects on an aggregate indicator represent an important first step, but policymakers may care about whether a policy is effective to reduce specific types of hardship such as eviction and nonpayment of rent or mortgage. The present study therefore draws on a larger sample from a broader population, examining outcomes over a longer time period to study effects of housing assistance on these two specific outcomes.

2 Data

Data come from the Fragile Families and Child Wellbeing Study (hereafter Fragile Families Study), a birth cohort panel survey which began with a probability sample of 4,898 children born in U.S. hospitals in cities with populations over 200,000 in 1998–2000.[3] Two aspects of the Fragile Families Study make it a good source for the study of housing hardship: the study design oversampled children born to unmarried parents, thereby producing a sample that tends to be disadvantaged and for whom housing hardship is common, and the study has consistently recorded indicators of eviction and nonpayment of rent or mortgage. Prior research has used these data to examine related questions such as the effect of paternal incarceration on housing hardship (Geller and Curtis, 2011; Geller and Franklin, 2014; Wildeman, 2014). Families were interviewed at the birth of the child and at approximately child ages 1, 3, 5, 9, and 15.

The full Fragile Families Study sample includes 4,898 families. We restrict to the subsample for whom the focal child resided with a responding mother or father at least half of the time at the age 9 interview (3,512, 72 % of baseline sample); this allows us to make use of the detailed set of variables asked of biological parents in every survey wave back to the birth of the child. Because homeowners are unlikely to seek assistance with rent, we restricted to parents who did not report owning their home (2,488, 71 % of those meeting prior criteria). Because families with consistently high incomes have no chance of qualifying for assistance, we restrict to those who report in at least one survey wave when the child was ages 1, 3, 5, and 9 that their income was 200% of the federal poverty threshold or less[4] (2,305, 93 % of those meeting prior criteria).

[3]The Fragile Families Study has two subsamples: a sample born in 16 cities that together form a probability sample of all newborns in the target population, and an additional sample born in 4 strategically chosen cities (Reichman et al., 2001). In the interest of statistical power, we use the full sample from all 20 cities to learn the association between pretreatment covariates $\vec{X}$ and housing hardship within each treatment group. We marginalize the resulting conditional average treatment effect estimates over the $\vec{X}$ distribution observed among the treated units in our sample (see Eq. 5). For an alternative version marginalizing over an estimate of the population $\vec{X}$ distribution, see Appendix C. Weighted estimates are substantively similar or suggest effects that are more protective, but are estimated with greater uncertainty.

[4]Because precision is a key concern for our analysis, we use a relatively high income threshold to maximize sample size and power. We also adjust for income relative to the poverty threshold in our regression models. Our use of a high threshold for inclusion also avoids omitting families who may have qualified for assistance with low incomes at a prior point and then maintained their assistance status despite income growth.

Finally, we restrict to families for whom the treatment variable (described below) is non-missing (2,219, 96 % of those meeting prior criteria). All missing values are multiply imputed with five imputations using the Amelia package in R (Honaker et al., 2011).

The treatment variable, housing assistance, is defined by the family's housing situation at the time of the age 9 interview and has three levels: public housing (321 families), receipt of federal, state, or local government support to help pay for a private rental (335 families), and no assistance (1,563 families). Throughout the paper, we treat receiving no assistance as the reference category and assess the causal effects of public housing and of other support relative to no assistance.[5]

We focus on two outcome variables[6] that capture events reported by the parent for the period between the interviews at child age 9 and 15: whether the parent ever (a) did not pay the full amount of rent or mortgage payments and (b) reported eviction from their home or apartment for nonpayment of rent or mortgage.[7] Missing reports of the outcome variable (384 cases for nonpayment and 383 cases for eviction) are excluded from model fitting but included in descriptive statistics with imputed values. By far the most common source of missingness is that the child's caregiver changed (377 cases) so that housing hardship at age 15 is unknown for the parent who

[5]Many but not all families in the no assistance category are eligible for assistance. We cannot determine eligibility with certainty because some determinants of eligibility, such as legal immigration status, are not available in our data. Further, eligibility guidelines vary across local jurisdictions (Moore, 2016). For instance, the federal government allows that "a PHA may adopt local policies permitting the admission of additional categories of low-income families to address essential local housing needs," (U.S. Department of Housing and Urban Development, 2001, ch. 5, p. 3). Area differences in eligibility represent a potentially exogenous source of variation in whether a given family receives assistance, so the inclusion of some families that are not eligible in the reference category may be beneficial for identification.

[6]One may additionally wonder whether housing assistance reduces these outcomes solely by reducing rental costs. Our data are not optimally suited to answer this question because rental costs are only recorded for a subset of respondents. Nonetheless, we do find that those receiving either type of assistance report lower rental payments than those not receiving assistance. See Appendix E for details.

[7]All variables, including the outcomes, are self-reported. We suspect that some respondents ignore the end of the question and report eviction regardless of whether the reason was for nonpayment of rent or mortgage. Among those reporting nonpayment of rent, 132 (26 %) report being evicted for nonpayment. Among those reporting that they always paid the rent, 14 (1 %) report that they were evicted for nonpayment. The latter group is logically impossible, and we suspect these individuals either misreported their consistent payment or misunderstood the eviction question and were evicted for other reasons. We nonetheless accept all responses and note that our outcomes are self-reported. We also note that the wording of the question (eviction for nonpayment of rent or mortgage) negates the utility of investigating whether nonpayment mediates the effect of assistance on eviction. By construction, this type of self-reported eviction is logically impossible (and empirically rare) among those who pay the rent each month.

was the caregiver at age 9. Because this could be a consequence of eviction, we conduct sensitivity analysis to address the possibility of non-ignorable missingness (see Appendix A).

We condition on a set of pre-treatment variables; we discuss the measurement of these variables here but focus on their use in causal identification in the next section. All variables are reported by the parent with whom the child lived at the age 9 interview. We include lagged indicators of the outcome variables: whether the parent reported not paying the full amount of rent or mortgage and whether they reported eviction, in (a) the 12 months immediately preceding the age 9 survey when treatment is defined and (b) at any survey wave prior to this.[8] We measure family income relative to the poverty threshold in the year preceding interviews at child ages 1, 3, 5, and 9. A disability variable indicates whether this parent reported a health condition that limited the type or amount of work they could do at age 9. We measure criminal conviction by whether the parent reported being convicted of any charges beyond minor traffic violations between the child's 1st birthday and the age 9 interview. We include whether the parents were married at the child's birth, the parent's race (black, white, Hispanic, or other), and the parent's education (less than high school, high school, some college, or a college degree). We include two scaled scores from assessments of the parent when the child was approximately three years old: a cognitive score from a modified version of the Weschler Adult Intelligence Scale (Wechsler, 1981) and a score on a modified version of Dickman's (1990) impulsivity scale. We refer the reader to the survey documentation (Fragile Families and Child Wellbeing Study, 2006) for further details on these scales as implemented in the study.[9]

[8]We keep the age 9 eviction indicators separate because they may be especially important sources of confounding. Because eviction and nonpayment are rare and prior indicators may be less relevant to confounding, aggregating prior measures may improve precision with minimal harm to identification.

[9]There are variables we explicitly do not include in the conditioning set. Our principle for deciding what to condition on was based on the notion that there are two types of variation in housing assistance: confounding variation that leads to biased estimates (i.e. X in Fig. 2) and identifying variation which helps us capture the causal effect (i.e. Z in Fig. 2). We aim to net out the confounding variation but not the identifying variation. We condition on income because this is likely a source of confounding variation: it affects assistance and also affects eviction directly. We do not condition on city of residence because this may be primarily a source of identifying variation: the available stock of public housing and vouchers may produce variation in assistance but do not directly affect whether a family would be evicted without assistance. One can argue the opposite: some cities have stronger rent markets, for instance, which might induce landlords to evict at higher rates in order to replace tenants with higher-paying alternatives. In this case, city of residence may be confounding. Despite this possibility, we believe that differences in assistance across cities help us to identify the effect more than they hurt, so we do not condition on city of residence.

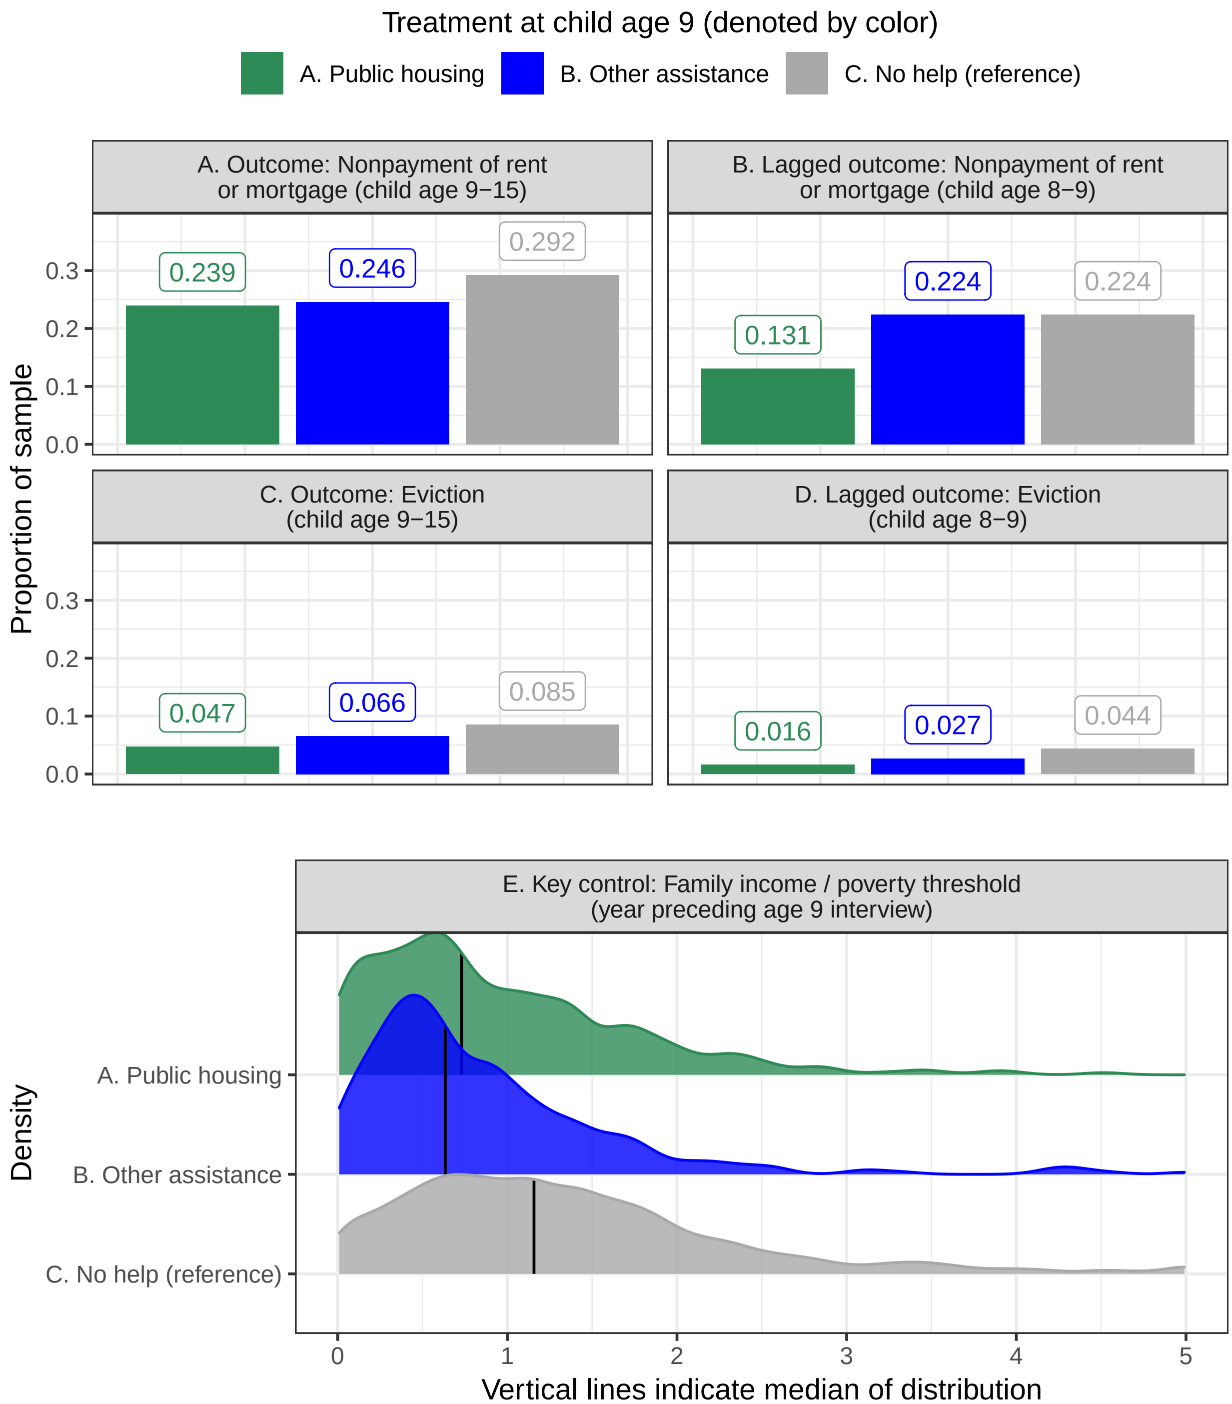

Fig. 1. Descriptive statistics for analytic sample (N = 2,219). Missing values are imputed. All estimates are unweighted. Incomes are top-coded at five times the poverty line.

Figure 1 provides descriptive statistics for key variables. The two outcome variables, nonpayment of rent or mortgage and eviction at child ages 9-15, are most common among those not receiving housing assistance at age 9 (Panels A and C). Countervailing selection processes, however, make it unclear how these raw associations relate to the causal effect of housing assistance on housing hardship. On one hand, families residing in public housing are markedly less likely to have not paid the full amount of rent or mortgage in the *previous* twelve months than those receiving no assistance (Panel B), and families receiving either form of assistance are less likely to have been evicted (Panel D). These associations may arise if housing authorities screen tenants based on their records with previous landlords when allocating units in public housing and if private landlords are less willing to accept a voucher tenant if that tenant has a history of eviction. These possibilities would produce positive selection into housing assistance: those who receive assistance are less likely to be evicted even in its absence. On the other hand, those receiving either form of assistance have lower family incomes relative to the poverty threshold (Panel E). Because assistance programs are targeted at low-income families, selection into housing assistance may be negative with respect to one's potential to avoid eviction or nonpayment in the absence of help. Overall, the imbalance on pre-treatment covariates suggests that adjustment may be necessary to support a causal interpretation of the association between housing assistance at age 9 and subsequent housing hardships.

3 Causal identification

The assumptions required for causal inference from observational data are strong, yet evidence-based policymaking demands an answer to the causal question of how housing assistance affects housing hardship. If housing assistance is merely associated with reduced hardship but the association is not causal, then expansion of housing assistance programs would represent a waste of resources that may be better spent elsewhere. If a strong causal effect exists, then expansion of assistance programs would produce real benefits for American families

For this reason, our discussion of the required assumptions and estimation strategies is substantial. We begin by defining the causal goal: the average *effect* of each type of housing assistance on housing hardship. We refer to this quantity as $\tilde{\tau}_d$. It is well-known that a causal effect cannot be identified from data alone. We therefore separately define the statistical quantity we estimate: the average *association* between each type of assistance and housing net of other variables. We call refer to this quantity as $\tilde{\theta}_d$. Formal definitions of these two quantities enable us to speak clearly about potential situations that would undermine our interpretation of the statistical evidence as supporting a causal conclusion. Finally, we discuss the particular strategy we use to *estimate* those associations. To estimate the association, we (1) fit an OLS regression model of housing hardship among those not receiving assistance, (2) fit the same model among those receiving assistance, and (3) estimate the difference between the predicted values from (1) and (2) at the predictors observed among those receiving assistance. To be clear that our estimates may deviate from the truth due to random chance, we distinguish our estimate of the association with a hat ($\hat{\tilde{\theta}}_d$ vs. $\tilde{\theta}$). This level of formality is standard in statistics but may be new to social scientists. By following this practice, we aim to be transparent about the gaps between a precise causal goal $\tilde{\tau}_d$, the quantity our statistical procedure is able to estimate $\tilde{\theta}_d$, and the actual number reported in the paper $\hat{\tilde{\theta}}_d$. We hope that this rigor of discussion will make transparent how our estimates could be wrong and will serve as an example for substantive scholars researching causal questions in the social sciences.

3.1 Defining estimands

We define the causal effect of housing assistance on eviction using the potential outcomes framework, which formalizes causal effects in terms of what would have happened to an individual if that person received different types of assistance (Neyman, 1923; Rubin, 1974; Imbens and Rubin, 2015). We define the treatment variable D_i based on the family's residence at age 9: $D_i = (\text{Public housing})$ for those residing in public housing, $D_i = (\text{Other assistance})$ for those receiving another form of federal, state, or local assistance for rent, and $D_i = 0$ for those receiving no government assistance for rent. We denote the number of families in each treatment category d in the population by $n_d = \sum_{i=1}^{n}(D_i = d)$. Each family i has three potential outcomes: whether the family would be evicted by the following interview if they lived in public housing Y_i(Public housing), if they received other assistance Y_i(Other assistance), or if they received no help $Y_i(0)$. The potential outcomes are deterministic functions of the treatment variable D_i; the observed outcome is $Y_i = Y_i(d)$ if $D_i = d$ for each treatment value d.[10] We assume that each individual's potential outcomes are a function of their own treatment assignment only (no interference, the Stable Unit Treatment Value Assumption, Imbens and Rubin 2015). For each individual, treatment effects are the differences between these potential outcomes. For instance, the effect of public housing for person i is Y_i(Public housing) $- Y_i(0)$.

Eq. 1 defines the conditional causal effect of housing assistance on eviction as a function of a set of pre-treatment covariates $\vec{x}$. Among families with pre-treatment covariates $\vec{x}$ who receive assistance type d, the causal effect captures the difference between probability that a randomly sampled family is evicted with assistance $Y(d)$ compared to the probability that would persist if they did not receive assistance $Y(0)$.[11] The function notation $\tau_d(\vec{x})$ allows that treatment effects may be heterogeneous as a function of the pre-treatment variables $\vec{x}$.

$$\tau_d(\vec{x}) = \text{P}\left(Y(d) \mid D = d, \vec{X} = \vec{x}\right) - \text{P}\left(Y(0) \mid D = d, \vec{X} = \vec{x}\right) \tag{1}$$

The conditional causal effect is not directly estimable because no families receiving assistance are ever observed in the condition of no help, making the data alone uninformative about $\text{P}\left(Y(0) \mid D = d, \vec{X} = \vec{x}\right)$ when $d \neq 0$. Instead, we define a statistical function $\theta_d(\vec{X})$ which can be estimated from data.

[10]This assumption is sometimes called an assumption of consistency. Because the "other assistance" category may include many types of assistance, consistency is less certain in this case. For this reason, our policy recommendations will focus on public housing, for which the treatment is more sharply defined.

[11]Although the potential outcomes are fixed, we use probability notation as a shorthand to state that sampling variability creates some probability that a given sampled unit with covariates $\vec{x}$ and treatment d would be evicted under a given treatment assignment.

$$\theta_d(\vec{x}) = \text{P}\left(Y(d) \mid D = d, \vec{X} = \vec{x}\right) - \text{P}\left(Y(0) \mid D = 0, \vec{X} = \vec{x}\right) \tag{2}$$

The statistical function $\theta_d(\vec{x})$ is similar to the causal function $\tau_d(\vec{x})$ but differs in a critical way: it replaces the potential outcomes $Y(0)$ among the treated ($D = d$) with the observable outcome $Y(0)$ for families who are not treated ($D = 0$). The well-known adage that correlation is not causation is true because the functions $\theta_d(\vec{x})$ and $\tau_d(\vec{x})$ are not equivalent. Our theoretical arguments about identification focus on the conditions under which $\theta_d(\vec{x}) = \tau_d(\vec{x})$. In a separate section on estimation, we discuss our empirical approach to estimate $\theta_d(\vec{x})$.

The conditional causal effect and the statistical function may take different values at different pre-treatment covariates $\vec{x}$. Our ultimate goal is to provide a single-number summary of the causal effect, so we marginalize these quantities over the distribution of covariates $\vec{x}$ observed among the treated cases ($D_i = d$) in the sample (indicated by $S_i = 1$). We refer to the resulting summaries as the causal estimand $\tilde{\tau}_d$ and the statistical estimand $\tilde{\theta}_d$. For an alternative marginalization that incorporates survey weights to draw population inferences, see Appendix C.

$$\underbrace{\tilde{\tau}_d = \frac{1}{n_d} \sum_{i:D_i=d,S_i=1} \tau_d\left(\vec{x}_i\right),}_{\substack{\textbf{Causal estimand} \\ \text{Average treatment effect} \\ \text{on the treated}}} \qquad \underbrace{\tilde{\theta}_d = \frac{1}{n_d} \sum_{i:D_i=d,S_i=1} \theta_d\left(\vec{x}_i\right)}_{\substack{\textbf{Statistical estimand} \\ \text{Average treatment association} \\ \text{for the treated}}} \tag{3}$$

3.2 Point identification

Under what conditions are the causal estimand ($\tilde{\tau}_d$, the *effect* of assistance on housing hardship) and the statistical estimand ($\tilde{\theta}_d$, the *association* between assistance and housing hardship) equivalent? We reason about these conditions in a framework that combines potential outcomes (Imbens and Rubin, 2015) with causal graphs (Pearl, 2009): Single World Intervention Graphs (SWIGs, Richardson and Robins 2013). The SWIG in Figure 2 depicts a hypothetical world in which we have intervened to assign all families to the condition of no assistance. The intervention (denoted by the broken oval) breaks the causal effect of the naturally-occuring treatment D

(e.g. public housing) on the outcome Y (e.g. eviction) so that any remaining association arises purely from confounding. The key identification question is whether assistance $[D]$ and hardship in the absence of assistance $[Y(0)]$ are spuriously related after the causal effect has been broken by intervention.

In Panel A the causal effect is identified. By conditioning on the pre-treatment variables $\vec{X}$, we can block all backdoor paths linking housing assistance D and housing hardship $Y(0)$. Conditional on these variables, we must assume that all determinants of housing assistance receipt D operate like Z; they do not directly affect housing hardship. If Z were measured, it could serve as an instrumental variable. In this study, we believe that an important source of variation Z remains after conditioning on $\vec{X}$: limited availability and lotteries. Because the demand for housing assistance far exceeds the availability, whether a given family receives assistance may be due in part to luck rather than to unobserved factors which also affect the outcome directly. Our theoretical confidence that Z exists adds credibility to our selection-on-observables design.

The assumptions in Panel A, however, may be optimistic because they omit unobserved variables that confound treatment assignment. Drawing causal inferences from observational data requires transparent acknowledgment of how the required assumptions could be violated. For this reason, we next discuss two concrete situations in which the effect is not identified (Panels B and C). We later report results with sensitivity analysis to assess the possibility that associations are driven by unobserved confounding.

Our first example (Panel B) is one of classic omitted variable bias. Suppose that families with the skill set to navigate the housing authority and secure assistance (D) are also be well-equipped to negotiate with a landlord and avoid an eviction even without assistance $Y(0)$. We refer to this skillset as cultural capital (U_1). In this case, we might find a statistical association between assistance (D) and housing hardship (Y) even if this association is not causal. Because cultural capital would be very difficult to measure, we instead treat it as an unobserved variable to be considered as a possible violation of our assumptions in sensitivity analysis.

Panel C clarifies how unobserved variables may produce confounding in more complex

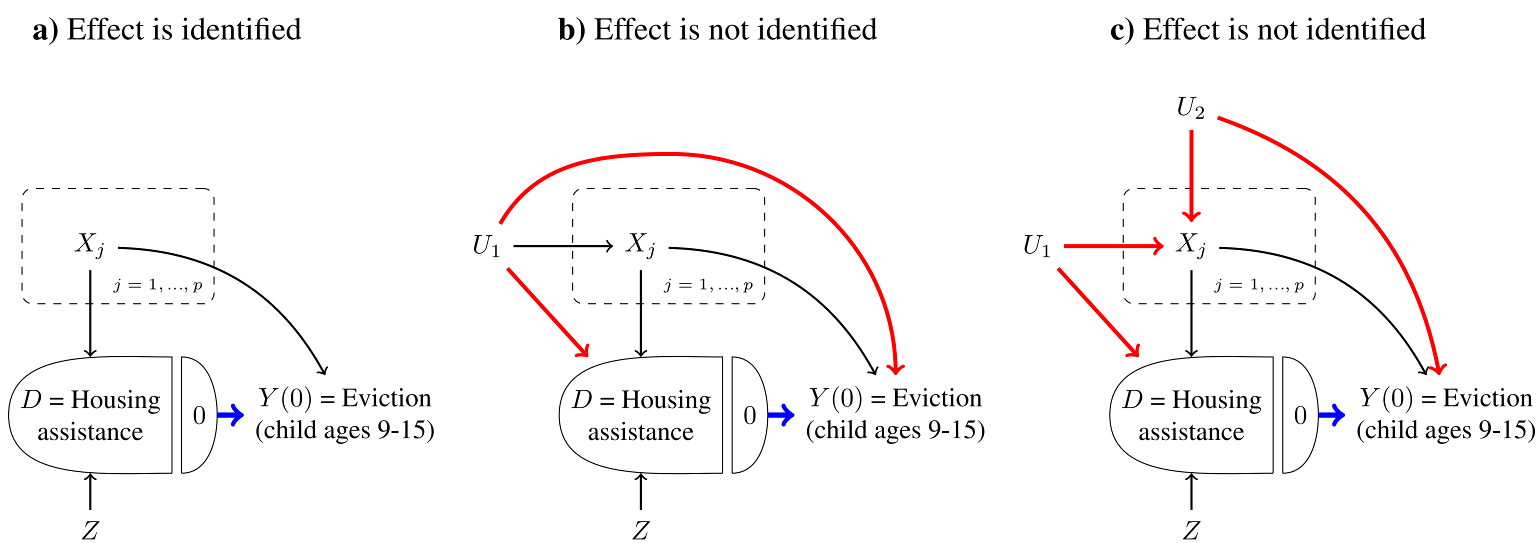

Fig. 2. Single World Intervention Graph (SWIG) under which we seek to identify the causal effect of housing assistance on eviction (thick blue edge). $\vec{X}$ variables include whether parents were married at the child's birth, the parent's race, education, impulsivity, cognitive score, criminal conviction, disability, family income, nonpayment in past 12 months and at any prior report, and eviction in past 12 months and at any prior report. Although we omit to simplify notation, all of the $\vec{X}$ variables could affect each other in arbitrary ways without threatening the identification assumptions. A third violation ignored in this figure may arise if missingness on the outcome variable is non-ignorable; for a discussion of this problem, see Online Supplement Part A.

ways. Suppose that cultural capital (U_1) affects receipt of housing assistance directly, but only affects $Y(0)$ through its effect on measured variables $\vec{X}$ such as nonpayment of rent in the 12 months preceding the age 9 interview. Suppose that gentrification in a family's neighborhood (U_2) makes rent go up so that they miss a payment (an element of $\vec{X}$), and gentrification also makes the landlord more likely to evict in the future $Y(0)$ in order to lease to a higher-paying tenant. In this case, prior nonpayment of rent (an element of $\vec{X}$) is a collider variable, and conditioning on it induces an association between cultural capital (U_1) and gentrification (U_2). Gentrification and cultural capital may be independent in the population, but holding prior nonpayment constant, those who keep up with rent despite gentrification must have high cultural capital.[12] Conditioning on nonpayment in the set $\vec{X}$ is essential because this variable causes the treatment D and the outcome $Y(0)$ in the absence of assistance, yet conditioning on $\vec{X}$ creates a spurious association through U_1 and U_2. Without measuring an instrument like Z, there is no solution to this problem aside from sensitivity analysis.

We outline the two cases above to be clear about how our estimates could be misleading, but in this application violations from variables like U_1 and U_2 are likely small. We believe that measured sources of confounding $\vec{X}$ (i.e. family income) and unmeasured instruments Z (i.e. limited availability) are the overwhelming determinants of housing assistance, so an estimate that adjusts for $\vec{X}$ approximately identifies the causal effect. Providing this estimate is important for evidence-based policy, even if identification is imperfect. Nonetheless, the violations depicted in Figure 2 Panels B and C are likely to exist, at least to a small degree. We therefore present sensitivity analyses to assess the size of these violations that would be required to undermine our causal conclusions.

[12]This example is a case of M-bias; see Greenland et al. (1999) and Pearl (2009) for similar examples.

4 Estimation

Even if the statistical estimand $\tilde{\theta}_d$ identifies the causal effect $\tilde{\tau}_d$, additional assumptions are required to produce an estimate $\hat{\tilde{\theta}}_d$ from available data. To do so, we first estimate the conditional treatment effect function $\hat{\theta}_d\left(\vec{x}\right)$ by replacing the conditional probabilities with estimates, denoted $\hat{\mathrm{P}}(Y \mid \bullet)$.

$$\hat{\theta}_d(\vec{x}) = \hat{\mathrm{P}}\left(Y \mid D = d, \vec{X} = \vec{x}\right) - \hat{\mathrm{P}}\left(Y \mid D = 0, \vec{X} = \vec{x}\right) \tag{4}$$

Any statistical or machine learning approach can be used to estimate $\hat{\mathrm{P}}$. We focus on the simple case of a linear probability model estimated by ordinary least squares (OLS).[13] OLS produces unbiased predicted probabilities only if the parametric specification is correct; the Online Supplement Part D reports substantively similar results estimated by a nonparametric machine learning approach that relaxes this assumption. The coefficients of the OLS models capture the association between pre-treatment variables and the outcome within each treatment group; because these coefficients have no causal interpretation, they are omitted from the main text and presented in Appendix B.

Model-based imputation of the outcome under various treatment conditions is sometimes called the parametric g-formula (Hernan and Robins, 2018) or the imputation estimator (Hahn, 1998; Abadie and Imbens, 2006, 2011). We note that this procedure is distinct from the more common practice of estimating a regression model conditioning on both the treatment and control variables. The more common regression approach is optimally efficient if the treatment effect is constant across the population. If the effect of housing assistance varies by pre-treatment variables, however, then the coefficient on housing assistance in such a model would represent a weighted average treatment effect with weights proportional to the variance of housing assistance that remains conditional on these variables (Angrist and Krueger, 1999). It is not clear why one would care theoretically about this weighted average. The procedure outlined above instead yields a consistent estimator for the average treatment effect as long as the conditional expectation function is

[13] We use OLS instead of a logit or probit model because it is an unbiased estimator, and because in this application it is not consequential if some predicted probabilities fall outside the [0,1] interval.

linear in family income, even if treatment effects are heterogeneous.

Finally, we marginalize over the distribution of $\vec{X}$ observed among the treated units in our sample (with sampling indicators denoted by S_i).

$$\hat{\hat{\theta}}_d = \frac{1}{\sum_{i:D_i=d} S_i} \sum_{i:D_i=d,S_i=1} \hat{\theta}_d\left(\vec{x}_i\right) \tag{5}$$

Our decision to marginalize over the observed distribution is driven by two considerations. First, as in a randomized experiment, our aim in this paper is to maximize internal validity and clearly state the conditions under which a causal effect can be inferred in our sample. By not inferring for a population, we avoid additional complications related to sampling. Second, statistical precision is a prime concern in our setting. Because a weighted estimator may be less efficient, we prefer an unweighted estimator. For an alternative specification that marginalizes over an estimate of the distribution of $\vec{X}$ in the population, see Appendix C.

5 Results

After adjusting for differences in pre-treatment variables by the OLS imputation estimator, residence in public housing is conditionally associated with reductions of 2.4 and 3.2 percentage points in the probability of nonpayment and eviction, respectively, and receipt of other assistance is conditionally associated with reductions of 5.2 and 3.0 percentage points in the probability of nonpayment and of eviction. These associations are depicted graphically in Figure 3.

Although the 95 % confidence intervals for the estimates include zero in three out of four cases (Fig. 3), it would be incorrect to interpret these results as null. Substantively, the point estimates are large, albeit estimated in a small sample (see Table 1). Public housing and other assistance reduce the prevalence of nonpayment of rent or mortgage by 10 and 18 %, respectively; in our sample, we expect that 26.1 more families would experience nonpayment if these programs did not exist. Public housing and other assistance reduce the prevalence of eviction by 41 % and 32 %, respectively; in our sample, we expect that 20.7 more families would be evicted if these

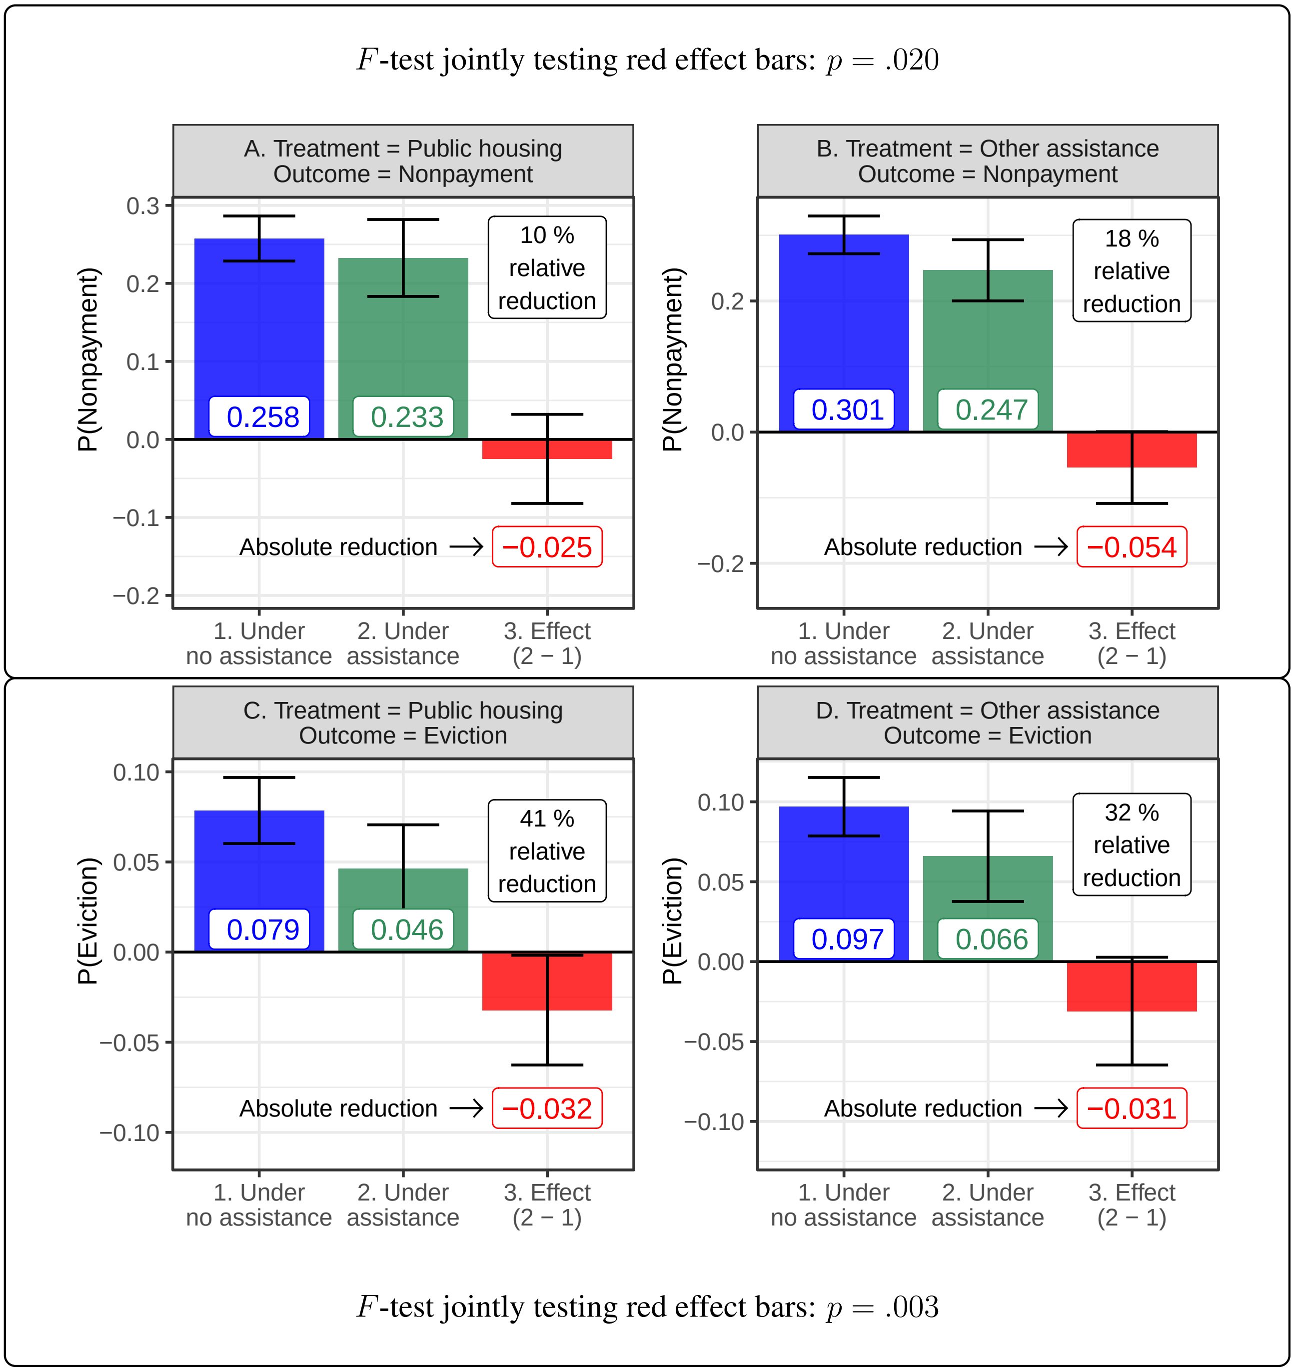

Fig. 3. Conditional associations between housing assistance and eviction. Associations can be interpreted causally under the identification assumptions discussed above. Although only one out of the four effects is statistically significant, for both outcomes we can reject the joint null hypothesis of no average treatment effect for either of the treatments.

policies did not exist. Housing assistance does not eliminate housing hardship, but it may place a serious dent in this problem.

Outcome	Treatment	Number treated	Expected # with hardship under no help	Expected # with hardship under treatment	Number protected	Percent reduction due to treatment
Nonpayment	Public housing	321	82.7	74.7	8.0	10 %
Nonpayment	Other assistance	335	100.2	82.7	18.1	18 %
Eviction	Public housing	321	25.2	14.9	10.3	41 %
Eviction	Other assistance	335	32.2	22.1	10.4	32 %

Table 1. The effects of housing assistance on rent nonpayment and eviction are substantively large. Future research is needed with larger samples, but point estimates suggest that each type of housing assistance reduces the number of evictions by more than a third. Values in table are rounded.

Statistically, it is unlikely that such consistent negative associations would appear across both treatment conditions for each outcome. A joint F-test shows that the probability of finding the set of estimates $(\hat{\theta}_{\text{Public}}, \hat{\theta}_{\text{Assistance}})$ this extreme if both effects are zero is small: $p = .020$ for effects on nonpayment and $p = .003$ for effects on eviction. The substantively large size of these protective effects and the unlikeliness that such consistent results would arise by chance suggest that housing assistance programs my be remarkably effective policy levers to reduce housing hardship.

5.1 Sensitivity

As discussed in the section on identification, a causal interpretation of $\hat{\hat{\theta}}_d$ is warranted only if the conditioning set (the pre-treatment variables) blocks all backdoor paths linking treatment assignment to the outcome of interest. Because this may not be the case, we discuss the sensitivity of results to unobserved confounding. For this exposition, we focus on the effect of public housing receipt on eviction. The logic to assess the sensitivity of claims about the effect of other assistance (e.g. vouchers) is analogous. We conduct the following thought experiment. Imagine selecting at random (a) a family residing in public housing and (b) a family not receiving assistance whose pre-treatment variables $\vec{X}$ match those of family (a). Suppose we intervene to revoke the assistance of (a), sending the family into the private market with no assistance. By what percent is the resulting probability of eviction for (a) less than for (b)?

We refer to this percent as $\Delta_d(\vec{x})$ and assume for simplicity that it is constant regardless of pre-treatment variables: $\Delta_d(\vec{x}) = \tilde{\Delta}_d$ for all $\vec{x}$. How big would $\tilde{\Delta}_d$ have to be to undermine our conclusion that public housing protects against eviction?

$$\Delta_d\left(\vec{x}\right) = \frac{\overbrace{\text{P}\left(Y(0) \mid D=0, \vec{X}=\vec{x}\right)}^{\text{(b)}} - \overbrace{\text{P}\left(Y(0) \mid D=d, \vec{X}=\vec{x}\right)}^{\text{(a)}}}{\underbrace{\text{P}\left(Y(0) \mid D=0, \vec{X}=\vec{x}\right)}_{\text{(b)}}} \tag{6}$$

Figure 4 presents this sensitivity analysis. The x-axis shows the sensitivity parameter $\tilde{\tau}_d$: the difference (b) - (a) discussed above. The appropriate value of this parameter is not estimable from the data, and each reader may have a distinct belief about an appropriate value. The y-axis shows the true causal effect $\tilde{\tau}_d$ implied by the point estimate from the data and the sensitivity parameter on the x-axis. In all plots, the line has a positive slope: the true effect of housing assistance implied by the data is closer to zero if the sensitivity parameter is greater (i.e. if those residing in public housing are more positively selected). In each panel, we highlight two values: the point estimate reported previously, which assumes a sensitivity parameter of zero, and the critical value at which the sensitivity parameter is large enough to nullify our causal conclusions. We argue that it is implausible that the sensitivity parameter is as high as this critical value.

Our most robust result is the effect of public housing on eviction. If we intervened to remove families from public housing, their probability of subsequent eviction would need to remain $\tilde{\Delta}_d = 41$ % lower than the probability among similar families who we observe without assistance for our point estimate to be consistent with no causal effect. To argue against our causal interpretation, one must believe that unobserved variables such as cultural capital are enormously important. We find this implausible, but we note that the plausibility of this degree of confounding is a subjective judgment. Our least robust result is the effect of public housing on nonpayment of rent or mortgage. If those receiving public housing would continue to have $\tilde{\Delta}_d = 9$ % lower probabilities of nonpayment than those receiving no assistance even if we removed these families from public housing, then our estimate is consistent with zero causal effect. This amount of confounding

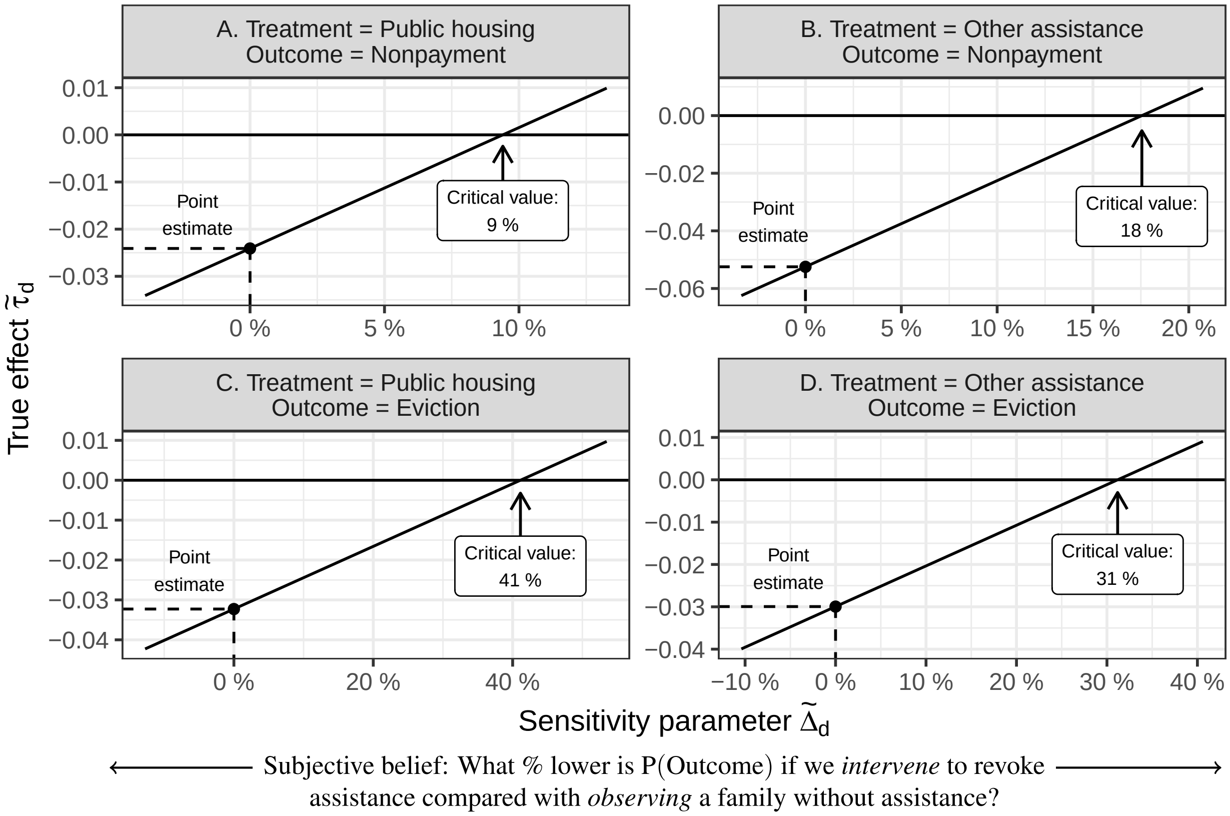

Fig. 4. Causal effect $\tilde{\tau}_d$ as a function of the degree of confounding $\tilde{\Delta}_d$ (see Eq. 6). The critical value represents the degree of confounding required to undermine a causal interpretation of our point estimate.

may be possible. The sensitivity of claims about other assistance fall between these two extremes. Overall, the sensitivity analysis in Figure 4 suggest that housing assistance is protective at a wide range of beliefs about the degree of confounding.

6 Limitations

Our claims are limited by the possibility of unobserved confounding, by limited statistical power, and by potential measurement problems. In addition, our claims apply to a period of particular interest — the aftermath of the Great Recession — but may not generalize to other periods.

The most well-known limitation of any study seeking to establish a causal effect with observational data is the possibility of unobserved confounding. Because we suspect that some unobserved variables which affect housing assistance may also directly affect the probability of hardship in the absence of assistance, we present sensitivity analyses that allow readers to draw their own conclusions based on their beliefs about the extent of such confounding. We believe that the amount of confounding required to invalidate our claims is implausible, but we can never rule it out. We therefore call for future research with randomized treatment assignment to provide more definitive evidence. Although observational data are limited, they offer important insights that represent a critical first step toward a more complete understanding of the effect of housing assistance programs.

Second, our estimates suffer from limited statistical power. Although for each outcome we can reject the null hypothesis that the effects of public housing and other assistance are both zero, our estimate of each individual effect is noisy. Because questions about eviction are not commonly included in household surveys or studies of housing assistance, future research to provide more precise estimates will require new large-scale data collection efforts.

Third, our measurement (a survey question) may understate the prevalence of nonpayment and of eviction. Previous research has demonstrated that tenants who are evicted may not report their experience as such in a survey, possibly due to disagreement about what constitutes an eviction, social desirability bias, or a desire to portray maximum control over their lives (Desmond et al., 2015; Desmond and Shollenberger, 2015; Desmond and Gershenson, 2017). Future research using new data sources asking multiple questions about a variety of forced moves (i.e. Desmond and Shollenberger 2015) will be needed to overcome this limitation. The question wording in our sample also limits us to a focus on how housing assistance affects whether families ever miss the rent or are ever evicted; future research on the number of such events could provide important insights into effects on serial nonpayment and serial eviction. Finally, the survey question includes all evictions for nonpayment of rent or mortgage; future research is needed to assess effects that are specific to eviction from rental units.

Finally, the results from these analyses may be particular to the aftermath of the Great Recession. In the period after the Great Recession, the housing affordability crisis in the United

States worsened and the gap between the supply and demand for affordable housing widened (Joint Center for Housing Studies, 2011). Even among families with housing assistance, rent burden was common during this period, with about 59 percent paying at least a third of their income towards housing and a third paying more than 50 percent (Joint Center for Housing Studies, 2011). In this environment, even housing assistance may not have been enough to prevent nonpayment of rent. It is possible that housing assistance may more effectively prevent housing hardship in other periods in which it more effectively reduces the proportion of income spent on rent.

7 Discussion

Evidence-based policymaking demands an answer to the question of how housing assistance affects housing hardship. Is housing assistance merely associated with reduced housing hardship, in which case expanding assistance may be a waste of resources? Or is the association causal, in which case expanding assistance may produce real benefits for American families struggling to make ends meet? Randomized trials have assessed effects of housing assistance on more distant outcomes, such as child behavior and adult earnings, with mixed results. By focusing on two more proximate outcomes, we shift attention to domains of family well-being that are important in their own right, and for which housing assistance programs may be especially effective.

We find that families receiving housing assistance have lower risk of nonpayment of rent or mortgage and lower risk of eviction than similar families not receiving assistance. Receipt of public housing and of other government support are associated with 41 % and 32 % reductions in the probability of eviction, respectively. These programs are associated with 10 % and 18 % reductions in the probability of nonpayment of rent or mortgage. A secondary contribution of this paper is to clarify the conditions under which these statistical associations can be interpreted causally. Under a range of reasonable assumptions, our results support the claim that housing assistance protects against housing hardship. Our estimates are substantively large: we find that public housing reduces the probability of eviction, for instance, by more than a third. Future research should re-evaluate these effects with randomized trials, but in the absence of this information we believe policymakers should be informed by the best available evidence from an observational study.

Our results point directly toward policy recommendations. Because the treatment category of other housing assistance includes both person-based (vouchers) and place-based assistance, we do not make policy suggestions for this group. Future research better delineating between vouchers and place-based housing other than public housing can help establish the basis for policy recommendations for those programs (i.e. which specific programs should be expanded). We can, however, advocate the expansion of public housing to prevent nonpayment and eviction. Clearly, the need for public housing far outstrips availability: only 4 percent of all public housing authorities in 2012 did not have waiting lists for public housing (Public and Affordable Housing Research Corporation, 2016). If public housing were expanded to serve more families, then nonpayment of rent or mortgage and eviction may both become less common. While a single observational study cannot establish this effect with certainty, we believe the evidence points toward important protective effects. We therefore recommend that researchers continue to build the evidence base on this question and that policymakers begin to act on the evidence available today, expanding public housing to better serve the needs of American families.

Appendix

A Sensitivity to missing data on the outcome variable

The identification assumptions in Fig. 2 ignore the fact that the outcome variable is not observed for some cases. For this discussion, we will use eviction as the outcome, but the same discussion applies when nonpayment of rent or mortgage in the outcome. Estimating among observed cases only may induce bias (Fig. 5) if housing assistance and eviction both cause missing responses. The vast majority of missing cases are missing because the child's primary caregiver changed between the age 9 and age 15 interviews, and only the caregiving parent was interviewed at age 15.

It is possible that receipt of housing assistance increases the probability that a parent remains the caregiver at age 15, perhaps by enabling residential stability for this parent. Eviction (the outcome) may also cause the caregiver to change, perhaps because the child stops living with the evicted parent while that parent finds a new apartment. Because missingness is caused by the outcome value and also directly caused by the treatment, it is a collider and conditioning on it can induce a non-causal association between the treatment and the outcome (Pearl, 2009; Elwert and Winship, 2014). To address this concern, we construct the situation most detrimental to our claims: among those not receiving assistance, missingness is not associated with the outcome conditional on pre-treatment variables, but among those receiving assistance type d the probability of the outcome is $\tilde{\eta}_d$ percentage points higher if the outcome is missing. Figure 6 plots the true effect as a function of the sensitivity parameter $\tilde{\eta}_d$.

Discussing effects on nonpayment, a missing outcome would have to be associated with at least a 16 percentage point increase in the probability of nonpayment in order to nullify the finding of a protective effect of public housing (Panel A) and a 50 percentage point increase to nullify the finding of a protective effect of other assistance (Panel B). Given that less than 30 % of each treatment group reports nonpayment, these higher rates of nonpayment among those with missing responses would be implausibly large.

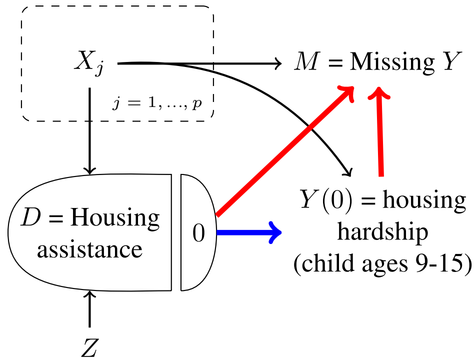

Fig. 5. SWIG including concerns about missing values of the outcome. If missing responses are a consequence of eviction Y and are also affected by housing assistance D directly, then M is a collider and models estimated on non-missing observations do not identify the effect of D on Y.

For effects on eviction, a missing outcome would have to be associated with a 21 percentage point increase in the probability of eviction to nullify the evidence that public housing is protective (Panel C). The association would have to be 29 percentage points to nullify evidence on the effect of other assistance. Given that fewer than 10 % of each treatment group reports eviction, these increases in the probability of eviction among those whose outcome is missing are also implausibly large. Because non-ignorable missingness would have to be extreme to undermine our claims, we conclude that our evidence is reasonably robust to violations of this assumption.

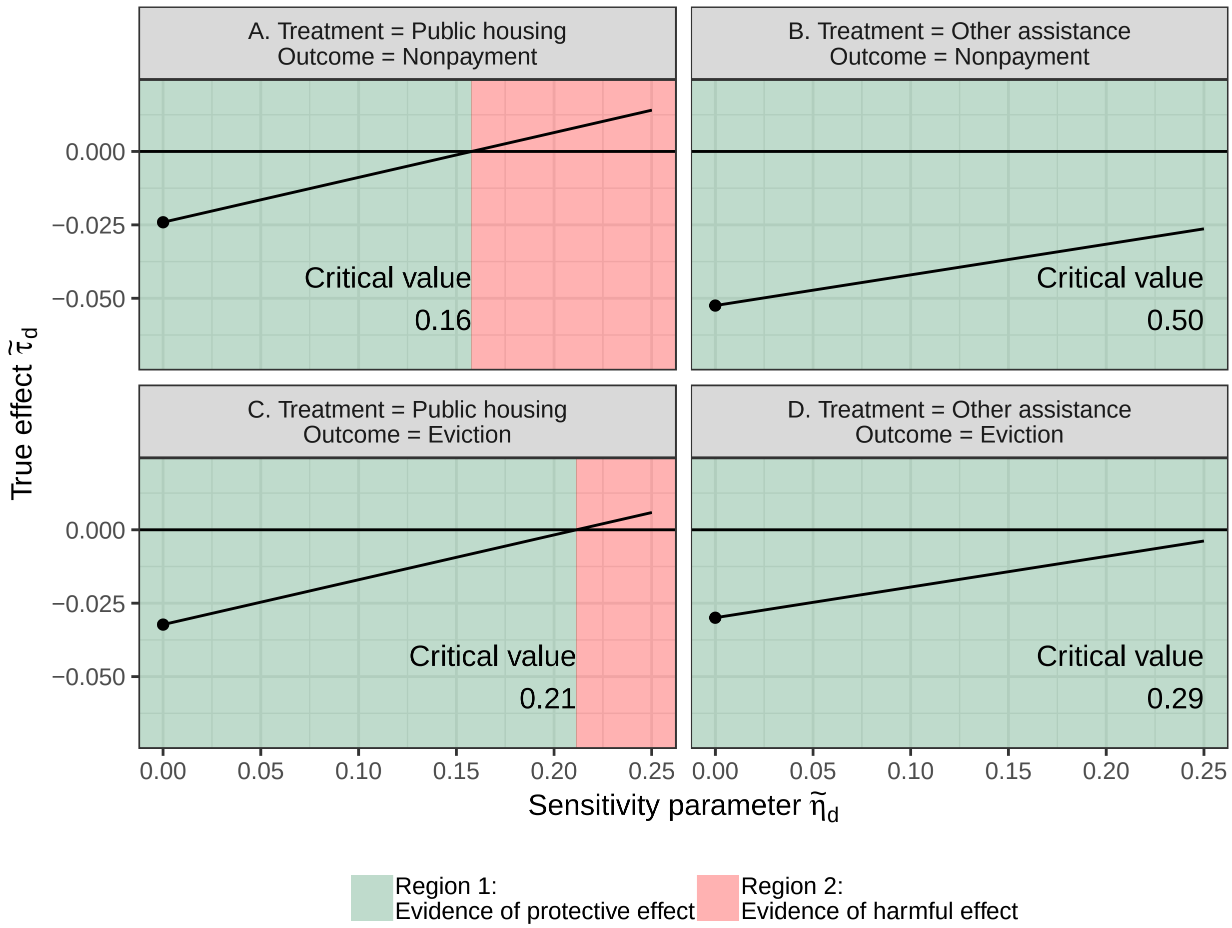

Fig. 6. Causal effect $\tilde{\tau}_d$ as a function of the unknown missingness parameter $\tilde{\eta}_d$. The x-axis represents the extent to missing observations have higher probabilities of eviction than non-missing observations, among those receiving assistance. I assume missingness is ignorable given pre-treatment variables among those not receiving assistance. Dashed lines represent a 95% confidence interval for the true effect at each sensitivity parameter.

B OLS estimates of conditional expectations

Because our aim is to estimate the effect of housing assistance on housing hardship, the OLS regression coefficients by which we estimate the conditional expectation functions are ancillary parameters with no clear interpretation. We therefore omit these estimates from the main text. For interested readers, Tables 2–3 present these estimates. The coefficients represent the association between pre-treatment variables and the outcome, within groups defined by the treatment variable at age 9.

	A. No help	B. Public housing	C. Other assistance
Evicted in past 12 months	0.11	-0.12	-0.08
	(0.06)	(0.22)	(0.14)
Evicted at age 1, 3, or 5	0.20***	0.19	0.06
	(0.05)	(0.13)	(0.09)
Nonpayment in past 12 months	0.22***	0.20*	0.33***
	(0.03)	(0.08)	(0.06)
Nonpayment at age 1, 3, or 5	0.12***	0.13*	0.04
	(0.03)	(0.06)	(0.06)
Income / poverty threshold in past 12 months	0.01	-0.01	-0.02
	(0.01)	(0.04)	(0.03)
Income / poverty threshold averaged over ages 1, 3, and 5	-0.05**	-0.03	0.02
	(0.02)	(0.05)	(0.05)
Disability	-0.04	-0.06	-0.02
	(0.04)	(0.08)	(0.06)
Conviction	0.05	0.06	0.06
	(0.05)	(0.11)	(0.13)
Education: High school	0.03	0.01	0.05
	(0.03)	(0.06)	(0.06)
Education: Some college	0.03	0.04	-0.04
	(0.03)	(0.08)	(0.07)
Education: College	0.13	0.19	0.29
	(0.08)	(0.25)	(0.25)
Parents married at birth	-0.02	-0.03	-0.08
	(0.04)	(0.10)	(0.12)
Race: Hispanic	-0.09**	-0.05	-0.03
	(0.03)	(0.06)	(0.07)
Race: White/other (black omitted)	0.06	-0.12	0.08
	(0.03)	(0.10)	(0.08)
WAIS-R cognitive score	0.01*	0.02	0.01
	(0.01)	(0.01)	(0.01)
Impulsivity (Dickman 1990)	0.01	0.02	0.00
	(0.02)	(0.04)	(0.04)
Intercept	0.13	0.04	0.08
	(0.07)	(0.15)	(0.14)
N	1263	272	300

Table 2. OLS coefficients for models of nonpayment of rent or mortgage as a function of pre-treatment variables within each treatment group. * $p < .05$, ** $p < .01$, *** $p < .001$.

	A. No help	B. Public housing	C. Other assistance
Evicted in past 12 months	0.23***	0.17	0.13
	(0.04)	(0.11)	(0.09)
Evicted at age 1, 3, or 5	0.09**	0.28***	-0.02
	(0.03)	(0.06)	(0.06)
Nonpayment in past 12 months	0.09***	0.06	0.08*
	(0.02)	(0.04)	(0.04)
Nonpayment at age 1, 3, or 5	0.04*	0.04	-0.01
	(0.02)	(0.03)	(0.03)
Income / poverty threshold in past 12 months	0.00	0.01	-0.02
	(0.01)	(0.02)	(0.02)
Income / poverty threshold averaged over ages 1, 3, and 5	-0.02	0.01	0.00
	(0.01)	(0.02)	(0.03)
Disability	0.02	-0.02	-0.04
	(0.02)	(0.04)	(0.04)
Conviction	0.09**	0.06	0.11
	(0.03)	(0.05)	(0.08)
Education: High school	0.01	0.00	0.00
	(0.02)	(0.03)	(0.03)
Education: Some college	-0.03	-0.03	0.04
	(0.02)	(0.04)	(0.04)
Education: College	0.05	-0.04	-0.10
	(0.05)	(0.13)	(0.15)
Parents married at birth	0.00	0.01	-0.06
	(0.02)	(0.05)	(0.07)
Race: Hispanic	-0.07***	-0.04	-0.05
	(0.02)	(0.03)	(0.04)
Race: White/other (black omitted)	-0.02	-0.07	-0.01
	(0.02)	(0.05)	(0.05)
WAIS-R cognitive score	0.00	0.00	0.00
	(0.00)	(0.01)	(0.01)
Impulsivity (Dickman 1990)	0.01	0.02	0.00
	(0.01)	(0.02)	(0.03)
Intercept	0.03	-0.07	0.08
	(0.05)	(0.07)	(0.08)
N	1264	272	300

Table 3. OLS coefficients for models of eviction as a function of pre-treatment variables within each treatment group. * $p < .05$, ** $p < .01$, *** $p < .001$.

C Weighted estimates

Given an estimate of the conditional average treatment association function $\hat{\theta}_d(\vec{x})$, our estimates in the main text marginalize these effects over the $\vec{X}$ distribution observed among treated families in the sample. Because we do not observe the full population distribution of $\{\vec{X}_i\}_{i:D_i=1}$, inferences about the population are more difficult. To provide population estimates in this section, we instead (1) fit the models on the subsample of 1,507 families drawn probabilistically from the sampling frame and (2) marginalize over an estimate of the distribution of $\vec{X}$ from survey weights w_i provided by the Fragile Families Study for mothers interviewed when children were age 9. This serves as an estimator for the average effect of assistance type d, compared with no help, on mothers who received this type of assistance and who gave birth in 1998-2000 in a U.S. city with population over 200,000 who lived with their child when the child was 9 years old.

$$\hat{\hat{\theta}}_d = \frac{\sum_{i:D_i=d,S_i=1} w_i \hat{\theta}_d\left(\vec{x}_i\right)}{\sum_{i:D_i=d,S_i=1} w_i} \tag{7}$$

To estimate the uncertainty of our estimator, we (1) rely on the asymptotic normality of OLS coefficients to capture uncertainty about the conditional average treatment effects and (2) use 26 sets of replicate weights from the Fragile Families Study to capture uncertainty about the population distribution of $\vec{X}$ among the treated. For each set of replicate weights (2), we simulate our uncertainty about (1) with 400 draws of the OLS coefficients from their sampling distribution. This results in 26 $\times$ 400 = 10,400 simulated draws capturing sampling uncertainty about $\hat{\hat{\theta}}$. We report the 2.5 % and 97.5 % quantiles of this distribution to provide 95 % confidence intervals.

Fig. 7 compares the unweighted estimates from the main text, which apply to the sample, and the weighted estimates which are designed to yield inferences about the population. As expected, the weighted estimates are less precise because they involve uncertainty about the distribution of $\vec{X}$ in the target population. Nonetheless, statistical evidence against the null hypothesis that housing assistance does not protect against housing hardship remains, with the joint F-test rejecting the null for effects on nonpayment (p = .036) and eviction ($p = .021$). Point estimates from

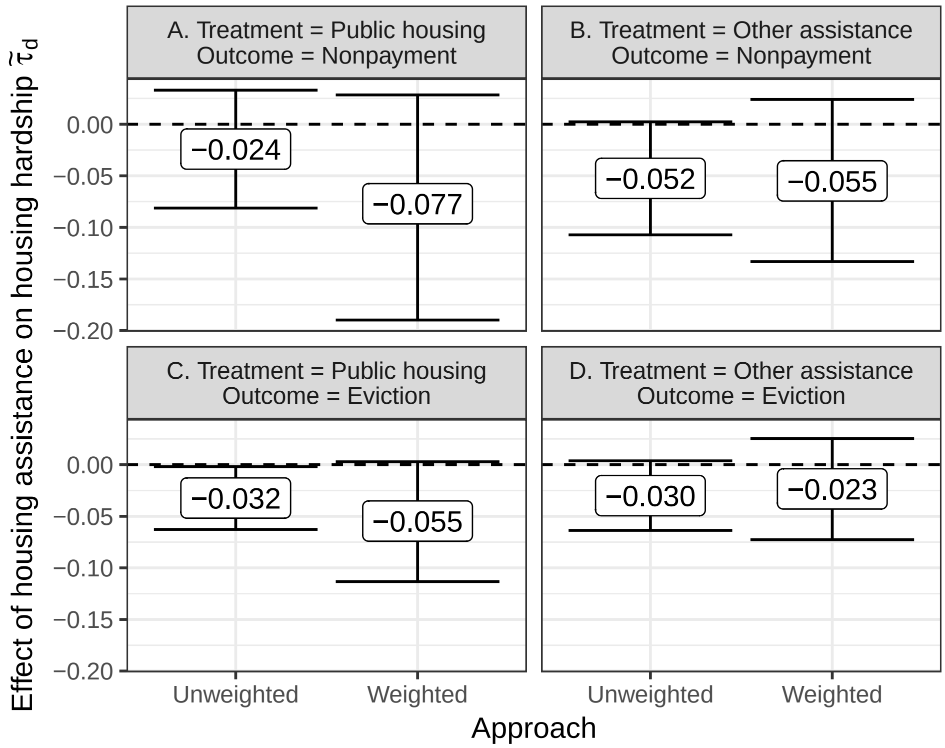

Fig. 7. Comparison of unweighted and weighted estimates of the effect of housing assistance on housing hardship.

both approaches are also similar to the main text and point toward a protective effect of housing assistance on housing hardship, both in the sample and in the population. If anything, the population estimates suggest effects that are slightly more protective (in three out of four cases) than those estimated for the sample.

D Flexible machine learning estimation by causal forests

The main text uses OLS for estimation, but any machine learning approach could be used to learn the unknown functions $\mathbf{E}(Y \mid D = d, \vec{X} = \vec{x})$. An appeal of machine learning is that one need not make parametric assumptions; rather than assuming that the $\vec{X}$ variables are associated with Y in a linear, additive functional form, flexible nonparametric approaches can learn the appropriate functional form from the data. Nonparametric methods can therefore serve as an omnibus robustness check on errors of model specification.

We assess the robustness of results to one specific machine learning approach: causal forests (Wager and Athey, 2017). Our discussion of this procedure is brief and we point readers toward the original articles for a fuller discussion. Causal forests repeatedly grow causal trees to identify subsets of $\vec{X}$ for which the treatment effect $\tau_d(\vec{x})$ is both (a) reasonably homogenous and (b) can be estimated with reasonable precision. At the core of the procedure are causal trees (Athey and Imbens, 2016). Trees proceed by first finding a value along some variable such that the sample can be split into two parts (branches), where the treatment effect estimate in each part is better than the unadjusted difference in the full sample. The tree continues splitting until the data are partitioned into many "leaves" in which treatment effects can be estimated. This partition is learned in a training sample of observations. A separate estimation sample is used to learn the treatment effect associated with each leaf. The entire procedure is designed to optimize accuracy in a held-out test sample. Because individual trees are unstable, causal forests (Wager and Athey, 2017) grow many causal trees on bootstrapped versions of the full data, akin to the procedure of random forests (Breiman, 2001). Forests generally have good predictive performance in many settings, and the flexibility of the splitting approach enables them to approximate any functional form. We use the `causal_forest` implementation in the R package `grf`, which (1) automatically tunes hyperparameters by cross-validation and (2) incorporates a de-biasing procedure to reduce bias from regularization (Athey et al., 2018).

Figure 8 presents estimates from the random forest approach, alongside comparable estimates from the OLS approach. The two approaches yield nearly identical results. We focus on

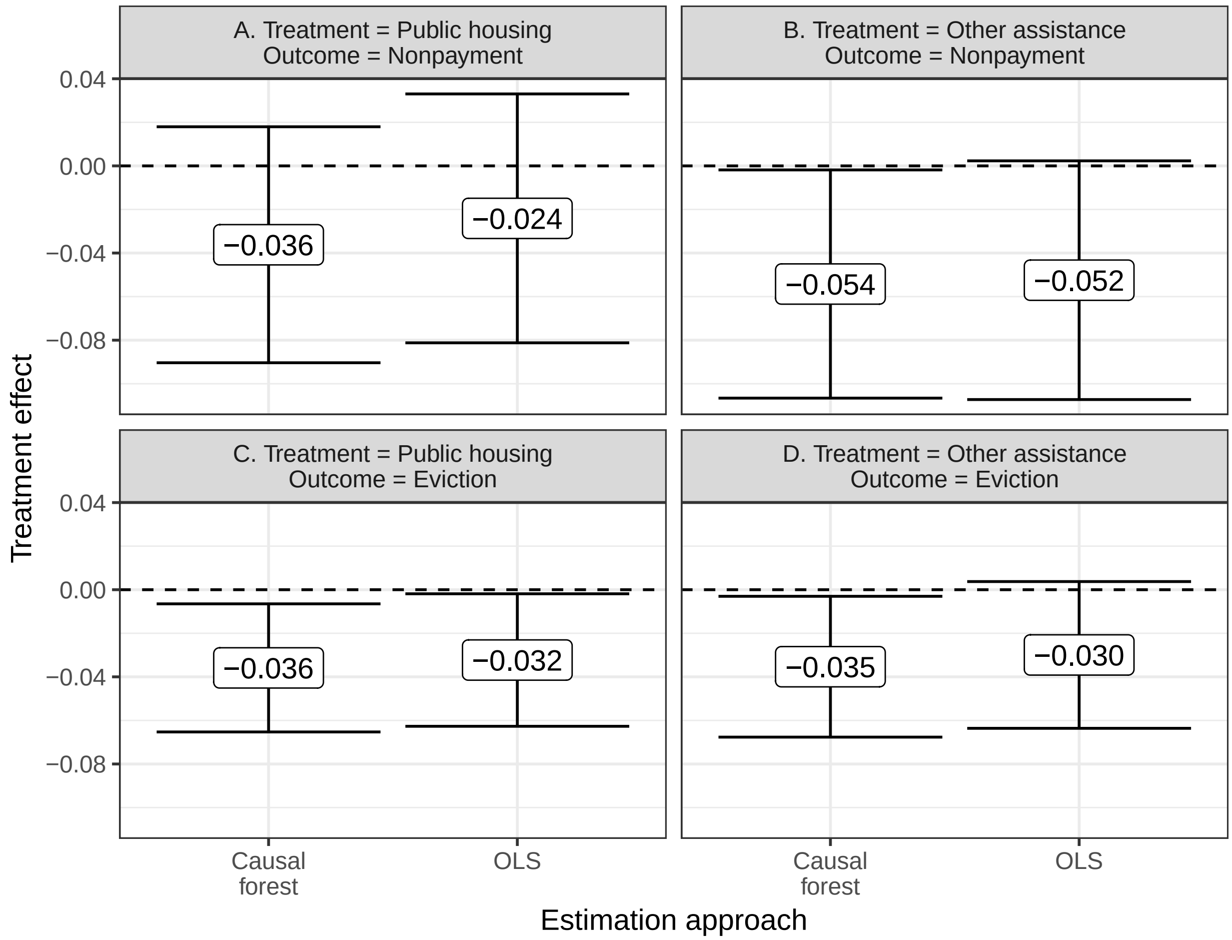

Fig. 8. Estimates are robust to the use of causal forests, a nonparametric machine learning estimation approach.

the OLS results in the main text because they are likely to be more familiar to our social science audience, but the consistency of the results with a more flexible estimation approach demonstrates that claims are not dependent on the particular functional form specification of the OLS model.

E Alternative outcome: Amount of rent payments

It is possible that receipt of public housing or other assistance reduces the risk of nonpayment and of eviction by reducing the amount of money that families must devote to rent. Our data are not optimal to answer this question because rental costs are only recorded at age 9 for respondents who moved between the age 5 and age 9 interviews and are skipped for some families (e.g. those living in a shelter). Whether one moves in this period may be a consequence of receiving assistance, so subsetting to the sample that moves may induce difficult selection problems. Further, we cannot adjust for a lagged outcome because rental costs at age 5 are also recorded only for the subsample that moved in the period preceding that interview. Rental costs at age 15 are recorded for all responding families, but unlike the nonpayment and eviction questions do not cover a cumulative report back to the age 9 interview; this makes rental costs at age 15 a more distant outcome of less interest when the treatment is defined at age 9. We therefore report estimates for rental costs at age 9 only in this section of the appendix, and we warn that estimates should be interpreted cautiously. This exercise is primarily a sanity check that those who receive assistance at age 9 indeed report lower rental costs.

Among the 2,219 families in the analytic sample, we restrict to those who moved between the age 9 and 15 interviews (1,533), who rent their own apartment or house or live with friends or family but contribute to the rent (1,413), and whose monthly rent payment at child age 9 is nonmissing (1,387). This subsample includes 170 families in public housing, 226 families receiving other assistance, and 991 families receiving no help. We top-code monthly rents at $1,500.

Descriptively, those residing in public housing or receiving other assistance report lower mean monthly rental payments ($409 and $486, respectively) compared with those receiving no help ($733). Following the same procedure as for the primary estimates, we estimate that net of other covariates residence in public housing is associated with rental costs that are $250 lower (95% CI: –296, –203) than similar families receiving no help. Likewise, those receiving other assistance have rent payments that are $132 (95% CI: –184, –80) lower than similar families receiving no help. These results agree with our expectation that those receiving assistance have lower rental costs. Because rental costs are only available for a subgroup of the overall sample, though, we hesitate to draw firm conclusions about whether reduced costs mediate the effect of assistance on nonpayment and eviction.

References

Abadie, A. and G. W. Imbens 2006. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267.

Abadie, A. and G. W. Imbens 2011. Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics*, 29(1):1–11.

Angrist, J. D. and A. B. Krueger 1999. Empirical strategies in labor economics. In *Handbook of Labor Economics*, volume 3, Pp. 1277–1366. Elsevier.

Athey, S. and G. Imbens 2016. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360.

Athey, S., J. Tibshirani, and S. Wager 2018. Generalized random forests. *arXiv preprint arXiv:1610.01271*.

Author redacted for peer review 2019. A research note on the prevalence of housing eviction among children born in U.S. cities. *Demography*, Forthcoming.

Breiman, L. 2001. Random forests. *Machine learning*, 45(1):5–32.

Chetty, R., N. Hendren, and L. F. Katz 2016. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902.

Code of Federal Regulations 1976. 24 CFR Part 247: Evictions from certain subsidized and HUD-owned projects. `https://www.law.cornell.edu/cfr/text/24/part-247`.

Code of Federal Regulations 1994. 24 CFR Part 982: Section 8 tenant-based assistance: Housing Choice Voucher Program. `https://www.law.cornell.edu/cfr/text/24/part-982`.

Desmond, M. 2012. Eviction and the reproduction of urban poverty. *American Journal of Sociology*, 118(1):88–133.

Desmond, M., W. An, R. Winkler, and T. Ferriss 2013. Evicting children. *Social Forces*, 92(1):303–327.

Desmond, M. and C. Gershenson 2017. Who gets evicted? Assessing individual, neighborhood, and network factors. *Social Science Research*, 62:362–377.

Desmond, M., C. Gershenson, and B. Kiviat 2015. Forced relocation and residential instability among urban renters. *Social Service Review*, 89(2):227–262.

Desmond, M., A. Gromis, L. Edmonds, J. Hendrickson, K. Krywokulski, L. Leung, and A. Porton 2018. Eviction lab national database: Version 1.0. `www.evictionlab.org`. Accessed: 16 September 2018.

Desmond, M. and R. T. Kimbro 2015. Eviction's fallout: housing, hardship, and health. *Social Forces*, 94(1):295–324.

Desmond, M. and T. Shollenberger 2015. Forced displacement from rental housing: Prevalence and neighborhood consequences. *Demography*, 52(5):1751–1772.

Dickman, S. J. 1990. Functional and dysfunctional impulsivity: Personality and cognitive correlates. *Journal of Personality and Social Psychology*, 58(1):95.

Elwert, F. and C. Winship 2014. Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology*, 40:31–53.

Fragile Families and Child Wellbeing Study 2006. *Scales Documentation and Question Sources for Three-Year Questionnaires*. Center for Research on Child Wellbeing, Princeton University.

Geller, A. and M. A. Curtis 2011. A sort of homecoming: Incarceration and the housing security of urban men. *Social Science Research*, 40(4):1196–1213.

Geller, A. and A. W. Franklin 2014. Paternal incarceration and the housing security of urban mothers. *Journal of Marriage and Family*, 76(2):411–427.

Greenland, S., J. Pearl, and J. M. Robins 1999. Causal diagrams for epidemiologic research. *Epidemiology*, Pp. 37–48.

Hahn, J. 1998. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, Pp. 315–331.

Hernan, M. A. and J. M. Robins 2018. *Causal Inference*. Unpublished book manuscript.

Honaker, J., G. King, M. Blackwell, et al. 2011. Amelia II: A program for missing data. *Journal of Statistical Software*, 45(7):1–47.

Imbens, b. G. W. and D. B. Rubin 2015. *Causal Inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.

Jacob, B. A., M. Kapustin, and J. Ludwig 2014. The impact of housing assistance on child outcomes: Evidence from a randomized housing lottery. *The Quarterly Journal of Economics*, 130(1):465–506.

Joint Center for Housing Studies 2011. Rental market stresses: Impacts of the great recession on affordability and multifamily lending. `https://www.urban.org/research/publication/rental-market-stresses-impacts-great-recession-affordability-and-multifamily-lending`. Accessed 27 September 2018.

Joint Center for Housing Studies 2018. The state of the nation's housing 2018. `http://www.jchs.harvard.edu/state-nations-housing-2018`. Accessed 18 September 2018.

Kim, H., S. A. Burgard, and K. S. Seefeldt 2017. Housing assistance and housing insecurity: A study of renters in southeastern michigan in the wake of the great recession. *Social Service Review*, 91(1):41–70.

McCarty, M. and A. Siskin 2012. Immigration: Noncitizen eligibility for needs-based housing programs. *Congressional Research Service Report RL31753*.

Mills, G., D. Gubits, L. Orr, D. Long, J. Feins, B. Kaul, M. Wood, and A. Jones 2006. Effects of housing vouchers on welfare families. `https://www.huduser.gov/portal/publications/commdevl/hsgvouchers.html`. Accessed 27 September 2018.

Moore, M. K. 2016. Lists and lotteries: rationing in the housing choice voucher program. *Housing Policy Debate*, 26(3):474–487.

Neyman, J. 1923. Translation published 1990. On the application of probability theory to agricultural experiments. Essay on principles. Section 9. *Statistical Science*, 5:465–472.

Park, J. M., A. Fertig, and S. Metraux 2014. Factors contributing to the receipt of housing assistance by low-income families with children in twenty American cities. *Social Service Review*, 88(1):166–193.

Pearl, J. 2009. *Causality*. Cambridge University Press.

Public and Affordable Housing Research Corporation 2016. Housing agency waiting lists and the demand for housing assistance. https://www.housingcenter.com/wp-content/uploads/2017/11/5k1S1v-waiting-list-spotlight.pdf. Accessed: 17 September 2018.

Reichman, N. E., J. O. Teitler, I. Garfinkel, and S. S. McLanahan 2001. Fragile Families: Sample and design. *Children and Youth Services Review*, 23(4-5):303–326.

Richardson, T. S. and J. M. Robins 2013. Single world intervention graphs (SWIGs): A unification of the counterfactual and graphical approaches to causality. *Center for the Statistics and the Social Sciences, University of Washington Series. Working Paper*, 128(30):2013.

Rubin, D. B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688.

Sanbonmatsu, L., L. F. Katz, J. Ludwig, L. A. Gennetian, G. J. Duncan, R. C. Kessler, E. K. Adam, T. McDade, and S. T. Lindau 2011. *Moving to Opportunity for Fair Housing Demonstration Program: Final Impacts Evaluation*. US Department of Housing and Urban Development.

Sanbonmatsu, L., J. R. Kling, G. J. Duncan, and J. Brooks-Gunn 2006. Neighborhoods and academic achievement results from the moving to opportunity experiment. *Journal of Human Resources*, 41(4):649–691.

U.S. Department of Housing and Urban Development 2001. Housing choice voucher guidebook. http://www.lb7.uscourts.gov/documents/15c6523.pdf. Accessed: 27 September 2018.

U.S. Department of Housing and Urban Development 2018. HUD's public housing program. https://www.hud.gov/topics/rental_assistance/phprog. Accessed: 19 August 2018.

Wager, S. and S. Athey 2017. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*.

Wechsler, D. 1981. *Wechsler Adult Intelligence Scale — Revised (WAIS-R Manual)*. The Psychological Corporation. Harcourt Brace Jovanovich.

Wildeman, C. 2014. Parental incarceration, child homelessness, and the invisible consequences of mass imprisonment. *The ANNALS of the American Academy of Political and Social Science*, 651(1):74–96.