The Insurance Value of Federal Rental Assistance

Nick Grasley*

November 13, 2025

Link to latest version

Abstract

Each year, the U.S. federal government spends \$48 billion providing large rental subsidies to a relatively small share of low-income households, many of whom wait years to receive assistance. Unlike other safety net programs, rental assistance uniquely subsidizes households when area rents rise or when household income falls, smoothing the share of income spent on rent. Motivated by this design and the high fixed costs of adjusting housing consumption, I conceptualize federal rental assistance as insurance against joint rent-income risk. To analyze this risk, I use the Panel Survey of Income Dynamics (PSID) and rich HUD administrative data, showing that rents exacerbate consumption risk beyond income risk alone. I quantify how rental assistance insures this particular consumption risk using a sufficient statistics framework, finding that rental assistance generates \$1.51 in benefits per dollar of government spending, significantly larger than the previous literature's estimate of \$0.66. I attribute this large insurance value to insuring rent volatility for fixed-income households and insuring income in high-rent locations. Using a structural lifecycle model, I conduct counterfactual analysis that demonstrates that the program's dynamic subsidy design and rationing through waitlists are crucial features that enhance its insurance value.

*Stanford University, (email: ngrasley@stanford.edu). I thank Neale Mahoney, Luigi Pistaferri, and Caroline Hoxby for their mentorship. I thank Matthew Brown, Dante Domenella, Mark Duggan, Mariana Guido, Brendan Moore, Petra Persson, Charlie Rafkin, Stephen Redding, Isaac Sorkin, Alessandra Voena, and seminar participants at Stanford University for valuable comments and suggestions. I thank Rae Winegardner and other staff at HUD for their incredible support in obtaining the HUD data. I thank the Stanford Institute for Research in the Social Sciences for graciously providing the office space to conduct this research. I gratefully acknowledge support from the Bradley Graduate and Post Graduate Fellowship of the Stanford Institute for Economic Policy Research. The views expressed here are those of the author and should not be construed as representing those of the US Department of Housing and Urban Development

1 Introduction

Government rental assistance programs aim to subsidize excessive rent costs of households in poverty, with a quarter of households in poverty spending over three quarters of their income on rent (Desmond, 2015). While some worry about these excessive rent costs, others worry that large rent subsidies, averaging \$8,000 per year per recipient, are an inefficient means of reducing poverty. They argue for flexible cash transfers that households would prefer over restrictive rent subsidies.

We lack a rationale for why these large rent subsidies are welfare improving over cash welfare programs. The prevailing view is that rent subsidies may incentivize households to move to better neighborhoods where their children will have improved outcomes (Chetty, Hendren, and Katz, 2016). While true for those who move, most households in these programs either do not move to better neighborhoods or do not have children to benefit from the move (Carlson et al., 2012; Bergman et al., 2024). Even special housing vouchers that condition aid on moving to better neighborhoods reduce voucher use by a third (Chetty, Hendren, and Katz, 2016). Without migration, we lack an economic argument for why we should provide transfers through rental assistance, and therefore the current argument finds a marginal expansion in rent subsidies is only worth \$0.66 per dollar of government spending (Hendren and Sprung-Keyser, 2020).

In this paper, I rationalize rental assistance subsidies as valuable insurance against risk that cash welfare fails to address and quantify its insurance value. This insurance motive arises because rent prices pose a large risk to the consumption of low income households. These households consume large portions of their budget on housing and, unlike other goods, face high adjustment costs to substitute away from high rent prices. I show that rent subsidies directly insure households against the extra consumption risk from rent prices by smoothing the share of their income they pay as rent rather than solely smoothing income. I quantify that the insurance benefits alone—without appealing to neighborhood benefits—imply that marginally expanding federal rental assistance is worth the costs.

Federal rental assistance insures both rent and income risk through its rent subsidy design. The rent subsidy is structured as an income share agreement, where households generally pay 30% of their income as rent. Since households face no market rent costs, this intuitively insures rent volatility. I show how marginal utility directly relates to volatility in rent prices, and households value insurance against rent volatility if they inelastically spend large shares of their budget on housing.

Federal rental assistance can even provide more valuable *income* insurance than traditional cash welfare because it conditions subsidies on rent prices. The key insight is that rent prices constrain budgets more in low income states than high income states. When income falls, housing becomes a larger and inscreasingly inelastic share of household consumption. This leads to rent prices becoming an important constraint on

total household consumption when low income. Therefore, those in high rent locations will value income insurance more to protect them against the risks of exorbitant rent costs. Rental assistance directly insures income more in high rent locations by fixing the share of household budget spent on rent. Cash welfare misses this important heterogeneity in income insurance demand.

Insuring joint rent-income risk is further targeted by rationing rental assistance through a wait list, in contrast to entitlement cash welfare programs. Because of fixed funding, households must often wait years to receive assistance (Acosta and Gartland, 2021). Years-long waits select households whose shocks are persistent enough to last until they can receive the subsidies. This targeting is beneficial if households demand more insurance for persistent shocks over temporary ones.

While federal rental assistance provides targeted insurance against joint rent-income risk, this leads to moral hazard in *both* labor supply and housing consumption. Because households do not pay market rent, they may overconsume housing within the limits of the rental assistance program. In addition, since they swap paying rent for an income share agreement, they also are incentivized to reduce their labor supply, leading to fiscal externalities on both rent subsidies and government tax revenues (Jacob and Ludwig, 2012). These tradeoffs must be evaluated against the insurance value generated by the program.

To estimate the net insurance value of the program, I combine the Panel Survey of Income Dynamics (PSID) with rich HUD administrative data on recipient households from 2003-2022. The HUD administrative data contains quarterly information on household demographics, income, and benefits while in the program. For calculating program benefits, income is certified by local housing authorities for accuracy and broken down by where the income was received (e.g., work, social security, pensions, etc.). HUD also tracks how long recipient households had to wait to receive assistance, the amount of rent subsidies they receive, and basic details on the quality and location of the housing. These administrative data play a pivotal role in identifying and estimating how much households receive in benefits and their behavior while in the program.

Using these data, I first estimate the net insurance value of the program using the sufficient statistics framework (Chetty and Finkelstein, 2013). Using elasticity estimates from the literature and new evidence on consumption and benefits from the PSID and HUD, I find that a marginal expansion of the program generates \$1.51 in benefits per dollar of government spending, relative to the literature's estimate of \$0.66 (Hendren and Sprung-Keyser, 2020). This greatly changes the implications for whether the program should expand or contract.

I decompose this estimate into the value of insuring rent prices and insuring income shocks conditional on rent. I find that the average household solely values the program for its income insurance component. This income insurance value scales with local rent prices, suggesting that the program better targets con-

sumption risk from income shocks than income-based cash welfare programs. While the average house-hold values the income insurance component, this masks heterogeneity in the insurance value among other household types. I find that fixed-income elderly households do value rent price volatility insurance, as their pension incomes are uncorrelated with local labor market dynamics that change rent prices.

To improve identification of the key parameters governing the insurance value and understand the mechanisms behind it, I estimate a lifecycle model where households choose consumption, housing, and labor supply in the presence of joint rent-income risk. Because the choice of rental assistance is inherently dynamic, the lifecycle model provides structure to the interactive dynamics of wages, rent prices, and rental assistance wait lists. When eligible, households may choose to enter rental assistance after a waiting period. In the model, federal rental assistance reduces rent burdens and thereby smooths consumption between high and low rent burden states. However, the incentives of the program encourage households to reduce labor supply and distorts the choice between rent subsidies and private market housing.

I rely on the PSID and HUD data to identify the key parameters of the model. Both datasets contribute to understanding the joint rent-income risk that households face. The HUD data especially aids in accurately identifying wait times, rent subsidies, and household dynamics while in rental assistance. I then estimate the lifecycle model using indirect inference, connecting the model's parameters to HUD and PSID data using an auxiliary model.

After estimating the model, I examine proposed reforms that change the key insurance features of the program. Several proposals advocate for changing to a flat rent subsidy or allocating rental assistance funds to cash welfare. First, I find that changing from an income-based rent subsidy to a flat rent subsidy decreases the welfare without any labor supply benefits. By shifting benefits away from low income households, a flat rent decreases the insurance value of the program. In addition, the flat rent does not improve labor supply because it changes the selection of households who participate in the program.

Changing rental assistance to cash welfare can modestly improve welfare, and keeping the core screening features of rental assistance helps control costs. The most important screening mechanism is the wait-lists, which limits the receipt of cash welfare to those who are persistently in poverty. In addition, indexing the cash transfers to both income and local rent prices improves targeting of benefits over income-based cash welfare.

My paper contributes to the social insurance literature by reconceptualizing rental assistance as a unique form of insurance. There are a growing number of papers evaluating the net insurance value of common social insurance programs, including unemployment insurance (Chetty, 2006, 2008; Ganong and Noel, 2019; Landais and Spinnewijn, 2021), disability insurance (Low and Pistaferri, 2015; Autor et al., 2019; Deshpande and Lockwood, 2022), health-related insurance (Lieber and Lockwood, 2019; Finkelstein, Hen-

dren, and Luttmer, 2019; Lockwood, 2025), and the insurance value of take-up/targeting (Rafkin, Solomon, and Soltas, 2025; Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019). My paper formalizes how rent prices lead to uninsured consumption risk and how rental assistance insures that risk. This closely complements recent results from Gadenne et al. (2021) that show how households value in-kind transfers in the presence of commodity price risk, focusing on the development context. In their case, households value insurance against volatile commodity prices. In my context, households value insurance against volatile income that is exacerbated by commodity prices, in this case rent prices.

My paper also relates to a literature evaluating the benefits and costs of federal rental assistance programs. Papers either focus on the moral hazard costs of the program Jacob and Ludwig (2012); Collinson and Ganong (2018) or the locational effects of the program (Chetty, Hendren, and Katz, 2016; Cook, Li, and Binder, 2025; Chyn, 2018). A related paper to mine examines how the nuances of the wait list design for specific properties impact targeting (Waldinger, 2021). My paper conceptually differs by focusing on measuring the consumption risks that cause households to value a general rental assistance program.

Finally, my paper contributes to a literature evaluating household consumption risk. These papers develop lifecycle models where households face primarily wage risk (Low, Meghir, and Pistaferri, 2010; Low and Pistaferri, 2015; Low et al., 2022). There are a small number of papers examing the consumption risk of rent prices, both the risk of renting versus owning (Kueng et al., 2023) and the equilibrium effects of a theoretical rent guarantee insurance contract (Abramson and Van Nieuwerburgh, 2024). My paper contributes to this literature by examining how rental assistance acts as one of the primary forms of consumption insurance against rent prices in the United States.

In the rest of the paper, I show how rental assistance insures joint rent-income risk and quantify the insurance value, using a lifecycle model to simulate the counterfactuals. In Sections 2 and 3, I provide background on my data sources and the institutional design of federal rental assistance. With this background, I then develop a conceptual framework on the insurance value of rental assistance in Section 4. In Sections 5 through 8, I develop a lifecycle model of joint wage-rent risk and show how the rental assistance program insures consumption under various counterfactuals.

2 Data

I first describe the main data sources that will appear throughout the analysis of federal rental assistance programs, including HUD administrative data, the PSID, and the Census data.

2.1 HUD Administrative Data

The HUD administrative data provides rich information on households in public housing and voucher programs. This paper uses a 50% quarterly sample of all households who participate in federal rental assistance from 2003-2022. This dataset contains records of recipient household demographics, income, and housing characteristics in the program.

HUD collects these administrative records to monitor program requirements, such as recipient rent payments and the housing quality of recipients. Local housing authorities regularly certify income of households via pay stubs (or other methods) in order to determine accurate rent payments. Because program rules can adjust rent payments differently for different streams of income, HUD collects detailed records on multiple household income streams, includling earned income, asset income, and social safety net cash payments. HUD also monitors many other demographic characteristics that determine eligibility for HUD programs, including family size, age, disability, and homelessness status.

HUD also collects some basic housing quality indicators as an audit of housing units used by recipients. This information includes variables such as location, the type of structure, the number of bedrooms of the apartment, and the age of the building. Generally, this is comparable to housing characteristics found in general US Census data.

These data play an important role in identifying the insurance value of federal rental assistance, and to estimate a model of household behavior in the context of federal rental assistance. These data provide precise information on how rent subsidies vary across households and locations, a key moment in identifying the insurance value. Additionally, model identification relies on many moments from the HUD data that determine when households receive assistance and their revealed preference for that assistance.

For the baseline estimation, I restrict the sample to 2012-2022 unless stated otherwise. After 2011, variable definitions are consistent within the HUD data between years, allowing for a consistent panel.

2.2 Panel Survey of Income Dynamics

To obtain more data on households outside of rental assistance, I complement the HUD administrative data with a restricted sample of the Panel Survey of Income Dynamics (PSID). By its nature, the HUD administrative data only captures a snapshot of households while they are in rental assistance. The PSID zooms out and provides a fuller picture of household dynamics both in and out of rental assistance, at the cost of less information on households in rental assistance. Specifically, the PSID data contains information about household annual tenant payment in rental assistance, but it does not have reports on the specific rent subsidies that households receive nor the total contract rent.

Generally, the PSID relies on self-reports of rental assistance status, which can lead to misreporting of program receipt (Meyer and Mittag, 2019). To overcome this, I access a restricted PSID sample that links households to available administrative records¹ on rental assistance addresses. These addresses represent known rental assistance units, either public housing or other voucher programs. This improves the precision of rental assistance status when using the PSID.

2.3 Census Data

I use Census data, including the Census and American Community Survey (ACS), to measure aggregate wages and rents in locations. The HUD data and the PSID are ill suited for measuring aggregate wages and rents, either because it's a select population (in the case of HUD) or the sample size is small (in the case of the PSID). To address this, I use the Census data to estimate Commuting Zone statistics on wages and rents, which I then merge to each of the previous two datasets.

Commuting Zones (CZs) allow me to define consistent geographic boundaries between Censuses, crucial for measuring wage-rent shocks over time. The Census often limits geography to course Public Use Microdata Areas (PUMAs) that can change between Censuses. CZs allow me to specify a consistent geography with which to measure wages and rents, at the cost of losing some precision on hyper-local wages and rents. To measure CZ wages and rents, I follow Kueng et al. (2023), who use CZ wages and rents for very similar purposes. Their method involves estimating a quality-adjusted, normalized CZ wage/rent using available characteristics in Census data. For more details on this method, see Appendix B.1.

3 Background: Federal Rental Assistance Programs

Federal rental assistance provides in-kind housing as a safety net to ensure access to housing. Because there are a multitude of programs all with different designs, I focus on the two traditional HUD programs, public housing and the voucher program—which I use as a catch-all for any Section 8 or similar program². While the supply-side differs between public housing and vouchers, both programs have similar demand-side rules, rationing subsidies via wait lists and subsidizing household rent based on income.

3.1 What Programs?

To tractably analyze the insurance value of federal rental assistance, I focus on the two traditional rental assistance programs: public housing and voucher programs. These programs began with the Housing Act of 1937, starting with public housing and expanding to the modern voucher program in 1974. These

¹These administrative records are distinct from my HUD administrative records. Please see Newman and Schnare (1997) for a complete description of the collection of addresses to create this dataset.

²There are some smaller voucher programs specifically legislated for specific groups, such as those with disabilities. These programs have very similar designs save for the targeted group receiving it. These programs are also generally small relative to Section 8

Table 1: Number of HUD occupied units by program type

Program	1990 No. Households	2022 No. Households
Public Housing	1,141,068	834,946
Voucher Programs	2,227,612	3,617,292
LIHTC	850,812	2,366,324
All Programs	4,246,540	6,875,768

Notes: Public HUD data on total occupied units by program type, 1990-2022. Voucher programs include Housing Choice Vouchers (Section 8 vouchers), project-based Section 8, Section 202 (elderly vouchers), and Section 811 (disabilities vouchers). LIHTC refers to the Low Income Housing Tax Credit program. Data collected through reports from local housing authorities. Occupied units exclude, at the time of data collection, vacant units and units that are unreported by a local housing authority.

programs make up a large portion of assisted housing in the United States, as shown in Table 1, albeit with a shrinking share of the overall assisted housing stock.

The key difference between public housing and the voucher program is that public housing is owned and operated by the government, while vouchers provide subsidites to lease on the private market. Households may choose to apply to both programs, but each program is managed separately from the other. When a household receives public housing, they cannot choose to convert public housing into a Section 8 voucher and vice versa.

3.1.1 When Do Households Receive Assistance?

Unlike the majority of welfare programs, rental assistance is rationed, even if the household is eligible for assistance. Congress allocates HUD a fixed number of housing units to distribute between local areas. To do so, they apportion subsidies to local public housing authorities, who provide the day-to-day management of the subsidies.

To ration subsidies, receipt follows a two-step process. First, households must first be income-eligible before even applying to receive one of the rationed subsidies. Income eligibility varies by the local housing authority based on the local income distribution. For each housing authority, HUD will estimate an Area Median Income (AMI), which is the upper limit on eligible household income after deductions and household size adjustments. AMI is generally equivalent to a metro's median income measured in the ACS, with a floor of the poverty level. Households must be below 80% of the AMI in order to receive rental assistance. In Table 2, I show that the average 80% AMI cap is relatively high for a safety net program, reaching above \$50,000 per year for the average household.

The AMI itself is often not the binding constraint on eligibility. Statutorily, 75% of a local housing authority's recipients must be below 30% of the AMI, explicitly targeting subsidies to the needlest eligible households. With this design, households are generally admitted when income falls below 30% AMI,

with the exception of households with particular housing distress. After admission, household income can naturally rise above 30% AMI without being removed from the program.

The second step to receive rental assistance is to apply for the waitlist. While a household may be statutorily eligible for rental assistance, this does not guarantee that they will receive housing subsidies. Because each housing authority often has less funding than eligible households, authorities must ration available subsidies through a waitlist. Housing authorities have wide discretion in how they manage waitlists, often creating priorities based on household characteristics, offering lotteries, or even closing the waitlist entirely. In Appendix A.1, I show examples of these priorities and how they are incorporated into waitlists.

Housing authorities also have wide variation in expected wait times conditional on being on the wait-list. In Table 2, I show that while the mean housing authority has a wait time of around 2 years, there is a 2.98 standard deviation in these wait times. This distribution of wait times reflects HUD recommendations to keep the wait below 2 years, but often authorities struggle to manage the excess demand when they do open the waitlists.

3.1.2 How Much Assistance do they Receive?

When households do receive assistance, households pay 30% of their adjusted income as rent, with the rest of the market rent subsidized by the housing authority. To estimate rent, housing authorities use expected annual income³ and verify income by collecting pay stubs from employers. Income adjustments include family size deductions and many household expenses, such as child care, disability expenses, and uninsured medical expenses. In Table 3, I show both annual income and the adjusted income of households in the program.

Rent payments change at the extremes of adjusted income in the program. At the low end of income, Housing authorities may set a minimum rent that households must pay in the program, typically set to \$50 per month, with exceptions for hardship. At the high end of income, households may opt for a flat rent as opposed to 30% of their income, effectively capping their rent payments to the authority. In practice, few households opt for a flat rent because it is much higher than the typical rent formula.

To mitigate moral hazard in voucher rent subsidies, local housing authorities take various actions to limit the rents landlords can charge. They first verify that the charged rent is similar to other apartments in the area. If the charged rent is above other comparable apartments, the housing authority may refuse to allow the voucher to be used for that apartment. In addition, HUD caps rent subsidies at a measure called the Fair Market Rent (FMR), an estimate of the median rent in an area. This further prevents landlords from charging exhorbitant rents, and also effectively limits the quality of apartments that voucher holders can

³to account for expected seasonal variation in work.

Table 2: Voucher program characteristics

Characteristic	Mean	St. Dev.
Annual Rent	3,851.01	2,718.29
Utility Subsidy	117.53	398.13
Annual Subsidy	8,374.64	4,914.07
Wait Time (Years)	2.12	2.98
Fair Market Rent	13,290.84	5,560.15
80% AMI	50,291.04	13,290.84
Structure Age (Years)	48.02	30.64
Number Bedrooms	1.81	0.945
Observations	37,066,888	

Notes: Table of voucher program characteristics, HUD data 2012-2022. Programs include public housing and housing choice vouchers. Annual rent subsidy only applicable for voucher holders since public housing has no private market rent amount. All dollar amounts deflated to 2016 dollars. Variables winsorized at 1st and 99th percentiles to remove outliers. Area Median Income is adjusted for household size.

Table 3: Assisted demographics.

Statistic	Mean	St. Dev.
Head White (%)	50.2	50.0
Head Female (%)	77.7	41.6
Head Disabled (%)	41.2	49.2
Number Dependents	0.980	1.369
Prior Homelessness (%)	5.1	22.0
Ann. Adj. Income	12,914.580	11,066.710
Observations	58,741,803	

Notes: Demographics of households in rental assistance, HUD data 2012-2022. Household head derived from the individual on the lease. Disability is a self-report of disability status in the HUD data. All dollar amounts deflated to 2016 dollars.

rent⁴.

In Table 2, I show how these rules translate into household rent payments and rent subsidies for Section 8 programs⁵. Households pay around 31% of the total rent costs of their housing when in rental assistance, receiving subsidies of around \$8,400 per year. The mean total rent costs are only \$1,065 below the mean Fair Market Rent, suggesting many voucher apartments rent near the limits of the program.

The interaction of these rules generally select either young households with children or older disabled households. In Table 3, I show the demographics of households in rental assistance. Over 40% of households declare that they have some disability⁶. The average adjusted household income is well below the poverty line for a 3-person household.

⁴For more details on the supply-side of vouchers, see Collinson and Ganong (2018)

⁵It is much more challenging to estimate the implicit subsidy for public housing since it is owned and operatedby the government.

⁶Note that this is reported disability and not whether households are receiving SSDI.

4 Conceptual Framework: What Does Federal Rental Assistance Insure?

If rental assistance functions as insurance, it matters whether the value of benefits received from the program is correlated with the marginal utility of consumption. To make the insurance argument for rental assistance precise, I outline a simple conceptual framework of joint wage-rent risk and targeted rental assistance subsidies⁷. In contrast to income-based cash transfers, federal rental assistance targets the share of household income spent on rent, which I define as the rent expenditure share. If the rent expenditure share is correlated with marginal utility, then rental assistance necessarily provides insurance value. Decomposing the insurance value shows that rental assistance insures both rent volatility and also better insures income shocks in high rent locations. I show descriptively that how households may value insurance against both of these risks and quantify that decomposed value.

4.1 The Insurance Value of Federal Rental Assistance

Federal rental assistance has insurance value when it delivers the largest benefits when households need them most—that is, when households have high marginal utility of consumption. I define a simple conceptual framework where households face risky wages and rents, and the social planner can smooth these risks by marginally expanding rental assistance subsidies.

Household i with ex ante risk type θ^8 faces the wage-rent distribution $(w, p^h) \sim F_{\theta}(w, p)$. In the face of this risk, households choose labor supply P, numeraire consumption c, and housing consumption h. Housing consumption includes all amenities and, for exposition, there is no migration θ^9 .

When households receive a rental assistance offer, they may choose to enter rental assistance and pay a fraction τ of their income as rent. Assistance offers arrive randomly according to household characteristics. When they receive and accept an offer (Z=1), they receive subsidy $(p^hh-\tau wP)$, where household rent above share τ of income wP is subsidized. To theoretically analyze the value of these vouchers, I include a

⁷The additional advantage of this simple model is that the results generalize to any government program that subsidizes excess expenditure shares of a commodity. Any commodity with heterogeneous pricing across space and time and a similar subsidy design will have the same theory for the insurance value of that subsidy. Any particular nuances of rental assistance over other consumption category subsidies are addressed in the structural model.

 $^{^8}$ A common risk type is education level, as in Deshpande and Lockwood (2022). This risk type serves as a useful benchmark to literature estimates of insurance value of other government programs. It is normatively debatable what characteristics or risks are part of a household's risk type θ . For example, being born in a high rent location can either be a risk that a household faces or a part of θ . I plan to show robustness to the specific choice of risk type but default to the benchmark of education to make my estimates comparable to the literature.

⁹Modelling migration separately from housing consumption does not qualitatively change the resulting insurance value of federal rental assistance. Choosing housing consumption can be thought of as including migration decisions; when rent prices rise, the household may migrate to a location with worse amenities and thus choose lower housing quality *h*. Therefore, one can interpret any housing elasticity parameters as including migration elasticities. The structural model will incorporate migration separately to account for the quantitative differences between housing quantity choices and location choices. A limitation of both is this framework cannot evaluate targeted marginal expansions without a full structural model.

 β to allow for partial subsidies, or $\beta(p^hh - \tau wP)$. In practice, $\tau = 0.3$ (the 30% implicit income tax of rental assistance) and $\beta = 1$, meaning all excess rent is subsidized¹⁰.

The ex ante marginal benefit of expanding federal rental assistance is the ex post value of the subsidy plus the insurance value of targeting rental assistance to high marginal utility states. To derive this, consider the following ex post household utility:

$$v(w, p^h) = \max_{c,h,P} u(c,h,P)$$

$$\tag{4.1}$$

s.t.
$$c + p^h h \le wP + Y + Z\beta(p^h h - \tau wP)$$
 (4.2)

Similar to Lieber and Lockwood (2019), the ex ante marginal benefit of expanding rent subsidies through β^{11} is

$$MB = \frac{\partial E[v]/\partial \beta}{E[v_y]} = \frac{E[v_y Z(p^h h^* - \tau w P^*)]}{E[v_y]}$$

$$\tag{4.3}$$

$$= \underbrace{E[Z(p^{h}h^{*} - \tau wP^{*})]}_{\text{Transfer Value}} + \underbrace{\text{Cov}(\hat{v}_{y}, Z(p^{h}h^{*} - \tau wP^{*}))}_{\text{Insurance Value}}$$
(4.4)

where \hat{v}_y is the normalized marginal utility of income. The transfer value is how much the subsidy on average loosens the household budget constraint ex post, while the insurance value is how much the subsidy covaries with high marginal utility states and thus smooths utility. If household marginal utility is high when rent expenditures greatly exceed a fraction τ of income, then households will have high insurance value from expanding rental assistance.

This insurance value expression¹² is a black box about how insuring rent expenditure shares is valuable relative to traditional cash welfare. In comparison, traditional cash welfare only insures income drops, ignoring any targeting that comes from rent prices. What matters is if the marginal utility of income is significantly correlated with rent prices, or if income can account for most of the correlation alone¹³.

 $^{^{10}}$ In addition, there is a cap on p^hh in rental assistance. I omit the cap for expositional purposes.

 $^{^{11}}$ A more intuitive policy is marginally expanding the number of available rent subsidies, or in other words the probability of receiving offer R in state (w, p^h) . This does not change the implications. Marginally expanding the benefit through β is equivalent to marginally expanding the probability of receiving the benefit. In expectation, they both increase the expected rent subsidy a household receives in a given state.

¹²The expression is typical in the social insurance literature. See (Chetty, 2006; Lieber and Lockwood, 2019; Finkelstein, Hendren, and Luttmer, 2019; Gadenne et al., 2021; Deshpande and Lockwood, 2022).

 $^{^{13}}$ The first-best insurance contract is the Arrow-Debreu security $x(w, p^h)$, which perfectly smooths marginal utility between all income-rent states. Both federal rental assistance and income-based cash welfare are second-best. Federal rental assistance transfers in-kind while income-based cash transfers do not target rent prices. However, they are the current mechanisms by which the federal government can insure households. In the full structural model, I compare federal rental assistance to cash transfer indexed on both wages and rent.

4.2 Decomposition: Rent and Income Risks

To understand what makes rental assistance a different form of insurance from cash welfare, I decompose the covariance to highlight two distinct consumption risks: insurance against rent price risk and insurance again conditional income risk.

By the law of total covariance, the insurance value term can be decomposed into two components¹⁴:

$$Cov(v_{y}(w,p),Z(ph-\tau wP)) = \underbrace{Cov(E[v_{y}(w,\bar{p})\mid \bar{p}=p],E[Z(\bar{p}h-\tau wP)\mid \bar{p}=p])}_{\text{Rent Price Risk Insurance Value}}$$
(4.5)

$$+ \underbrace{\mathbb{E}_{p} \left[\operatorname{Cov}(v_{y}(w, \bar{p}), Z(\bar{p}h - \tau wP) \mid \bar{p} = p) \right]}_{\text{Income Insurance Risk Value}}$$
(4.6)

The first term is how much rental assistance insures between rent price states, which I call the rent price risk insurance value. If high rent prices generate both high average marginal utility and high average rent subsidies, then rental assistance is valuable rent price insurance. The second term is the average value of insurance *within* a given rent price state, which I call the income insurance risk value. If there are meaningful differences in income risk or the effects of income risk on marginal utility across rent price states, then rental assistance has scope to target more rent subsidies to those rent price states. In other words, rental assistance can target rent subsidies to those who value income insurance more if rent prices correlate with that income insurance value.

Below, I discuss in detail of each of these terms and provide concrete case studies of how rental assistance targets these risks.

4.2.1 Rent Price Insurance

Rental assistance naturally insures rent price risk by scaling benefits with the rent price. Households may value this insurance if rent price volatility itself affects the marginal utility of income, as captured in Equation 4.5.

Because rent prices enter the budget constraint, it's possible to derive how they affect the marginal utility of income. Adapting the commodity price risks of Turnovsky, Shalit, and Schmitz (1980) and Gadenne et al. (2021) to rent, the derivative of the marginal utility of income with respect to rent price is

$$v_{yp} = \frac{v_y}{p^h} \alpha_h \left(\gamma - \eta_y \right) \tag{4.7}$$

¹⁴This is similar to the decomposition done in Deshpande and Lockwood (2022), which examines how DI insures health versus non-health risks.

where h^* is the optimal housing choice, α_h is the share of the budget devoted to housing, γ is the coefficient of relative risk aversion, and η_y is the income elasticity of housing. In words, household marginal utility increases when rent prices increase if the household inelastically consumes a large amount of housing relative to their income¹⁵. If so, the rent price increase acts as a large budget shock since households cannot easily substitute away from the price increase. By subsidizing rent expenditure shares, rental assistance implicitly targets this change in marginal utility through targeting the rent expenditure share α_h . Households with large expenditures on housing naturally receive more subsidy, creating a covariance between v_y and the subsidy.

Despite a theoretical covariance, households may already have partial insurance against volatile rent prices through local labor markets. Because rent prices and wages are connected through labor demand shocks, wages can already hedge rent price risks that households face. This erodes the value of insuring households against volatile rent prices.

However, not all households have a correlation between wages and rent prices. As shown in Appendix C.1, rental assistance uniquely insures the rent prices of households whose income does not co-move with rent prices. Intuitively, as rent prices increase, households receive more in rent subsidies that is not offset by increases in their income due to wage-rent correlation. These are the households that will demand rent price insurance, often being those on fixed incomes such as the elderly or others reliant government programs.

As motivated by this theory, I use elderly households on Social Security as a case study of the consumption risks of rent prices. Since Social Security does not adjust benefits with rent prices, elderly households have lower correlation between income and rent prices than those in the general population. Therefore, elderly renters on Social Security are more likely to demand rent price insurance and better isolates the effect of volatile rent prices on household budgets.

To approximate the rent price risk these elderly households might face, I examine elderly renters on Social Security who initially face similar rent prices that later diverge. To choose elderly renters who face similar rent prices, I select those in the two middle quartiles of the CZ rent price distribution. I then approximate the rent risk these households face by selecting the top and bottom 20% of commuting zones that had the largest changes in rent prices between 2004 and 2019 in the ACS. In essence, each group represents the 20% probability tails of the elderly renter's rent price growth distribution if they choose to remain in their current CZ¹⁶. Following Equation 4.7, comparing these rent price shocks to the housing budget share

¹⁵It is actually theoretically ambiguous whether households prefer to smooth price volatility. Households may be sufficiently elastic relative to their risk aversion to *prefer* price volatility. This allows them to consume lots of the good when prices are low and substitute away when prices are high. This is unlikely to be the case with housing given the elasticity of housing consumption.

¹⁶The assumption that they remain in their CZ will overstate the potential rent price risk of these households. In the Appendix, I show evidence of the migration rates of elderly renters between CZs as an indication of their elasticity to these rent price changes.

provides descriptive evidence of whether households demand insurance against this rent price growth 17.

I document this case study in Figure 4.1, showing rent prices pose a moderate risk to consumption of elderly retired households. Elderly households in the top quintile of rent increases pay 21.5% more in rent 15 years later, or \$1,453 more, nearly 6 times the increase in rents for the bottom quintile. For comparison, the average Social Security benefit in 2019 was \$18,036, making the rent increase is roughly equivalent to one month of Social Security benefits. In terms of budget shares, the top quintile of households devote 3.9 percentage points more of their budget to rent, which is a 10.7% increase from the initial year. If anything, the bottom quintile has a slight decrease in their income devoted to housing. In total, elderly households who remain in their CZ face moderate consumption risks relative to their Social Security incomes, suggesting possible insurance benefits from smoothing this risk.

Rental assistance naturally subsidizes the rent expenditure share in Figure 4.1 Panel B. As the rent expenditure share increases, households in rental assistance will receive more rent subsidies to bring their rent expenditure share to 30%. In this example, for households who would receive positive subsidies from rental assistance, the rent subsidies would increase by an average of \$568 per year for the top quintle. This provides insurance against this type of consumption risk coming from these housing costs conditional on elderly households receiving this assistance.

Unlike elderly households, households who work do have natural insurance against these rent price shocks. Rent prices are directly related to labor demand for a location, leading to income as a hedge against rent price increases. Equation 4.5 incorporates this, given that the expectation is over the wage w. The example above shows the extreme, where elderly renter households have no hedge against rent price increases. The insurance value of rental assistance depends on whether eligible households benefits from the local income increases when their rents increase.

4.2.2 Income Insurance, Conditional on Rent Price

The conditional income insurance component (Equation 4.6) captures how rental assistance may provide targeted income insurance to households in high rent locations. The key insight is that rent prices may act as a progressive wealth shock, where income losses hurt more when you face expensive necessary housing costs. By providing higher subsidies in high rent locations, rental assistance better targets this income insurance demand than income-based cash welfare.

As an example, suppose that person A lives in San Francisco, California and person B lives in Saint Louis, Missouri. Both experience a work disability shock that leaves them unable to work, and both reduce their luxury consumption to mainly consume necessities like food and housing. However, person A may

 $^{^{17}\}mbox{For more}$ details on the construction of these figures, see Appendix C.2.1

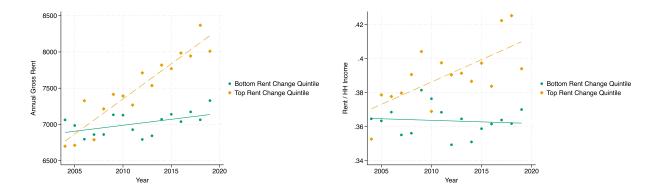


Figure 4.1: Changes in elderly rent (Panel A) and rent expenditure share (Panel B), ACS. *Notes:* Sample of elderly renters on social security is limited commuting zones where rent prices in 2004 are in the middle two quartiles. The rent change quintile bins elderly households based on the largest absolute changes in rent prices over time, comparing only the top and bottom quintile for graphical clarity. Elderly households are restricted to those who have non-zero Social Security income and have a household income (across all family members, both labor and non-labor income) of at least \$5,000.

be worse off because housing, one of their remaining consumption categories, is 165% more expensive in San Francisco than Missouri. Therefore, work disability may pose a larger consumption risk for person *A* than person *B*, leading person *A* to demand more income insurance against this type of shock.

Because rent prices enter the budget constraint, this example's extra income insurance demand conditional on rent prices can be formalized. To do so relies on showing how the marginal rate of substitution—which perfectly captures the insurance value of transfers between two states—changes with respect to rent prices. For clarity, suppose that there are two income states, high income state e and low income state u, and that households have constant relative risk aversion γ . The marginal rate of substitution, MRS = $\frac{v_y^u}{v_y^e}$, captures the value of transferring a dollar from state e to state u. As shown in Appendix C.3, differentiating the MRS with respect to price and transforming to elasticities gives

$$\frac{\partial MRS}{\partial p^{h}} \frac{p}{MRS} = \left[\alpha_{h}^{u} \left(\gamma - \eta_{y}^{u} \right) - \alpha_{h}^{e} \left(\gamma - \eta_{y}^{e} \right) \right]$$
(4.8)

where α_h^s and η_y^s are the housing budget share and income elasticity of housing in state s, respectively. This implies high rent price states have higher income insurance if households spend more of their budget on housing when low income. This is further exacerbated if households are relatively income inelastic in their housing consumption, making rent prices act more as a wealth shock. Once again, rental assistance naturally covaries with income insurance demand by targeting budget shares that directly enter into marginal utility. This theory implies that examing housing budget shares between income states is indicative of income insurance demand. As shown in Equation 4.8, if the difference in the ratio of housing to income rises with rent prices, then this is linked to potentially higher income insurance demand.

As a second case study, Figure 4.2 shows how unemployment impacts the ratio of rent to income conditional on your location's rent¹⁸, providing descriptive evidence for Equation 4.8. I map the income states e and u to employed and unemployed households, respectively, with the goal of measuring their budget shares α_h in each state by rent price. If the difference in the budget share increases as rent price increases, this suggests that income insurance demand is also increasing with rent price.

To carry out this case study, I cross-sectionally examine the budget shares of the employed and unemployed in the ACS between CZ rent price quintiles. In order to have a defined budget share, I restrict to households with non-negligible household income, which may come from non-labor or familial earned income. This descriptive example matches the logic of Equation 4.8 if the employed and unemployed in a given quintile are similar households facing idiosyncratic risk of unemployment.

Indeed, Figure 4.2 shows that those who are unemployed in high rent locations face a greater change in their rent expenditure share between employment states than those in low rent locations, suggesting that housing costs exacerbate income risk more in high rent locations. This demonstrates that income shocks, such as unemployment, may pose more of a consumption risk in high rent price locations because rent prices may pose a greater constraint on the household budget in the unemployed state. Rental assistance subsidizes the unemployed state in the high rent locations more than in the low rent ones, naturally subsidizing this greater consumption risk.

4.3 The Timing of Insurance

The final targeting component of this insurance is the timing of when households receive it. Because assistance is rationed and waits are long, it provides greater assistance to persistent idiosyncratic shocks rather than temporary aggregate shocks. By having a limited amount of assistance to distribute, assistance does not expand when demand increases and households must wait to receive the assistance. If household consumption shocks are temporary, they may not receive the assistance in time to smooth the shock. On the other hand, households with persistent shocks are able to wait for the assistance, and they receive larger benefits than if the limited assistance was divided between all eligible households.

As a third case study, I show how the receipt of rental assistance responds to a temporary versus a persistent income shock in the PSID. In Figure 4.3, I use the PSID to construct an event study of receipt of rental assistance in response to unemployment versus disability¹⁹. While an unemployment shock does not significantly increase the probability of receiving rental assistance, a disability shock does, with a slow entry into rental assistance after the disability shock. This is suggestive evidence that rental assistance selects on those shocks that are able to outlast the waitlist.

¹⁸For details on the construction, see Appendix C.2.2

¹⁹For details on the construction of the event study, please see C.2.3

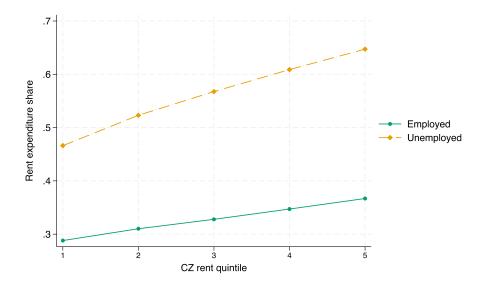


Figure 4.2: Average Ratio of gross rent to household income by employment status and CZ rent quintile, ACS 2005-2019.

Notes: Sample of employed and unemployed households by CZ rent quintile, ACS 2005-2019. CZ rent quintile is estimated using the constructed wages and rents in Appendix B.1. Households are limited to those who have household income, from all sources, greater than \$5,000 (82% of the sample).

In addition, the waitlists allow housing authorities to have discretion in targeting. Housing authorities can select those from the waitlist whose characteristics correlate with marginal utility. However, this can be abitrary and entirely depend on the local housing authority.

4.4 Government Costs

While providing insurance, a marginal expansion of rental assistance increases government costs according to the additional amount of benefits received plus the fiscal externalities of the subsidy. I consider the direct fiscal externalities relating from the rent subsidy itself and also the indirect fiscal externalities on government taxing and spending.

The direct fiscal costs of a marginal expansion are²⁰

$$MC = \underbrace{E[Z(p^{h}h^{*} - \tau wP^{*})]}_{\text{Transfer Value}} + \underbrace{E[Z(p_{h}\frac{\partial h}{\partial \beta} + \frac{\partial p^{h}}{\partial \beta}h - \tau w\frac{\partial P}{\partial \beta})]}_{\text{Moral hazard}}$$
(4.9)

The first term is the expected mechanical increase in rent subsidies to households. The second term is the direct fiscal externalities on the government budget. Since household rent is now tied to income and not rent prices, households behaviorally increase housing consumption and decrease labor supply.

 $^{^{20}}$ I omit the indirect fiscal costs on taxation and spending for clarity. They enter through reduced labor supply and depend on the tax/spending rates of the given programs.

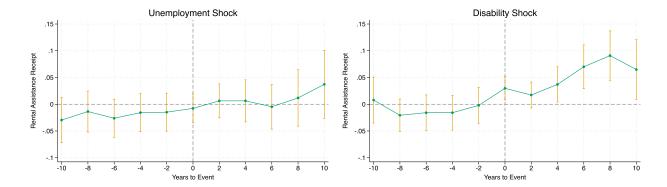


Figure 4.3: Event studies of rental assistance receipt in response to unemployment (Panel A) and disability (Panel B), PSID 1990-2022.

Notes: Sample of household heads who experience their first recorded unemployment or disability shock in Year 0. Disability is defined according to Low and Pistaferri (2015). Unemployment defined as households being employed for two consecutive surveys prior to a recorded unemployment. To maintain a consistent panel in the PSID, odd years are dropped. Event study created using the imputation method of Borusyak, Jaravel, and Spiess (2024).

In addition, if the government must pay a landlord to provide housing, the landlord may increase the rent price that the government must pay (Collinson and Ganong, 2018). This increases the costs of transferring subsidies to recipients, yet landlords still value the transfers that they receive from the program. The losses from this partial incidence depend on the marginal utility gains of landlords relative to the government costs of providing these funds. To allow for flexibility, I assume that the social planner place welfare weight λ_{ℓ} on the transfers T_{ℓ} that they receive.

The net benefits of the program will depend on if the extra insurance benefits from targeting rent prices outweigh the extra moral hazard costs of distorted housing consumption and prices. Relative to traditional cash welfare, the increased targeting of marginal utility comes at a cost. Traditional cash welfare only creates moral hazard on labor supply and generally avoids other fiscal externalities.

4.5 Sufficient Statistics Approach to Estimating the Net Insurance Benefits

Before placing more structure on the conceptual framework, I estimate the net insurance value given current estimates from the literature and the sufficient statistics approach. The sufficient statistics sidesteps specifying the full model by estimating (or calibrating) the "sufficient" moments necessary for knowing the net benefits of a marginal expansion of a program. The key empirical objects are the covariance between marginal utility and the rent subsidies along with the elasticities that determine moral hazard. I use data from the PSID and HUD administrative data to estimate the key rent subsidy and consumption distributions of households.

For this approach, I estimate the insurance value given the current state of the social safety net. Some of the value of rental assistance may come from the fact that households are not fully insured against various shocks through other programs. For example, if disability insurance does not adequately insure households against disability shocks, then rental assistance may provide additional insurance benefits against these costs. This does not imply that expanding or contracting rental assistance is necessarily more efficient than expanding other social welfare programs. However, by holding fixed other social safety net programs, this improves estimation of the insurance value for the program as is.

Insurance Value The insurance value is the covariance between marginal utility and the rent subsidies (Equation 4.4). The components needed to measure these objects some definition of normalized marginal utility with respect to consumption and how much households receive in rent subsidies.

I begin with marginal utility, using a functional form that is a transparent mapping from household consumption to marginal utility. I assume Constant Relative Risk Aversion (CRRA) and that numeraire consumption is linearly separable from labor supply and housing subutility:

$$u(c,h,P) = \frac{c^{1-\gamma}}{1-\gamma} + s(h,P)$$

By doing so, marginal utility is measurable solely through numeraire consumption. Assuming households maximize, households will equate the marginal utility of income to the marginal utility of numeraire consumption, or $v_y = c^{-\gamma}$ (Finkelstein, Hendren, and Luttmer, 2019). This simplifies the estimation of marginal utility at the cost of restricting interaction effects of housing and labor supply on the marginal utility of consumption.

For *c*, I use the non-housing expenditures in the PSID. I define non-housing expenditures as expenditures excluding rent, utilities, and mortgage payments. By using expenditures, I implicitly assume that all non-housing prices are uniform across the US. If untrue, this will likely bias the consumption of high rent price households upward, lowering their marginal utility. This would work against the insurance value of rental assistance since households in high rent locations receive larger subsidies. I adjust *c* for household size and winsorize outlier consumption. To adjust for household size, I use the square root of household size as an equivalence scale. Finally, to keep outlier low consumption values from dominating average marginal utility, I winsorize consumption at the 10th percentile.

As shown in Equation 4.4This marginal utility must be normalized, and I do so with respect to the non-assisted states of the given risk type. Theoretically, this represents how much a household would be willing to pay in non-assisted state consumption ex ante for the current design of rental assistance. This matches the reality of how these programs are funded and follows the conventions in the literature Chetty (2006); Finkelstein, Hendren, and Luttmer (2019); Deshpande and Lockwood (2022). Alternatively, marginal utility could be normalized with respect to all states, or with different definitions of risk types. I show robustness

to the normalization of marginal utility in Appendix C.5.

With a specified marginal utility in hand, I next need to measure how these covary with rent subsidies. The rent subsidy amount is not observed in the PSID, so I use the HUD data to predict the household's rent subsidy conditional on commuting zone rent and household income. I then use these expected subsidies to predict the household subsidy amount in the PSID data. This leads to a dollar amount of subsidies that a household receives given their household characteristics and location. I detail how I estimate the expected rent subsidy and the results in Appendix C.4.

For simplicity, I assume that all recipient households receive the equivalent of a voucher to avoid complications in measuring the costs of public housing. Since public housing is a durable good owned by the government, it is difficult to properly measure the depreciation of public housing over time. By assuming public housing functions as a voucher, I am following a similar strategy of Waldinger (2021) in assuming public housing costs are similar to market-rate costs. While public housing does not have private landlords that receive a portion of the benefits of rent subsidies, Olsen (2002) finds that there are additional cost overruns that may function as transfers to developers in building public housing or even as deadweight losses.

Moral Hazard Elasticities For moral hazard costs (Equation 4.9), I use elasticities from the literature. This involves the fiscal externalities from changes in labor supply, housing consumption, and use of other social safety net programs.

For the labor supply and social safety net fiscal externalities, I rely on Jacob and Ludwig (2012), who use a voucher experiment in Chicago to estimate a 10% change in earnings upon receipt. The experiment used a lottery to determine who received rental assistance, and labor supply was subsequently measured in unemployment data. While possibly not representative of the average voucher recipient, these estimates are the most credible to date on the labor supply effects of the program. Using this lottery, they also find significant effects on other government programs, including food stamps and TANF. Together, I estimate an average fiscal externality from labor supply and the safety net of \$917 per year conditional on receipt.

For the housing moral hazard, I combine Reeder (1985) and Collinson and Ganong (2018) to estimate the household ex post valuation of rental assistance housing net of both passthrough and in-kind consumption. Reeder (1985) finds that a dollar of rental assistance is valued at 83 cents by households. He measures the demand for rental assistance using household housing expenditures before and after receiving rental assistance. I use this as my baseline estimate since it is also used in other work that quantifies the value of rental assistance, making my estimates comparable (Hendren and Sprung-Keyser, 2020). For the incidence of rental assistance vouchers, I use the estimates of Collinson and Ganong (2018). They exploit changes in the FMR rent ceiling to measure changes in hedonic quality and rent price. They find that changing the rent

Table 4: Sufficient statistics estimates of the insurance value, government costs, and MVPF of federal rental assistance.

Risk Type	Prob Assist	EAWTP	Ins Val	Gov Cost	MVPF
Less Than HS	0.12	851.47	276.66	611.93	1.39
HS Grad	0.06	418.07	224.97	270.76	1.54
Average	0.07	518.33	260.06	349.68	1.51

Notes: Sample of non-college educated household heads in PSID with reported consumption, 1998-2022. Estimate of the ex ante willingness to pay and the insurance value follow Section 4.5. Government costs represent the average government cost per year given the probability of assistance, the transfer costs net of landlord passthrough, and the fiscal externalities of \$917 conditional on receipt. The MVPF is the marginal value of public funds, as detailed in Hendren and Sprung-Keyser (2020).

ceiling raises rents by \$0.46, while only \$0.05 is in improved measurable hedonic quality. While specific to the FMR rate setting policy, I assume that this incidence applies to the entire voucher program.

Results I report the net welfare benefits in Table 4. In the benchmark specification, I find that households value rental assistance at \$1.51 per dollar of government cost. The magnitude of these benefits relative to costs is due to the insurance value. For the average household, over half of the value of the program comes from insuring households against consumption risk. Previous literature estimates that use similar elasticities but exclude the insurance term find the program generates \$0.66 per dollar of government costs. By generating benefits larger than costs, this qualitatively implies we should switch from contracting rental assistance to expanding it.

I compare this result to other estimates of the insurance value of traditional government insurance programs. Given that most other papers in the literature use $\gamma=2$, I highlight that the MVPF of rental assistance is 1.77 when assuming $\gamma=2$. Deshpande and Lockwood (2022) find that disability insurance generates an MVPF of 1.76, nearly the exact same as rental assistance. The close estimates likely reflect that both programs have screening mechanisms that select those with persistent income shocks. For other insurance programs, such as unemployment insurance and Medicaid, the MVPF often lies between 0.5 and 1.2 (Hendren and Sprung-Keyser, 2020; Finkelstein, Hendren, and Luttmer, 2019). The evaluations often found that the consumption variation of households that may receive these programs is small relative to the moral hazard costs.

The fact that the ex ante willingness to pay outweighs the government costs holds under a variety of paramterizations. In Appendix C.5, I show robustness to the coefficient of relative risk aversion, the normalization of marginal utility, how much the landlord transfers count as a fiscal externality, and other key parameters of this derivation. The MVPF only falls below one if the risk aversion parameter is below one or if the entirety of the landlord transfer counts as a cost.

4.5.1 Insurance Value Decomposition

Following the conceptual framework, I then decompose the insurance value into whether it insures rent price risk or income risk. To carry out this decomposition, I discretize the CZ rent distribution into rent ventiles²¹, permitting an estimate of the covariance between rent price ventiles (Equation 4.5) and within rent price ventiles (Equation 4.6). Therefore, the income insurance component is any consumption shocks that occur within a rent price ventile, while the rent insurance component is any rent shocks that shift households between rent price ventiles.

I show the results in Table 5 for the average household first. More than all of the insurance value comes from insuring income conditional on rent price, with rent price insurance negative yet nearly zero. This result is theoretically motivated by urban spatial models, where the average household is likely indifferent between each rent price ventile because wages naturally adjust to offset the utility losses from rent prices. Therefore, they can only derive value from the program if their income diverges from the typical wages in the area.

For these average households, the income insurance value increases from \$144 in the bottom ventile to \$475 in the top ventile, theoretically following the argument laid out in Section 4.2.2. These results imply that households in high rent price states have lower numeraire consumption when they are low income than those in low rent price states. Rental assistance then targets these subsidies to those households in the high rent price states.

The average household may have different insurance value from the program than other types of households. As shown in the conceptual framework examples, households such as elderly renters may benefit more from the rent insurance component because they have a fixed income. In Panel B of 5, I show the insurance values and its decomposition for each of the household example in the conceptual framework. I define the risk set as all household-quarters that meet the criteria, such as elderly renters on Social Security. For simplicity, I keep maginal utility normalized to the whole population.

I find that the insurance value decomposition confirms the descriptive analysis of the case studies. First, I find that elderly households do indeed benefit from the rent insurance component, receiving 29% of their insurance value from insuring rent price risks. The other remaining risk comes from insuring differences in income, whether from differences in Social Security benefits due to the Social Security schedule or differences in other sources of family income. Additionally, disabled households have a much higher insurance value than the unemployed, supporting the evidence that the eligibility and subsequent waitlist screening mechanisms are more likely to select these disabled households.

 $^{^{21}\}text{I}$ show robustness to the choice of discretization in Appenix C.5.4

Table 5: Decomposition of insurance value into rent and income components.

Panel A: Main Insurance Value Decomposition Results

		1	
	Total Insurance Value (\$)	% Rent Insurance	% Income Insurance
Risk type			
Less than HS	275.15	-1.72	101.72
Graduated HS	176.15	-2.43	102.43
Average	199.05	-2.27	102.27

Panel B: Other Risk Types

		. J I	
	Total Insurance Value (\$)	% Rent Insurance	% Income Insurance
Risk type			
Elderly Renter on SS	334.34	29.05	70.95
Head Disabled on SSDI	469.51	-3.53	103.53
Head on UI	206.47	-14.73	114.73

Notes: PSID, 1998-2022. Normalization of marginal utility relative to all states conditional on a given risk type. Decomposition done discretely by CZ rent price ventile, estimating the between versus within covariance of marginal utility and rental assistance subsidies. The between covariance represents the rent price insurance value, while the within covariance represents the income insurance value conditional on risk type.

Together, these results imply that the average household values the persistent income insurance component of rental assistance, but that individual household types may value the rent insurance component. This suggests that expansion of rental assistance in persistently high rent locations would generate more welfare than a universal expansion, allowing for more income insurance for high rent households. However, a targeted expansion of rental assistance to particular groups, particularly those on fixed incomes, in volatile rent locations will also prove valuable at insuring both rent and income for these households.

While transparent, the sufficient statistics approach has several limitations in the context of federal rental assistance. These estimates rely on the accuracy of the parameter estimates from the literature, which may not be well calibrated to the exercise carried out here. In addition, the sufficient statistics approach assumes a flat marginal expansion in all states of the world. In reality, marginal expansion can target certain states and households. For example, reducing wait times will directionally expand rental assistance to households who have less priority in the wait list. It is ambiguous whether these households have higher or lower marginal utility of receiving a rental assistance offer than the average recipient in the sufficient statistics estimate. Finally, there are other marginal changes in the program not considered by the sufficient statistics estimate. For example, marginally adjusting the income share parameter τ affects the level of income insurance as well as who selects into the program. To correctly model these counterfactuals, we need a structural model of household behavior, wage-rent risk, and the federal rental assistance program.

5 Lifecycle Model: Household Dynamics and Federal Rental Assistance

The primary goal of this lifecycle model is to evaluate counterfactuals that inform which mechanisms of rental assistance generate insurance value for households. A lifecycle model is necessary to capture how persistent shocks interact with dynamic household choices regarding rental assistance. These wage and rent shocks often evolve over many years, affecting when households receive rental assistance and when they have the highest demand for assistance. In addition, households are forward-looking about when they will receive rental assistance. Their decision to work today or be in rental assistance depend on how long their wait is for assistance in the future²².

I choose counterfactuals that both reveal the mechanisms of the program's insurance value and are relevant for policy. First, I compare rental assistance to cash welfare, transfers that are weakly preferred to rental assistance but may lead to differential selection of recipients. I examine not only how cash transfers change the welfare of rental assistance, but also which design elements of rental assistance aid in the targeting of cash transfers. Second, I change rent payments in rental assistance to a flat rent, similar to LIHTC. Proponents argue that this improves the labor supply of households by decoupling rent payments from income, but it may also reduce the targeting of the program by reducing the amount of subsidies that go to households in need.

Finally, the model will contextualize previous reduced form estimates and provide alternative identification of the key parameters affecting the welfare of rental assistance. These parameters include the labor supply effects of rental assistance and the implied quality for households. I find the model closely matches the reduced form estimates used in Section 4.5, reinforcing the derived insurance value.

The key forces of the model are as follows. In the face of joint rent-income risk, households choose consumption, housing, and labor supply to maximize lifetime expected utility. When income eligible and after an uncertain waiting period, households can choose whether to enter rental assistance or stay in the private market. Each quarter, household i with characteristics X_{it} makes choices over housing consumption H, numeraire consumption C, rental assistance receipt D(conditional on application D), and labor supply D. They maximize lifetime utility

$$V_{it} = \max_{c,h,P,Z} \operatorname{E}_{t} \left[\sum_{s=t}^{T} \beta^{s-t} U\left(c_{is}, h_{is}, P_{is}\right) \right]$$

I first begin with the household choices of h, c, P, and Z(R).

²²As a complementary discussion of many of these dynamic considerations, see Jacob and Ludwig (2012).

5.1 Household Problem

I begin with the household budget constraint, then define household preferences, and finally the housing markets that households participate in.

Budget Constraint

Household income comes from their assets, labor earnings, and government programs. Their rent expenditures depend on whether they are in or out of rental assistance. Household i has the following intertemporal budget constraint:

$$A_{it+1} = r[A_{it} + w_{it}P + B(X_{it})]$$
(5.1)

$$-\left(c + \underbrace{Z_{it}(R)\tau\left(w_{it}P + B(X_{it}) + (r-1)A_{it-1}\right) + (1 - Z_{it}(R_{it}))p_{\ell t}^{h}}_{\text{Housing expenditures}}\right)\right]$$
(5.2)

where A_{it} are assets²³ with interest rate r, $w_{i\ell t}$ are wages, and $B(X_{i\ell t})$ are government transfers. The first three terms of the right hand side are the total income and assets that the person receives, with the remaining terms relating to household expenditures on numeraire and housing. Housing expenditures are a function of whether or not the household is receiving rental assistance Z_{it} , leading to different prices for the quality of housing that they receive. When in rental assistance $(Z_{it}=1)$, households do not face market rent and instead pay τ share of their total income, including asset income, as rent. When not in rental assistance $(Z_{it}=0)$, households purchase market rent housing h at price $p_{\ell t}$. All consumption is normalized by equivalence scale $e(k_i) = \sqrt{k_i}$ where k_i is the number of people within the household.

In addition to labor and asset income, I model a rich social safety net in $B(X_{it})$, including food stamps, unemployment insurance, social security, and an emergency housing program. I structure food stamps similar to the actual food stamps program, except treating it as a cash transfer. Importantly, I include food stamps' excessive housing cost deduction, where households may deduct housing costs above 50% of their income up to a cap. This is the second largest source of housing assistance in the US, providing an average of \$356 per month in deductions to households in the food stamps program in 2023 (Monkovic and Ward, 2025; Rosenbaum, Tenny, and Elkin, 2002), which roughly translates to \$70 per month in extra benefits. For unemployment insurance, households receive 75% of their labor income for a single period (Low and

 $^{^{23}}$ Assets provide the first means of self-insurance against wage-rent shocks. Households may also save for retirement but must spend down their assets in the final period. I assume that household cannot borrow $A_{it} \geq 0$. This prevents households from borrowing against future receipt of rental assistance This still allows households to preemptively save in anticipation of receiving rental assistance. This could occur if households anticipate reducing their labor supply while in rental assistance. Enforcing a borrowing constraint prevents unrealistic liquidity that low income households do not have.

Pistaferri, 2015). I model social security after the actual program, including all the schedule kinks. Finally, I assume an ad hoc emergency housing program that guarantees a minimum housing consumption floor above an estimated minimum housing constraint. This acts as a last-resort form of insurance intended to capture any other methods that households may insure against homelessness.

Household Preferences

Similar to the conceptual framework, households gain utility from numeraire and housing consumption and disutility from labor supply. The insurance value of rental assistance depends on their risk aversion and the relative inelasticity of housing when low income. Moral hazard also depends on their elasticity of housing consumption and their disutility of labor supply.

Subject to the budget constraint (Equation 5.1), household i splits consumption between a numeraire good and housing with a labor supply disutility. I parameterize this as

$$U_{it}(c,h,P) = \frac{\left(c^{\alpha}(h-\bar{h}_i)^{1-\alpha}\exp(\phi P)\right)^{1-\gamma}}{1-\gamma}$$
(5.3)

Crucially, this functional form places explicit structure on the housing Engel curve, and hence the rent expenditure shares of households. Unconstrained, households prefer to expend share α on numeraire consumption and $1-\alpha$ on housing consumption. However, they face a minimum housing constraint \bar{h}_i , and they become increasingly inelastic in housing consumption as they approach this housing consumption floor. This empirically approximates actual housing Engel curves and is used in other literature on rental assistance (Sieg and Yoon, 2020), at the cost of restricting the parametric form of the housing elasticity. As shown in the conceptual framework, this has important implications for the insurance value that households receive from rental assistance, explicitly parameterizing the marginal utility derived in Equations 4.7 and 4.8^{24} .

The other parameters are canonical in the lifecycle model literature. Preferences are standard CRRA governed by the curvature γ . Work disutility is the parameter ϕ , which I enforce to be negative. This disutility raises the marginal utility of consumption while working, but it does not change the marginal utility of numeraire versus housing consumption, leading the budget share of housing to be orthogonal to the choice of labor supply. For the details of how this functional form implies these preferences, see Appendix D.1.

²⁴For additional details, see Appendix D.1.1

Housing Markets

Households are either private market renters, homeowners, or in rental assistance. Households begin the lifecycle as renters and transition into the other housing markets.

To receive rental assistance, households must be eligible, the waitlist must be open, and they must spend time on the waitlist. I set the eligibility to get into the program at 30% AMI, allowing the AMI to differ by the location wage²⁵. Once in the program, households may remain in the program until they reach the 80% AMI limit.

The timing of receiving rental assistance depends on when waitlists open and how long the waitlist is. To model the opening and closing of the waitlists, households randomly apply to rental assistance with probability $p(X_{is}, w_{it}P_{it}, p_t^h)$, with $R_{i\ell t}$ as an indicator for application to rental assistance. This application probability accounts for both the eligibility rules—which can reduce the application probability to zero—and the actual decisions of the housing authority to close the waitlist.

After application, households must wait to receive assistance. Each period, they have a probability $\pi_{\ell t}(Z_{i\ell t}=1\mid Z_{i\ell t-1},R_{i\ell t}=1,t_0,X_i)$ of receiving rental assistance $Z_{i\ell t}$, conditional on the initial date t_0 of their application and characteristics X_i . When $Z_{i\ell t}(R_{i\ell t})=1$, households receive assistance and may remain in the program while their income is below the 80% AMI threshold.

Once in the program, households must choose housing consumption h_i^Z and pay τ ($w_{i\ell}P + B(X_{i\ell t}) + rA_{it-1}$) in rent. This take-it-or-leave-it approach is key to the moral hazard costs and selection into the program, as households must decide between distorted housing consumption h_i^Z or private market housing. This matches key aspects of the program that limit the choice of housing among recipients, and households may even consume less housing than desired in order to stay in the program. While an individual household's housing quality is fixed, I still allow housing quality to be heterogeneous across households, distributed $h_i^Z \sim \text{Beta}(\alpha_{h^Z}, \beta_{h^Z}, h_{min}^Z, h_{max}^Z)^{26}$. This reflects that, although an individual household faces a take-it-or-leave-it policy, there is still wide heterogeneity in the types of units that a household finds available to lease. I set the minimum quality above \bar{h}_i and the maximum quality at what is implied by the FMR at the median rent.

Because only renters are eligible for the program, I allow for homeownership to realistically control for the number of eligible households over the lifecycle. Households exogenously enter homeownership with probability $q(p_t^h, w_{it}, t)$, and subsequently household rent prices are fixed at p_t^h for the remainder of their

²⁵This is a simplification of the true eligibility criteria for entry. However, it simplifies the application decision by reducing the number of households who would be marginal to applying and starting their waitlist clock.

²⁶While non-standard, the Beta distribution theoretically matches the natural caps on quality in the program. HUD regularly certifies that apartments meet minimum quality standards, and, as discussed, HUD caps the maximum rent of units in the program. These limits imply that more households will be marginal in housing quality than preferences that are typically assumed to be distributed normal or extreme value.

life, approximating the price lock-in of a mortgage.

5.2 Wage-Rent Process

Following the conceptual framwork, I model a joint wage-rent process with shocks between rent price states and within rent price states. To do so, I specify a location wage-rent process and an idiosyncratic wage process. The location wage-rent process allows location wages to naturally hedge rent shocks, while the idiosyncratic component of wages generates consumption shocks within a given rent price state.

The location process is AR(1) in log wages μ_t and log rents p_t^h :

$$\mu_{\ell t} = \rho_{w\ell} \mu_{\ell t-1} + \varepsilon_{\ell t}$$

$$p_t^h = \rho_{h\ell} p_{t-1}^h + \nu_{\ell t}$$

where

$$(\varepsilon_{\ell t}, \nu_{\ell t}) \sim N(0, \Sigma_{\ell})$$

This AR(1) process aligns with previous research that focuses on estimating this process (Kueng et al., 2023). I crucially allow for correlation between wage and rent shocks at the location level, which allows for location wages to hedge against rent price risk.

At the household level, household rents do not differ from location rents, but household wages have an additional idiosyncratic component. I model this as

$$w_{it} = \mu_{\ell t} + f_i + X_{it}\beta + \eta_{it} + e_{it}$$
 (5.4)

where f_i is a household fixed effect, X_{it} are household characteristics (e.g., age, age squared, disability, etc.), e_{it} is a transitory component, and η_{it} is an AR(1) process:

$$\eta_{it} = \rho_{wi}\eta_{it-1} + \zeta_{it}$$

$$\zeta_{it} \sim N(0, \sigma_{wi}^2)$$

Households can also experience an idiosyncratic unemployment shock with probability λ and re-employment probability of δ , as in Low and Pistaferri (2015).

At the start of the lifecycle, households draw a location ℓ and idiosyncratic wage from

$$(\mu_{\ell 0}, p_{\ell 0}) \sim N(\xi_{\ell}, \Sigma_{\ell}^{0})$$
$$\eta_{i0} \sim N(\xi_{i}, \sigma_{wi0}^{2})$$

This wage-rent process matches the targeted insurance motives in the conceptual framework. House-holds experience rent price risk (Equation 4.5) that is partially insured by correlated wage risk. However, within a rent price state, households experience idiosyncratic income shocks that they would prefer to insure (Equation 4.6). The insurance value of rental assistance relative to cash depends on whether the dispersion in rent prices creates meaningful consumption risk.

5.3 Model Solution

The model has a large state space with no closed form solution, so I implement an adaptive sparse grid method necessary to solve and estimate the model. The model has 12 dimensions in total, 6 of which are uncertain and 3 of those uncertain dimensions are continuous. In addition, the solution for optimal consumption has no analytic solution, further increasing the number of evaluations of this large state space necessary to solve the model. If I wish to solve the model at a quarterly level, this large dimensionality makes the model intractable to compute with traditional gridded interpolation. In Table 6, I show that given a chosen resolution for the grid, a traditional grid would require 2.12 trillion value function evaluations to construct the interpolant. This model would be difficult to solve once, and intractable to perform indirect inference.

To overcome the dimensionality, I use an adaptive sparse grid to minimize the grid points needed to construct the interpolant. Adaptive sparse grids are an interpolation technique developed in the approximation theory literature (Griebel, 1998; Ma and Zabaras, 2009; Stoyanov, 2019), and first applied in the economics context in the macroeconomic modelling of Brumm and Scheidegger (2017). These grids seek to minimize the number of grid points needed by specifying simple basis functions that use only grid points that significantly reduce the error of the interpolation. To do so, the algorithm evaluates candidate grid points, and if the difference between the interpolated function and the actual function is outside some error tolerance, the point is included in the grid.

For my lifecycle model, I use an adaptive local linear grid for my interpolant, which only needs approximately 872 million value function evaluations for the same maximum resolution as the traditional grid (~3.4 orders of magnitude smaller). This grid constructs the interpolant using the chosen grid points and linear basis functions. Given this choice of grid, the parts of the state space that are roughly linear will often

Table 6: Model Dimensionality

Dimension	Resolution
Location wage (Working only)	9
Location rent	9
Idiosyncratic wage	17
Asset	33
Priority/receipt (non-homeowner only)	10
Assisted housing quality (non-homeowner only)	5
Homeownership	2
Employed	2
UI eligible (non-employed only)	2
Quarter	200
Brent Optimizer Average Evaluations	12
Expectation grid points	27
Full grid VF evaluations (Working only)	~2.12 trillion
Adaptive sparse grid VF evaluations (All quarters)	~872 million

Notes: Chosen maximum resolution for the interpolation grid along each dimension. For example, the idiosyncratic wage grid has a maximum resolution of 17 grid points. Full grid VF evaluations refers to the number of times the value function must be evaluated across all interpolation grid points, including all points needed for computing the expected value function for next period. The adaptive sparse grid method used in the paper reduces the number of value function evaluations by reducing the number of interpolation points.

need few points to approximate. Instead, the grid will adaptively concentrate points near the nonlinearities inherent in the model, such as consumption decisions near budget constraints or sharp program rule changes. Over course traditional grids, this improves the accuracy of the model by budgeting more points to key areas of the state space. Over fine traditional grids, this improves efficiency by reducing the number of points with only slight changes in interpolation error. With these changes, I reduce the number of value function evaluations by 99%, making the model computationally tractable²⁷.

To further speed up the solution, I implement the model solution on a GPU. Within a period of the life-cycle model, solving the model for each grid point of the interpolant is embarrassingly parallel. Therefore, solving the model is very amenable to GPU architecture. Given vast improvements in GPU performance in recent years, this makes model solutions much quicker than running in parallel across many CPUs.

For other model solution implementation details, please see Appendix D.3.

6 Identification and Estimation

I identify and estimate the model parameters through a combination of external estimation and indirect inference. The externally identified parameters are those that do not depend on dynamic household choices. All the other parameters are estimated by matching an auxiliary model between the simulated data from

²⁷The speed improvement is less than the reduction in the grid points because it is more computationally costly to evaluate the interpolant. Evaluation requires finding all grid points with local support of a desired interpolation point. Being sparse, lookup is costly compared to evenly spaced grids. However, the reduction in grid points more than pays for the extra computational costs of the interpolant.

the model and the true data.

6.1 Estimation

To estimate the model, I rely on indirect inference. In indirect inference, an auxiliary model is estimated on both the simulated and the real data. This auxiliary model is a combination of moments and reduced-form estimating equations that are informative of the underlying preferences. The auxiliary model need not be free of confounding variables, so long as the assumptions of the actual model correctly specifiy the confounding variables. Then, indirect inference minimizes the (weighted) distance between the simulated auxiliary model paramters and the true data. I estimate the model using BOBYQA, a derivate-free optimizer that creates a quadratic approximation of the optimization function (Powell, 2009).

6.2 Externally Identified Parameters

I externally calibrate or externally estimate parameters that do not interact with the dynamics of rental assistance. For some of these externally identified parameters, I rely on literature estimates from papers specifically designed to measure those parameters. For other parameters with less applicable literature estimates, I rely on estimating those parameters outside of the model solution.

6.2.1 Pre-Determined Parameters

I set the intertemporal discount rate β and risk aversion γ to values estimated in the literature, with $\beta = 0.9785$ and $\gamma = 1.5$ being the benchmark case. These values lie in the middle of the distribution of estimated ranges for these parameters in the literature. For the equivalence scale, I use the square root of family size like in Section 4.5. Finally, I set the interest rate r = 1.016 and assume the lifespan is 40 working years and 15 retirement years²⁸. I set the re-employment rate to 0.73 following Low, Meghir, and Pistaferri (2010).

6.2.2 Wage-Rent Process

I estimate the wage-rent process using Census data for the location distribution and the PSID for the idiosyncratic distribution. Importantly, this wage-rent process provides structure to the risks that households face that are outlined in Section 4. These estimates contextualize and complement the insurance value results of that section, further highlighting which primary risks households face.

For the location wages and rents, I estimate a panel autoregressive model using census data from 1970 to 2020, following the sampling and estimation procedure of Kueng et al. (2023). This procedure uses the quality-adjusted mean wage and rent in a CZ as a measure of the local wage and rent prices. For CZ wages

²⁸The current model has 10 retirement years, but I plan to expand to 15 in the full model. The length of retirement particularly affects the welfare costs of rent price risk since retirees have fixed incomes that do not hedge rent prices.

I regress (separately for each year) household wages on a fixed effect for the CZ and controls, such as household age, race, marrriage status, and other household characteristics. I also estimate separate wage distributions by education status to allow for differences in local labor demand. For CZ rents, I follow the same procedure but with controls for housing characteristics, such as the structure, the age of the building, and the number of rooms. The rest of the details of this estimation procedure is in the Appendix D.2.1. From this panel autoregressive model, I obtain the variance-covariance matrix of the location wage-rent process, along with the persistence of the wage-rent shocks.

For idiosyncratic income, I estimate the selection-corrected wage process using the PSID data following Low and Pistaferri (2015). The goal is to identify the variance of wage shocks, both between years and the initial distribution. To do so, I estimate the following first-differences model:

$$\Delta \ln w_{it} = \Delta X_{it} \beta_0 + \Delta \ln \mu_{\ell t} \beta_1 + \xi_{it} + \varepsilon_{it}$$

where X_{it} are household characteristics for quadratic age, marriage status, disability status, and gender of the household; ξ_{it} is measurement error in wages; and $\varepsilon_{it} \sim N(0, \sigma_{w_i})$ is an idiosyncratic error term. Importantly, including location wages removes the variance that comes from CZ wage fluctuations.

The true wage distribution is confounded by the fact that those who receive a low wage offer are more likely to choose to not work. To correct for this, I follow the process outlined in (Low and Pistaferri, 2015), which uses state variation in welfare benefits as a simulated instrument for employment. The idea is that exogenous variation in government benefits identifies the selection into working that can be controlled for in the wage estimation. The full estimation procedure uses GMM, deriving moments that account for the selection and measurement error in the wage process. For more details, see Appendix D.2.2.

I present the variance of the joint wage-rent process in Table 7. Panel A shows the parameter estimates of the annual wage-rent process. Location wages are strongly correlated with location rents for households with a high school level of education, but this correlation is cut in half for high school dropouts, implying their income process does not guarantee as much insurance against rent risk as higher education households. However, location wage risk is dominated by the idiosyncratic component of wage risk.

These wage-rent volatility estimates complement the findings on the insurance value in Section 4.5. The sufficient statistics estimates show that households primarily demand rental assistance as a form of conditional income insurance. Indeed, the parametric estimation of wage-rent risk show that households primarily face risk in their initial rent price and in the idiosyncratic volatility of their incomes. Households have a natural hedge against the remaining annual rent price volatility through the high correlation between wage and rent price shocks. This implies that the insurance value estimates match the actual price

Table 7: Joint Wage-Rent Proess Estimates.

(a) Panel A: Annual Wage-Rent Process

(b) Panel B: Initial Wage-Rent Distribution

Parameter	Estimate	Location Proces
Location Process		$\sigma_{\ln \mu_\ell}$
$\sigma_{\ln \mu_{\ell}}$ (Less than H.S.)	0.064	$\sigma_{\ln p_\ell^h}$
$\sigma_{\ln \mu_{\ell}}$ (H.S.)	0.025	$ ho_{\mu p}$
$\sigma_{\ln p_\ell^h}$	0.036	Idiosyncratic Pr
$\rho_{\mu p}$ (Less than H.S.)	0.250	$\sigma_{\ln w_i}$ Male
$\rho_{\mu p}$ (H.S.)	0.554	$\sigma_{\ln w_i}$ Female
Idiosyncratic Process	•	
$\sigma_{\ln w_i}$ Male	0.171	
$\sigma_{\ln w_i}$ Female	0.195	
-		

Notes: Estimation of location estimates adapted from Kueng et al. (2023) using Census data 1960-2010. Idiosyncratic process adapted from Low and Pistaferri (2015) using PSID 1989-2022. Location log wage-rent prices normalized to zero.

risks that households face.

6.2.3 Housing Consumption Parameters

The parameters that determine the division between numeraire and housing consumption have an analytic solution that does not depend on the dynamics of the model. Isolating the numeraire and housing consumption subutility, the first-order conditions imply

$$p_{\ell t}^{h}h_{it} = \alpha \bar{h}p_{\ell t}^{h} + (1-\alpha)\left(c + p_{\ell t}^{h}h\right)$$

This equation intuitively shows the functional form of preferences. When household total expenditures increase $(c + p_{\ell t}^h h)$, households allocate $(1 - \alpha)$ of their budget to housing expenditures. However, \bar{h} imposes a housing consumption floor for expenditures. This creates a linear expenditure system for housing.

I estimate this equation using the observed expenditures in the PSID and my constructed rent prices with non-elderly renters outside of rental assistance. I exclude elderly renters since their rent costs can include additional assisted-living costs, leading to higher estimates of the minimum rent than for the general population likely to use rental assistance. Importantly, housing expenditures of these households are the reported rent at the family level. These estimates do not represent minimum rent costs of the total apartment since households may adopt shared living arrangements as their expenditures fall. Therefore, the measurement of the minimum housing consumption implicitly captures forms of consumption sharing

and self-insurance against minimum apartment rent costs in a city²⁹.

To address attenuation, I use income as an instrument for total expenditures. Estimating this equation via a simple OLS may be biased for two reasons. First, adjustment frictions in housing can bias the slope $(1-\alpha)$ downward and the intercept \bar{h} upward. If households do not readily adjust housing expenditures to every consumption shock, then households appear more inelastic than without adjustment frictions. This will make it appear that households have a higher minimum housing constraint than in practice, making housing consumption seem more valuable. Second, there may be measurement error in the household expenditures themselves that attenuate the estimates of $1-\alpha$. Because the PSID data relies on self-reported consumption, measurement error is likely the case for the total expenditures.

Using household income as an instrument helps to address these issues. Income is correlated with total expenditures through the budget constraint, and, if housing consumption is correctly specified, it naturally increases housing expenditures solely through total expenditures. Income provides another report to correct measurement error in total expenditures, similar to the argument in Ashenfelter and Krueger (1994). In addition, households more readily adjust housing consumption in response to their income, alleviating issues with the short-run adjustment in housing consumption.

I present the estimates of the housing preference parameters in Table 10. When far from the minimum housing constraint, households prefer to spend 26% of their budget on housing, and they have a minimum housing consumption of \$548 per quarter (or \$183 per month) when facing the mean rent price. My consumption share parameter matches estimates found elsewhere in the literature, which often ranges from 0.22 to 0.3. For the minimum housing parameter, it is difficult to compare literature estimates to the one found here. Often, the literature will estimate a minimum housing consumption for a particular city, with the rent price in that city normalized to one. In addition, the literature will often estimate the relationship between *income* and housing expenditures. This relationship often implies a higher minimum housing expenditure because households can rely on assets to consume housing even when their income is near zero. Therefore, my estimates are conservative on the inelasticity of housing consumption when low income.

I estimate the homeownership transition probabilities using the PSID data. To do so, I estimate a probit of homeownership given household age, log wage, and household rent. Most importantly, this estimation leads households to select into homeownership as they age, placing a natural constraint on the number of households eligible for rental assistance. It also insures households against rent price risk by locking in rent prices, eliminating consumption risk through volatile rent prices for many in the population. I present how

²⁹This is especially important when comparing these estimates to others in the literature. These estimates will either rely on minimum reported rent in a city or rely on the housing-income curvature rather than housing-expenditure curvature. Both these methods lead to higher reported minimum housing levels because it ignores many forms of consumption sharing and self-insurance that households have. While appropriate for many settings not involving risk, this leads to much higher risks in housing consumption than in practice in my setting.

this estimated transition matches the PSID data in Appendix D.4.

6.3 Internally Identified Parameters

The other parameters of the model are identified via indirect inference. Many of these parameters govern the key elasticities affecting the welfare of the program, including the waitlists, quality of rental assistance, and the labor disutility.

6.3.1 Rental Assistance Parameters

The primary moments used to identify the rental assistance parameters rely on the HUD data. I use exit rates to estimate preferences for rental assistance housing, and I use the observed wait times and application rates in the HUD data to estimate the entry parameters.

Rental Assistance Preferences I use the exit rate in response to adjusted income and private market rents as a revealed preference. The key intuition is adjusted income income determines the price of assisted housing while private market rents affect the price of the outside option. I merge the CZ wage and rent prices to the HUD data and estimate exit rates via probit. The auxiliary model is

$$Pr(\text{exit}_{it} \mid \mu_{\ell t}, p_t^h) = \Phi(\beta_0 + \beta_1 \ln \mu_{\ell t} + \beta_2 \ln p_{\ell t}^h + \beta_3 (w_{it} P_{it} + B(X_{it})) + X_{it} \delta)$$
(6.1)

Identification relies on how many households are marginal to exit at a given price. For example, households facing low rent prices are more likely to experience a wage/rent shock that causes them to prefer the private market over rental assistance. Therefore, if the exit rates at given prices match those from the model, then this will pin down the distribution of households whose assistance quality is marginal at that price level. This jointly identifies the Beta distribution parameters.

It may seem intuitive that the price coefficients directly identify willingness to pay, but it is not the correct object for h_Z . In the model, h_Z determines the flow utility of assisted housing quality. On the other hand, the willingness to pay for rental assistance, in terms of prices, is the *present value* of both having rental assistance and receiving rental assistance in the future. Therefore, the willingness to pay implied by the price coefficient will be larger than the actual flow utility from quality h_Z , and the option value may be large enough that some households may be willing to pay more than they receive in quality. For example, knowing that you may not receive rental assistance for another 10 years if you exit can make you value the option of rental assistance as insurance. Therefore, one should not directly interpret the probit coefficients as estimates of h_Z , but instead as informative of households' expected present value of assistance.

One threat to identification is that the model does not capture all potential shocks that cause a house-

Table 8: Average marginal effects from the probit model of exit rates (Equation 6.1), HUD data.

Coefficient	Average Marginal Effect
Rent Price	-0.0555
Adj. Ann. Income	6.14×10^{-7}
Base Rate	3.3%

hold to move. If this were the case, then more households appear marginal to the prices than in reality, biasing the estimates of the value of assistance downward. As robustness, I also estimate a baseline exit rate from the HUD data of those who are most likely inframarginal to any price changes. This corrects for any non-price shocks that are consistent across households. I present the results in Appendix D.5.

Estimated from the HUD data, I present the average marginal effects of the probit model in Table 8. Baseline exit rates are low, averaging only 3.33% per quarter, implying either a high h_Z , high option value of rental assistance, or a mixture of both. Relative to baseline, rent prices and adjusted income have an economically significant impact on exit rates. Using the average marginal effects as a rough proxy, moving from the 50th to 75 percentile in rent prices decreases exit rates by 1.67 percentage points. However, these changes are small in absolute terms, implying strong preferences for remaining in rental assistance. The actual implied quality will later be inferred from the model.

Waitlist and Application For the waitlist parameters, I match baseline conditional wait times between the HUD data and the model. Because I only observe wait times conditional on entry, I miss those who select out of rental assistance because their household circumstances change between application and receipt. Therefore, I correct for this by matching the conditional hazard rate of households in the model to that in the HUD data. Theoretically, correctly specifying how household characteristics change, particularly wages and rents, controls for the unobserved unconditional hazard rates.

To allow for wait times to differ by household characteristics, I estimate a Cox proportional hazards model. This estimates a baseline hazard rate that is multiplicatively scaled by household characteristics. Importantly, this assumes that the hazard rate for any given household is proportional to other households across time. As expected, wait times decrease for those with lower incomes and lower local rents. I present the estimates of the proportional hazards model in Table 9.

I parametrize the application probability—or opening/closing of the waitlists—as a logistic function of the rent price:

$$p(R_{it} \mid p_{\ell t}^{h}) = \frac{1}{1 + \exp\left(-(\zeta_0 + \zeta_1^{p^h} \ln p_{\ell t}^{h})\right)}$$

This captures the primary channel that leads to the closing of the waitlist, excess demand from high rent

Table 9: Waitlist hazards rates from the Cox proportional hazards model, HUD data.

	Wait Time (Years)
	Hazard Ratio
log(CZ Income)	0.834***
log(CZ Rent)	0.499***
log(Adj. Ann. Income)	0.993***
Disabled	0.902***
Voucher	0.618***
Observations	4,571,347

prices, without internally estimating a large number of parameters. However, this comes at the cost of less flexibility if waitlist closings change drastically between rent price states. This may occur if HUD allocates assistance non-uniformly by rent price.

I identify the application probability parameters using the share assisted by rent price quartile. The intuition is that the stock of households in rental assistance is a function of the inflows and outflows. If outflows are correctly matched by the exit rate auxiliary model, and wait times are also correctly specified, then the remaining variation in the stock must come from application probabilities. Therefore, targeting share of households in rental assistance pins down the application probabilities. To allow for the application rates to differ by rent price, I estimate the share assistance by rent quartile from the HUD data³⁰.

6.3.2 Labor Disutility and Unemployment Rate

I identify labor disutility using employment moments from the PSID. The intuition is that if the wage-rent process, employment process, and rental assistance process are specified correctly, the remaining choice of labor supply is only attributable to the labor disuitility parameter, and the share of households employed is indicative of the labor disutility. I leave the non-employment shares in rental assistance untargeted and only target the employment shares of those outside rental assistance.

To identify the unemployment rate, I use differences in labor supply by household age among the non-assisted. As younger households have higher incentives to work, their non-employment rate is more informative about the unemployment rate in the economy than the non-employment rate of older workers. This allows separation between the estimates of labor disutility and the job separation rate.

³⁰It's also important to note that, in simulations, application probabilities are nearly one for households in the lowest rent quartile. Therefore, bottom rent quartile assist share also aids in identifying the quality of rental assistance. The distribution of quality must at least rationalize the number of households in rental assistance in the bottom rent quartile.

Table 10: Internally Estimated Parameters

Parameter	Value
Labor Supply Parameters	
Labor disutility ϕ	-0.629
Unemployment rate λ	0.060
Assisted Quality Shape Paramters	
$lpha_{h_z}$	3.214
$eta_{h_z}^{\bar{z}}$	4.673
mean h_z	1575.034
Application Probability Parameters	
ζ_0 (Const)	-2.109
$\zeta_1^{p^h}$ (Rent price)	-7.635

7 Results

I estimate the current model for male non-disabled household heads, using their specific moments in indirect inference. The model's simulated moments closely match many of the important targeted and nontargeted of the data. The model supports many of the key reduced-form parameter estimates used in Section 4 to estimate the insurance value of the program.

7.1 Internally Estimated Parameters

In Table 10 and Figure 7.1, I present the internally estimated parameters from indirect inference, including parameters influencing labor supply, the quality of assisted housing, and wait times. For labor supply, male household heads require a high labor disutility of -0.629 to rationalize the level of non-employment among older workers, and unemployment is well within the bounds expected. Section, the quality of housing implies that assisted housing is roughly half the quality consumed by the average renter household. The estimated distribution of assisted quality implies a mean of 1575 per quarter, or 1632 conditional on receiving. At the mean rent price, the average renter consumes roughly 3000 in housing quality per quarter. Finally, The application probabilities and baseline hazard rate imply long wait times for assistance, especially in high rent locations. The conditional hazard rate between the model and data closely match, albeit with noise in the far tail³¹.

7.2 Simulated versus Data Moments

I show that the simulated data from the model align closely with important data moments, including household consumption, the labor supply of non-assisted and assisted households, and rental assistance characteristics.

³¹These estimates could be made more precise with a larger number of simulations, but the benefits would likely be minor for the precision of the important parameters governing the insurance value of rental assistance

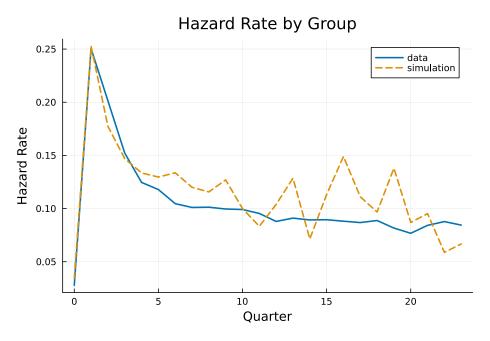


Figure 7.1: Conditional Hazard Rates

Notes: Line plot of simulated versus HUD data base conditional hazard rates implied by the Cox hazard model of wait times. The X axis is wait time (quarters), while the Y axis is the hazard rate.

Household Consumption Household consumption in the model follows housing consumption choices and consumption-savings choices over time. In Figure 7.2 Panel A, the share of expenditures that households devote to housing increase as household expenditures decrease. The simulation data closely matches the actual data, albeit with slightly less curvature in the relationship between housing budget share and total expenditures. This close match is independent of the internally estimated parameters and solely comes from the external estimation of the Stone-Geary parameters. This match between the simulation moments and data matters for the risk that housing poses for low-income households.

Examining total expenditures over time in Figure 7.2 Panel B, households also have a similar consumption profile to the data, with the characteristic humped-shaped consumption coming from the wage distribution. These results support that households make similar consumption-savings decisions as in the data, given their income profile over their lifecycle and the income-rent risk households face³².

Labor Supply The estimated disutility of labor supply and unemployment rate closely matches the targeted non-assisted labor supply of households. In Table 11, I show the labor supply moments, with very close matches between simulated and real data. For non-assisted households, I estimate labor supply moments from the PSID and match these moments via indirect inference. In addition, the implied labor disutility results in similar rent payments for assisted households. This is especially important for matching the

³²The slightly less curvature in consumption is likely due to no mortgage payoff in the model.

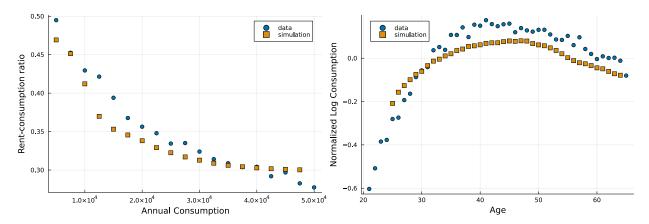


Figure 7.2: Panel A: The share of consumption devoted to rent expenditures, in \$2,500 bins. Panel B: normalized log consumption over the lifecycle by age.

Table 11: Data versus Simulated Moments

Moment	Data	Simulation
Employment: age < 45, non-assisted	0.923	0.925
Employment: age >= 45, non-assisted	0.904	0.911
Assisted rent payment	1120.548	1093.878
Assist Shares by Rent Quartile		
Quartile 1	0.034	0.034
Quartile 2	0.037	0.034
Quartile 3	0.031	0.034
Quartile 4	0.051	0.044
Mean Assist Shares	0.038	0.036

Notes: Employment data moments from the PSID, while all others come from administrative HUD data. Assisted rent payments are the household rent portion determined by their adjusted income. Rent quartiles estimated from constructed rent prices in Census and defined relative to Census populations.

costs of the program and subsidy benefits that households receive.

Rental Assistance Moments In Figure 7.3, I show the key moments that determine the preferences for and availability of rental assistance. First, the mean share assisted match between the data and the simulations, albeit the model underestimates the fraction of households in rental assistance in the highest rent quartile³³. Importantly, the share assisted in the lowest rent quartile match between the data and the simulations. Because these households face almost no wait times, this moment plays a large role in identifying preferences for rental assistance.

Second, the exit rates match between the data and simulations, although with a slightly steeper slope in simulations. These moments pin down the distribution of preferences for rental assistance, determining who is marginal at a given wage-rent price.

³³This occurs because the logistic function for application probabilities struggles to fit both the flat share assisted in the other rent quartiles and the jump in share assisted in the last rent quartile. Note that this also incentivizes the model to reduce the exit rates for high rent prices.

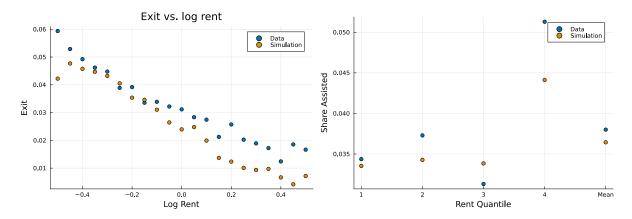


Figure 7.3: Panel A, exit rates from rental assistance by log CZ rent price. Panel B, share assisted by CZ rent quartile.

7.3 Comparison to Reduced Form Results

I use Jacob and Ludwig (2012) as a comparison of the labor supply effects between the structural model and reduced form results in the literature. Jacob and Ludwig (2012) use a lottery for vouchers in Chicago, Illinois in the late 90s to estimate the labor supply effects of voucher receipt. Including both the extensive and intensive margin, they find vouchers decrease earnings by \$385, or 10% of earnings. In my structural model, I find earnings decrease by \$495, or 13.6% of earnings within the structural model.

It's important to note the differences between Jacob and Ludwig (2012) and the structural model. The lottery involved an expansion on top of the existing vouchers already provided by the Chicago Housing Authority. Recipients who receive a voucher regardless of the lottery may have different labor supply effects than those on the margin of the lottery. In addition, this lottery occurred in Chicago during the early years of welfare reform in the late 90s. Welfare reform may have independently affected the labor supply of those in rental assistance during this period relative to my sample period. Finally, the structural model does not incorporate the intensive margin, generating courser-grained labor supply effects than in reality. In fact, Jacob and Ludwig (2012) find households make significant changes in work hours rather than solely the extensive margin, which may lead to more finely-tuned labor supply effects than my model can capture. Regardless of these caveats, the reduced form labor supply effects and structural model largely predict similar earnings effects.

I also compare the housing quality estimates to those in Reeder (1985) and Collinson and Ganong (2018). Using changes in hedonic housing consumption before and after obtaining rental assistance, Reeder (1985) estimates that households discount vouchers at \$0.83 per dollar of subsidy. Collinson and Ganong (2018) use exogenous variation in the Fair Market Rent to estimate that landlords obtain \$0.41-\$0.46 per dollar of subsidy, and the rest is incident on households. Both estimates provide a different view of how

households benefit from the assistance transfers, with Reeder (1985) focused on the household changes in housing consumption while Collinson and Ganong (2018) is more focused on the passthrough. Note that both of these estimates are specific to vouchers, while my estimates of housing quality are the average for both vouchers and public housing.

I find that the implied subsidy is \$0.56 per dollar of actual subsidy. These results are in line with Collinson and Ganong (2018), whose results imply \$0.54-\$0.59 per dollar of subsidy are incident on recipients. Because this is the implied subsidy and not how much households discount the transfer, this estimate is much more comparable to Collinson and Ganong (2018) than to Reeder (1985). These estimates support the existing reduced form literature on the actual subsidies incident on households in the program.

8 Counterfactuals

This insurance value depends on the current design of the program, and there is public debate over proposed reforms that impact the insurance-incentive tradeoff. Each component of the design of rental assistance plays a key role in who receives rental assistance, when they receive it, and how much the program distorts their labor supply and housing consumption.

Below, I implement several counterfactual reforms of the program. These reforms show both how the design influences the incentive-insurance tradeoff, and also how proposed reforms affect the welfare of the program. I demonstrate how converting to cash transfers and changing the rent design to a flat rent affect the incentive-insurance tradeoff.

In order to control of the general equilibrium effects of any counterfactual, I hold the share of house-holds assisted fixed by rent quartile. This prevents counterfactuals from expanding rental assistance and potentially affecting rent prices across markets. In order to hold the share assisted fixed, I optimize the application probabilities (i.e., the opening and closing of the waitlists) until the excess demand is balanced by the wait times. This simplifies the process of holding the share fixed, but it abstracts away from any potential improvements in targeting households through the waitlist. Note that, for any counterfactual, this implies that the reform accounts for the joint effects of the change in design itself plus the effects of the design on the availability of the program. If the change creates more demand for the program, then households may have to wait longer to receive assistance, reducing potential targeting benefits.

To estimate the welfare effects of these various counterfactuals, I estimate the consumption tax in the counterfactual program that would leave households indifferent the counterfactual and the baseline. The reported welfare effects are in terms of this tax. Currently, the consumption tax does not include the increased taxes necessary to fund the program if costs grow.

8.1 Converting to a Cash Transfer

As previously mentioned, a natural reform is to convert the rent subsidies to a cash transfer. A cash transfer improves the ex post value of the subsidies to households by eliminating the deadweight losses of in-kind transfers. However, cash transfers may reduce the targeting benefits of the program, especially if the cash transfer is not indexed to rent prices.

I convert the implied rent subsidies that households receive in the model to cash. This implies that I do not increase the transfers by the portion that the landlord receives, which would greatly improve the welfare of the program to low-income households. Since the model does not account for the welfare of landlords, including landlord transfers in the cash benefits biases the welfare effects upward. In addition, removing the transfers to landlords may have general equilibrium implications for rent prices. In this sense, this counterfactual should be thought of as transferring cash to households through landlords.

In Table 12, I implement four different variations of converting in-kind assistance to cash. First, I simulate a change to a regular cash transfer program with no waitlist, like other entitlement programs. For this counterfactual, I allow the costs to expand to meet the full demand for the program, but I still allow households to opt into the program. Second, I simulate a cash transfer that is either conditional on solely income or conditional on both income and rent. In both cases, I estimate the average subsidy that will keep the spending of the program fixed, conditional either on income or a household joint income and rent. I take interactions between these two policies to create four separate cash assistance policies.

All cash transfer programs improve welfare over in-kind transfers, albeit the entitlement programs greatly expand costs. First, the rent-indexed cash transfer with wait times reduces the housing distortions of rental assistance while maintaining nearly all of the targeting benefits of the program. The compensating variation measure for this counterfactual is 1.2% of lifetime consumption ³⁴. For entitlement cash transfers, the reduced wait times increase costs by more than 100% and only double the welfare benefits of the program relative to wait times. When comparing the income-base cash transfer to rent-indexed transfer, the rent-indexed transfer provides the same welfare benefits but at reduced costs. This comes from better targeting subsidies to households with welfare losses from high rent prices.

No cash welfare reform has a significant change on the labor supply of households in the program. Because each cash welfare reform keeps the implicit income tax, households do not face drastically different incentives to increase or decrease their labor supply. Any change in the labor supply of these reforms comes from changing the targeting of which households receive the assistance.

In summary, changing the program to cash welfare, while keeping the waitlist design, can modestly

 $^{^{34}}$ Note that optimizing application probabilities is not as successful for the income-based wait times. This is due to a large increase in demand for cash assistance in low income CZs, which the model struggles to fit with logistic probabilities.

Table 12: Proportional changes in key parameters when subsidies converted to cash.

Counterfactual	Welfare	costs	Labor Supply	Share Assisted
No wait; income-based	0.035	1.193	-0.009	1.432
No wait; rent scaled	0.035	1.063	-0.004	1.327
Wait; income-based	0.014	0.323	-0.006	0.130
Wait; rent scaled	0.012	-0.069	-0.002	0.006

Notes: Income-based cash assistance only conditions the transfer on income, while rent scaled also conditions on rent. For wait times, number assisted held fixed by adjusting application probabilities in the model. The costs implied by the model use the model's implied subsidy amount, given the identified assistance quality.

improve the welfare of the program. These changes largely keep the targeting benefits while reducing distortions in housing consumption. However, cash welfare may not target certain households not included in the model but aided by rental assistance, such as homeless households. Any reform change must weigh the benefits and costs for these other types of households in the program.

8.2 Income-based versus Flat Rent Subsidies

One of the most common proposals is to change rental assistance from an income-based subsidy to a flat rent subsidy. The rationale is that a flat rent subsidy encourages more work than the status quo implicit labor tax of 30%. However, changing to a flat rent subsidy reduces the insurance value, targeting less benefits to the lowest income households and especially reduces the income insurance of the program.

I set the flat rent payment to keep the average rent payments to HUD fixed by rent price. Therefore, a household in the program at a given rent price will pay the average rent of households in the program under the income-based rent design.

This has several effects on the selection into the program, the insurance value, and the moral hazard costs. First, low-income households face the largest losses from this change since they now face higher rents. On the other hand, high-income households in the program also face lower incentives to exit because their rents do not increase with their income. This reduces the targeting of benefits on income. However, by unpegging rents from incomes, this potentially removes moral hazard in labor supply. Households now face no implicit tax on earnings on the intensive side of rental assistance, albeit they still face income eligibility criteria.

In Table 13, I show the effects of flat rent on labor supply. I find that, contrary to the intent of the reform, changing to a flat rent reduces labor supply. This comes from how a flat rent changes demand for the program and the selection of households that receive the program. Since a flat rent incentivizes higher income households to participate in the program, excess demand increases, drastically reducing application probabilities. Due to the small probability of receiving rental assistance, receipt shifts to the

Table 13: The effect of flat rents on labor supply, application probabilities, and exit rates.

Variable	Value	Pct Change (baseline)
Labor Supply (LS)	0.907	-0.092
LS Assisted	0.750	-1.292
LS Assisted Age < 45	0.770	2.670
LS Assisted Age ≥ 45	0.451	-7.004
Application Prob	0.039	-81.620
Exit Rate	0.006	-80.612

older population and acts as an unexpected wealth shock for recipient households. In this case, the income effect of receiving this assistance dominates the incentives of a flat rent, pushing labor supply down for older households. This leads to small reductions in labor supply in the aggregate, but especially large reductions in labor supply of those older households in the program.

Even without considering the fiscal externalities of the labor supply effects, the flat rent reform reduces the welfare of the program by 12% relative to eliminating the program entirely. This comes from the both reduced targeting of benefits to the lowest income households and the reduced targeting that comes from the decrease in application probabilities. This suggests that while wait times have the potential to screen high marginal utility applicants, excessive wait times due to high excess demand can counteract these screening benefits. Therefore, the income-based rents of the program greatly improve the targeting of the program without the downsides of decreased labor supply.

9 Conclusion

This paper demonstrates that federal rental assistance functions as valuable insurance against joint rent-income risk, providing a new rationale for these programs beyond their established neighborhood effects. The core insight is that rent prices create consumption risk that extends beyond pure income volatility, particularly for low-income households who spend large, relatively inelastic shares of their budgets on housing. Federal rental assistance directly targets this risk through its income-share design, which fixes households' rent burden at 30% of income while subsidizing market rent costs.

My empirical analysis yields several key findings. First, using sufficient statistics methods, I estimate that marginal expansion of rental assistance generates \$1.51 in welfare benefits per dollar of government cost, relative to previous literature estimates of \$0.66. This suggests policy should expand rather than contract these programs. Second, decomposing the insurance value reveals that the vast majority comes from protecting income in high-rent locations rather than insuring rent price volatility directly. This targeting occurs because rent constraints bind more tightly when incomes fall, making income insurance more valuable in expensive housing markets. However, households on fixed incomes benefit from the rent insurance

component since their income is not adjusted with local rents. Third, the program's waitlist rationing mechanism effectively selects households with persistent rather than temporary shocks, enhancing the insurance value by concentrating benefits where they are most needed.

The lifecycle model provides additional insights into program design. Counterfactual analysis shows that converting to flat rent subsidies significantly reduces the welfare benefits of the program with even small negative effects on labor supply. This comes from reducing targeting of benefits to the lowest income households while also increasing the benefits of the program to those who are marginal to employment. On the other hand, changing the program to cash welfare modestly improves the welfare benefits of the program with little change in labor supply.

There are several directions for future research. First, this research abstracts away from the equilibrium effects of housing markets in general. These effects can have crucial impacts on the optimal federal rental assistance policies, particularly if federal rental assistance affects market rent prices. Second, there is no evidence on how households perceive their rent risk. Households may either be myopic to their future rent risks, or even have incorrect beliefs about how rent prices will evolve. Additionally, households may not internalize the risks that local rent prices create when they experience income shocks, exacerbating these income shocks because they do not self-insure.

References

Abramson, Boaz and Stijn Van Nieuwerburgh. 2024. "Rent Guarantee Insurance." SSRN Electronic Journal URL https://www.ssrn.com/abstract=4747216.

Acosta, Sonya and Erik Gartland. 2021. "Families Wait Years for Hous-Vouchers Inadequate Funding Center on Budget and Policy Priing Due to orities." Tech. rep. URL https://www.cbpp.org/research/housing/ families-wait-years-for-housing-vouchers-due-to-inadequate-funding.

Ashenfelter, Orley and Alan Krueger. 1994. "Estimates of the Economic Return to Schooling from a New Sample of Twins." *The American Economic Review* 84 (5):1157–1173. URL https://www.jstor.org/stable/2117766. Publisher: American Economic Association.

Autor, David, David Dorn, and Gordon Hanson. 2019. "When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men." *American Economic Review: Insights* 1 (2):161–178. URL https://www.aeaweb.org/articles?id=10.1257/aeri.20180010.

Autor, David, Andreas Kostøl, Magne Mogstad, and Bradley Setzler. 2019. "Disability Benefits, Consumption Insurance, and Household Labor Supply." *American Economic Review* 109 (7):2613–2654. URL https://pubs.aeaweb.org/doi/10.1257/aer.20151231.

Autor, David H. and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5):1553–1597. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.5.1553.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer. 2024. "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice." *American Economic Review* 114 (5):1281–1337. URL https://www.aeaweb.org/articles?id=10.1257/aer.20200407.

Bhargava, Saurabh and Dayanand Manoli. 2015. "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11):3489–3529. URL https://www.aeaweb.org/articles?id=10.1257/aer.20121493.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *The Review of Economic Studies* 91 (6):3253–3285. URL https://doi.org/10.1093/restud/rdae007.

- Brumm, Johannes and Simon Scheidegger. 2017. "Using Adaptive Sparse Grids to Solve High-Dimensional Dynamic Models." *Econometrica* 85 (5):1575–1612. URL https://onlinelibrary.wiley.com/doi/10.3982/ECTA12216. Publisher: John Wiley & Sons, Ltd.
- Carlson, Deven, Robert Haveman, Thomas Kaplan, and Barbara Wolfe. 2012. "Long-term effects of public low-income housing vouchers on neighborhood quality and household composition." *Journal of Housing Economics* 21 (2):101–120. URL https://www.sciencedirect.com/science/article/pii/s1051137712000253.
- Chetty, Raj. 2006. "A general formula for the optimal level of social insurance." *Journal of Public Economics* 90 (10):1879–1901. URL https://www.sciencedirect.com/science/article/pii/s0047272706000223.
- ———. 2008. "Moral Hazard versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy* 116 (2):173–234. URL https://www.journals.uchicago.edu/doi/full/10.1086/588585. Publisher: The University of Chicago Press.
- Chetty, Raj and Amy Finkelstein. 2013. "Chapter 3 Social Insurance: Connecting Theory to Data." In *Handbook of Public Economics, handbook of public economics, vol.* 5, vol. 5, edited by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez. Elsevier, 111–193. URL https://www.sciencedirect.com/science/article/pii/B9780444537591000030.
- Chetty, Raj, Nathaniel Hendren, and Lawrence 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. URL https://www.aeaweb.org/articles?id=10.1257/aer.20150572.
- Chyn, Eric. 2018. "Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children." *American Economic Review* 108 (10):3028–3056. URL https://pubs.aeaweb.org/doi/10.1257/aer.20161352.
- Collinson, Robert and Peter Ganong. 2018. "How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality?" *American Economic Journal: Economic Policy* 10 (2):62–89. URL https://pubs.aeaweb.org/doi/10.1257/pol.20150176.
- Cook, Cody, Pearl Z Li, and Ariel J Binder. 2025. "Where to Build Affordable Housing? Evaluating the Tradeoffs of Location.".

- Deshpande, Manasi and Lee M. Lockwood. 2022. "Beyond Health: Nonhealth Risk and the Value of Disability Insurance." *Econometrica* 90 (4):1781–1810. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA19668. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA19668.
- Desmond, Matthew. 2015. "Unaffordable America: Poverty, housing, and eviction." In *The Affordable Housing Reader*. London: Routledge, 2 ed., 389–395. URL https://www.taylorfrancis.com/books/9780429299377/chapters/10.4324/9780429299377-34.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127 (6):2836–2874. URL https://www.journals.uchicago.edu/doi/full/10.1086/702238. Publisher: The University of Chicago Press.
- Finkelstein, Amy and Matthew J Notowidigdo. 2019. "Take-Up and Targeting: Experimental Evidence from SNAP*." *The Quarterly Journal of Economics* 134 (3):1505–1556. URL https://doi.org/10.1093/qje/qjz013.
- Gadenne, Lucie, Samuel Norris, Monica Singhal, and Sandip Sukhtankar. 2021. "In-Kind Transfers as Insurance." URL https://www.nber.org/papers/w28507.
- Ganong, Peter and Pascal Noel. 2019. "Consumer Spending during Unemployment: Positive and Normative Implications." *American Economic Review* 109 (7):2383–2424. URL https://pubs.aeaweb.org/doi/10.1257/aer.20170537.
- Griebel, M. 1998. "Adaptive sparse grid multilevel methods for elliptic PDEs based on nite di erences." Computing 61:151–179.
- Hendren, Nathaniel and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics* 135 (3):1209–1318. URL https://doi.org/10.1093/qje/qjaa006.
- Jacob, Brian A. and Jens Ludwig. 2012. "The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery." *American Economic Review* 102 (1):272–304. URL https://www.aeaweb.org/articles?id=10.1257/aer.102.1.272.
- Kueng, Lorenz, Lee M. Lockwood, Pinchuan Ong, and Scott R Baker. 2023. "Renting Insures Wage Risk."

 URL https://www.dropbox.com/s/9ss92aw8aax89rk/RTCpublic.pdf?e=1&dl=0.

Landais, Camille and Johannes Spinnewijn. 2021. "Value of Unemployment Insurance | The Review of Economic Studies | Oxford Academic." URL https://academic.oup.com/restud/article/88/6/3041/6372704.

Lieber, Ethan M. J. and Lee M. Lockwood. 2019. "Targeting with In-Kind Transfers: Evidence from Medical Home Care." *American Economic Review* 109 (4):1461–1485. URL https://www.aeaweb.org/articles?id=10.1257/aer.20180325.

Lockwood, Lee M. 2025. "Health Insurance and Consumption Risk.".

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. "Wage Risk and Employment Risk over the Life Cycle." *American Economic Review* 100 (4):1432–1467. URL https://www.aeaweb.org/articles?id=10.1257/aer.100.4.1432.

Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2022. "Marriage, Labor Supply and the Dynamics of the Social Safety Net."

Low, Hamish and Luigi Pistaferri. 2015. "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off." *American Economic Review* 105 (10):2986–3029. URL https://pubs.aeaweb.org/doi/10.1257/aer.20110108.

Ma, Xiang and Nicholas Zabaras. 2009. "An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations." *Journal of Computational Physics* 228 (8):3084–3113. URL https://www.sciencedirect.com/science/article/pii/S002199910900028X.

Meyer, Bruce D. and Nikolas Mittag. 2019. "Using Linked Survey and Administrative Data to Better Measure Income: Implications for Poverty, Program Effectiveness, and Holes in the Safety Net." *American Economic Journal: Applied Economics* 11 (2):176–204. URL https://www.aeaweb.org/articles?id=10.1257/app.20170478.

Monkovic, Mia and Ben Ward. 2025. "Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2023." Tech. Rep. SNAP-23-CHAR, U.S. Department of Agriculture, Food and Nutrition Service, Evidence, Analysis, and Regulatory Affairs Office, Alexandria, VA. URL https://fns-prod.azureedge.us/sites/default/files/resource-files/snap-FY23-Characteristics-Report.pdf.

Newman, Sandra J. and Ann B. Schnare. 1997. ""... And a suitable living environment": The failure of housing programs to deliver on neighborhood quality." Housing Policy Debate 8 (4):703–

- 741. URL https://doi.org/10.1080/10511482.1997.9521275. Publisher: Routledge _eprint: https://doi.org/10.1080/10511482.1997.9521275.
- Olsen, Edgar O. 2002. "The Cost-Effectiveness of Alternative Methods of Delivering Housing Subsidies." SSRN Electronic Journal URL http://www.ssrn.com/abstract=296785.
- Powell, M J D. 2009. "The BOBYQA algorithm for bound constrained optimization without derivatives." .
- Rafkin, Charlie, Adam Solomon, and Evan Soltas. 2025. "Self-Targeting in US Transfer Programs." .
- Reeder, William. 1985. "The Benefits and Costs of the Section 8 Existing Housing Program." *Journal of Public Economics* 26 (3):349–377. URL https://www.sciencedirect.com/science/article/pii/0047272785900143?via%3Dihub.
- Rosenbaum, Dorothy, Daniel Tenny, and Sam Elkin. 2002. "The Food Stamp Shelter Deduction: Helping Households with High Housing Burdens Meet Their Food Needs." Tech. rep., Center on Budget and Policy Priorities.
- Sieg, Holger and Chamna Yoon. 2020. "Waiting for affordable housing in New York City." *Quantitative Economics* 11 (1):277–313. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/QE1160. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/QE1160.
- Stoyanov, Miroslav. 2019. "User Manual: TASMANIAN Sparse Grids." Tech ORNL/TM-2015/596, Oak Ridge National Laboratory. URL https://mkstoyanov.github.io/tasmanian_aux_files/docs/TasmanianMathManual.pdf.
- Turnovsky, Stephen J., Haim Shalit, and Andrew Schmitz. 1980. "Consumer's Surplus, Price Instability, and Consumer Welfare." *Econometrica* 48 (1):135–152. URL https://www.jstor.org/stable/1912022. Publisher: [Wiley, Econometric Society].
- Waldinger, Daniel. 2021. "Targeting In-Kind Transfers through Market Design: A Revealed Preference Analysis of Public Housing Allocation." *American Economic Review* 111 (8):2660–2696. URL https://pubs.aeaweb.org/doi/10.1257/aer.20190516.

A Appendix

A.1 Waitlist Preferences

Housing authorities have wide discretion over waitlists and the priorities they set for selecting recipients from the waitlist. Many housing authorities assign points based on household characteristics that deter-

Waitlist Preference

The Authority will give waitlist preferences as follows:

<u>RAD</u>

- Veteran/Surviving Spouse of Veteran (+1)
- · Families with a Right to Return (16 points)
- · RAD Emergency Referral (15 points)
- Residents of Residential Care Facilities for the Chronically III (RCFCI)/ Transitional Residential Care Facilities (TRCF) (14 points) c
- · Mixed Families currently residing at non-RAD units at HOPE SF sites (13 points)
- Involuntarily Displaced with Residential Certificate of Preference (COP) (11 points)
- Department of Homeless and Supportive Housing Referral (HSH) (7 points)
- · Families with minor children living in SROs with a referral from DBI (7 points)
- Involuntarily Displaced from San Francisco residence (5 points)

Figure A.1: San Francisco Housing Authority waitlist preferences.

mine priority, often favoring homeless households, veterans, or even those with jobs.

In Figure A.1, I present the San Francisco Housing Authority's waitlist preferences from their website. The San Francisco Housing Authority often chooses households who are involuntarily displaced, have experienced homelessness, or have participated in other housing government programs. Higher points lead to shorter waitlists for housing.

B Data Construction

B.1 Location Wage-Rent Price Construction

I construct wage-rent prices similar to Kueng et al. (2023), who use the mean quality-adjusted wages and rents in a given commuting zone (CZ). The advantage of using commuting zones is that the boundaries are consistent over time, allowing for a clean panel to estimate the wage-rent distribution. This comes at the cost of measurement error in the true location wages and rents that households face within a CZ. The resulting distribution is shown in Figures B.1 and B.2. Below I provide a summary of the construction of these variables.

For brevity, I only summarize the wage-rent construction and refer the reader to Kueng et al. (2023) for a more detailed description of the construction of CZ wage-rent prices. I explicitly state any deviations that I take from their method.

The goal is to use Census data to construct an aggregate measure of CZ wages and rents. While Kueng et al. (2023) use the years 1940 to 2010, I shift the sample to 1970 to 2020. This concentrates the wage-rent estimation on years where I have data either from the PSID or from HUD. Where possible, decennial census data are used, substituting for the ACS when not available. First, I assign households to CZs using existing

crosswalks of Census locations to CZs (Autor and Dorn, 2013; Autor, Dorn, and Hanson, 2019). Because some Census locations cross CZ boundaries, these crosswalks assign each household a probabilistic weight of being in each CZ that their location spans. I combine these probabilistic weights with the existing Census weights to create the full weight for each household-CZ-year.

I now turn to rent estimation. The chosen measure of rent prices is the mean quality-adjusted rent price in a given CZ, with the rent price normalized to one. This measure attributes any rent prices differences not attributable to the houing quality variables in the Census as true differences in the rent price $p_{i\ell t}^h$. These differences in rent prices could be due to demand for the location and/or differences in housing supply elasticities. In the future, I will test for whether these rent prices are correlated with these two underlying structural causes, or whether my measure is capturing unobserved housing quality.

Households are restricted to those with less than college education³⁵ who rent a typical unit in private rental markets. This drops rental properties with large acreages, that are used for commercial use, or have amenities like food or labor payments included in the rents. In addition, some unit types like condos are removed to ensure consistency across years. This creates a sample that represents what a typical renter that would use rental assistance may rent on the private market.

To estimate the aggregate CZ log rent price by year, log household rents are regressed on hedonic characteristics of the unit $X_{i\ell t}$ and CZ fixed effects $\delta_{t,CZ}$:

$$\ln p_{i\ell t}^h = \delta_0 + X_{i\ell t} \beta_t + \delta_{t,CZ} + \varepsilon_{i\ell t}$$

The resulting CZ fixed effects $\delta_{t,CZ}$ represent quality-adjusted mean log rent prices for that CZ-year, normalized relative to the other CZs in that year. Included housing characteristics are indicators for the year built, indicators for the number of rooms, the unit structure, whether the unit has plumbing, and whether the unit has a kitchen–the main consistently provided housing characteristics in Census data. The resulting distribution of CZ rents (weighted by population) is presented in Figure B.1 Panel A.

Using these normalized rent prices, where the mean log CZ rent is zero, does not change the results if I were to add the mean back in. Since housing consumption is also unobserved, the units of housing consumption are inferred from the rent price. Adding the mean rent back only shifts the units of housing consumption, and households still make the same housing consumption decisions relative to the rent price.

Wage estimation follows a very similar process. Households are restricted to those working full time and non-college graduates in order to have an accurate measure of wages. Then, CZ log wages are regressed on household characteristics and CZ fixed effects. These household characteristics include the head's age

 $^{^{35}}$ Kueng et al. (2023) do not make this restriction as their goal is to estimate the full wage-rent-housing price distribution.

and education. Figure B.1 Panel B presents the distribution of CZ wages. For robustness and heterogeneity, I add to Kueng et al. (2023) by also estimating separate CZ wages by education group.

Figure B.2 shows the resulting relationship between CZ wages and rents for a given CZ-year. There is a very strong correlation between these two measures of 0.72. As expected, the largest CZs generally have the highest prices in both wages and rents.

For the purposes of the lifecycle model, I ignore aggregate variation in log wage/rent prices over time. The tradeoff is that while I ignore some meaningful variation that rental assistance insures, this reduces the state space of the problem to focus more richness on important state variables for rental assistance. However, for the sufficient statistics approach, aggregate effects can still impact household consumption and generate insurance value.

One concern is that the CZ wage-rent distribution does not adequately capture the true variation that households face at the local level. This is particularly important for the program features that depend on metro wages and rents, such as the Fair Market Rent and the Area Median Income. In Figure B.3, I show how my constructed CZ wage-rent distribution is correlated with the FMR and AMI. At the individual level, I find a correlation of 0.71 between CZ rent prices and the FMR, which increases to 0.93 when aggregated to the CZ-year level. This implies that when aggregated, the CZ rent price largely matches the constructed FMR, but there is a moderate degree of local variation not captured by CZ rent prices.

For wages, I compare the CZ wage price to the AMI. Note that this is comparing a mean wage of non-college household heads to the median household income across all households, so this comparison is less applicable. I find a correlation of 0.64 for CZ wages and the AMI at the individual level, which increases to 0.85 at the CZ level.

Note that even if the CZ rent and wage prices perfectly captured local prices, the correlation would not be perfect. The FMR and AMI are themselves constructed prices, sometimes relying on specified floors to ensure a minimum of benefits. Even so, the aggregate prices closely follow each other at the CZ level.

B.2 Data Sample Construction

Below, I detail the sample construction of both the PSID and HUD data.

B.2.1 PSID Sample Construction

First, the PSID sample variables must be constructed and standardized across the sample years. To maintain consistency with the literature, I follow the cleaning procedures of the disability literature, primarily using Deshpande and Lockwood (2022) as the main source. This involves construction of income, asset, employment, household expenditures, and measures of disability that are used throughout the literature.

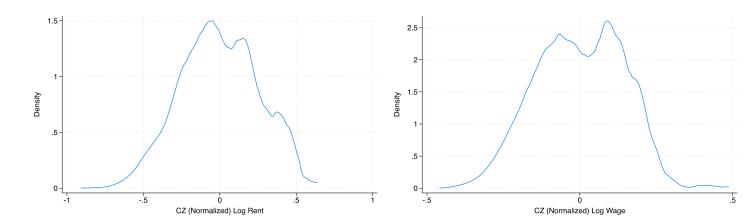


Figure B.1: Kernel densities of constructed CZ log rents and log wages, weighted by population.

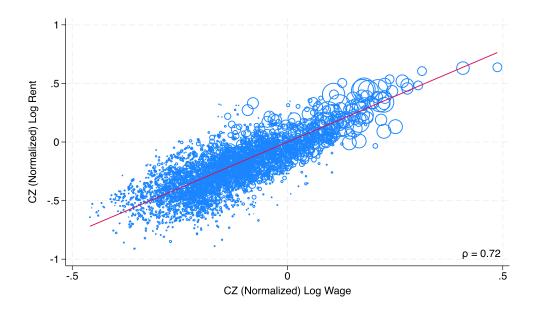


Figure B.2: Relationship between constructed CZ log rents and log wages, weighted by population. Correlation coefficient of 0.72.

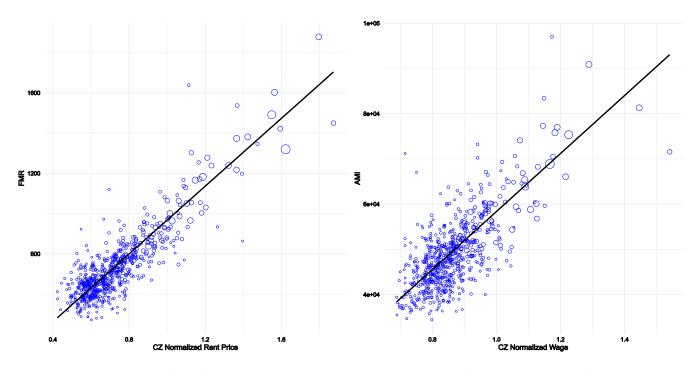


Figure B.3: Relationship between HUD constructed prices and CZ constructed prices following Kueng et al. (2023).

For further details, see Deshpande and Lockwood (2022).

The most important additional construction is rental assistance receipt. The PSID consistently began asking households about government assisted housing in the late 1980s. In the main survey, the PSID breaks this down into whether the household received either public housing or some other form of government rent subsidy.

The PSID restricted sample also links address records of known rental assistance properties to the PSID data through 2009. These records provide additional detail on the type of subsidy that the unit is. The process for this linkage is detailed in Newman and Schnare (1997).

B.2.2 HUD Sample Construction and Description

For the baseline results, I limit the sample to 2012-2022. These years have consistent data construction that reduces potential mismeasurement of key variables.

I employ several restrictions to ensure a consistent sample. First, I drop any households in U.S. territories, such as Guam and Puerto Rico, to ensure consistency with the PSID data. Second, I drop any households in programs that have sufficiently different designs from typical public housing and voucher programs. This includes "incentive" programs such as Family Self-Sufficiency (FSS) and Welfare-to-Work housing authorities that can change the implicit rent payments of rental assistance. I also eliminate some

smaller programs, such as the Moderate Rehabilitation and Coordinated Entry, that are small enough to not play a large role in the rental assistance landscape. This limits the sample to households in either public housing or a voucher program, such as Housing Choice Vouchers (HCVs).

I drop any households that are only partially eligible for the program. Partial eligibility can occur if specific members of the household do not eligibility requirements. This can relate to the legal status of some members of the household. In this case, the household receives prorated subsidies, which complicates the incentives of these households.

I limit the analysis to households who actually enter the rental assistance program. There is a sizeable portion of households that receive a voucher but do not lease up before the expiration date for finding an apartment. I consider these households as never having received a voucher in the first place.

To maintain consistency with the Census and PSID samples, I aggregate household location to the CZ level. I do so by merging the household's reported county to the crosswalks provided by Autor and Dorn (2013); Autor, Dorn, and Hanson (2019).

C Conceptual Framework Details

C.1 Further Rent Price Insurance Decomposition

Theoretically, income-based cash transfers also can insure high rent price states if income covaries negatively with rent prices. To better capture how rental assistance insures rent price volatility differently, we can decompose the second term again on income (omitting the receipt term for brevity):

$$Cov (E [v_y \mid p], E [ph - \tau y \mid p]) = E_y [Cov (E [v_y \mid p], E [ph - \tau y \mid p] \mid y)]$$

$$+ Cov_y (E_p [E [v_y \mid p] \mid y], E_p [E [ph - \tau y \mid p] \mid y])$$

The first term is the average within-income value of insuring rent price volatility. For income-based cash transfers, this term is zero since it does not vary conditional on y. For rental assistance, those in higher rent locations receive more subsidies conditional on income. The value is whether, conditional on income, those in high rent locations are worse off than those in low rent locations. The second term is the between-income value of insurance rent price volatility. This captures whether those with lower income receive higher subsidies. For both income-based cash welfare and rental assistance, subsidies increase as income falls. Whether one is more valuable insurance than the other depends on how closely the increase in subsidies follows the increase in marginal utility as income falls.

C.2 Conceptual Framework Case Study Construction

C.2.1 Rent Price Risk of Elderly Renters

The goal is to approximate the rent price risk that an elderly household on social security may face. To do so, I examine elderly households with similar rent prices in an initial year and examine the distribution of rent prices after 15 years.

Using the one-year ACS samples from 2004-2019, I place several restrictions to focus on elderly renter households on social security. I first select households whose head is between the ages of 65 to 85, receives positive social security benefits, and is out of the labor force. These elderly households must additionally be renters whose housing characteristics satisfy the sample restrictions of B.1 and similarly assign elderly households to CZs. To reduce noise in measuring the ratio of rent to income, I restrict to households who declare an annual household income of at least \$5,000, which is well below the mean social security benefits. Finally, to examine renters facing comparable rent prices, I then select elderly renters who reside in CZs in the middle two quartiles of the rent distribution. These restrictions result in 1.94 million household-year-CZ observations³⁶.

To demonstrate the rent price risk that these households face, I group CZs based on changes in rent prices over 15 years. I estimate rent price growth in a CZ over 15 years and divide CZs into quintiles. I then compare the top quintile CZs in rent price changes to the bottom quintile CZs in rent price changes. Assuming that CZs near the median rent price face the same rent price risk distribution, this approximates the 20% tails of the rent price risk distribution over time.

I then estimate rent expenditures and rent expenditure ratios using gross rent and household income. Gross rent maintains consistent rent prices across households since some household rent may include utilities. Household income includes not only social security income, but any other income sources that these households may receive. Therefore, households may still have non-zero income correlation if others in the household work.

C.2.2 Conditional Income Risk of Unemployment

This case study aims to highlight how rent expenditures differentially adjust in high and low rent locations. To do so, I compare rent expenditures relative to income in the cross section of those employed and unemployed in each CZ rent quintile.

Using the one-year ACS samples from 2004-2019, I restrict to working-age household heads (ages 25-

³⁶Note that the CZ assignment can duplicate households if the CZ they reside in is uncertain, but that these households are reweighted according to the likelihood of being in that CZ.

61) in the labor force who reside in units that satisfy the sample restrictions of Section C.2.1, similarly assigning these households to CZs. Similar to Section C.2.1, I restrict to households whose total income is above \$5,000, which is more restrictive for unemployment than for elderly households. This eliminates around 18% of unemployed households, of which 7% have no household income whatsover and hence an undefined rent expenditure share. However, this prevents very low income households, whose rent expenditure ratios can be an order of magnitude larger and noisy, from dominating the estimation of rent expenditure ratios. The resulting sample is 55.5 million household-year-CZ observations, with the same caveat on duplicated households as in Section C.2.1.

I estimate the mean rent expenditure ratio by CZ rent quintile-employment status. Because this is done using the ACS, this is a cross-sectional comparison of the potential changes in rent expenditure ratio between employment and unemployment. In order for this to be the case, the key assumption is that employment is an idiosyncratic shock that all households in a CZ rent quintile similarly face.

C.2.3 Shock Persistence and Receipt of Rental Assistance

The goal is to examine how the persistence of income shocks affects the probability of entering rental assistance. I estimate an event study of rental assistance receipt by either an disability shock or unemployment shock in the PSID.

To define a disabled household, I follow the same process as Low and Pistaferri (2015), that has become standard in the disability literature using the PSID. I rely on household self-reports of unemployment to PSID among those in the labor force. I drop odd years to ensure a consistent panel because the PSID switched from annual to biennial in the middle of my panel. This creates a sample of 35,088 household-year observations for the disability event study and 31,403 observations for the unemployment event study.

Using the event study imputation method of Borusyak, Jaravel, and Spiess (2024), I implement the following event study design:

$$Z_{it} = \alpha_i + \gamma_t + \sum_{t=-10}^{10} D_{it} \tau_{it} + \varepsilon_{it}$$

where α_i are household fixed effects, γ_t are time fixed effects, D_{it} are indicators for whether the household is experiencing a given income shock in that year (i.e., disability or unemployment). For each household, I select the first shock that they experience and define all time periods relative to that shock. Standard errors are clustered at the household level. The event study results for both the disability and unemployment shocks are presented in Table X.

C.3 Proof of MRS Derivation

For income states *u* and *e*, the marginal rate of substitution is

$$MRS = \frac{v_y^u}{v_y^e}$$

Taking the derivative with respect to p, we have

$$\frac{\partial MRS}{\partial p} = \frac{v_{yp}^{u}v_{y}^{e} - v_{y}^{u}v_{yp}^{e}}{\left(v_{y}^{e}\right)^{2}}$$
$$= \frac{v_{yp}^{u} - v_{yp}^{e} \cdot MRS}{v_{y}^{e}}$$

Next, assume that utility is CRRA, with $\gamma = -y \frac{v_{yy}}{v_y}$. Also, note an alternative expression for v_{yp} from the main text:

$$v_{yp} = -v_{yy}h^* - v_y \frac{\partial h^*}{\partial y}$$

Plugging these into the derivative of MRS and rearranging, we obtain

$$\frac{\partial MRS}{\partial p} = MRS \left[\gamma \left(\frac{h^u}{y^u} - \frac{h^e}{y^e} \right) - \left(\frac{\partial h^u}{\partial y} - \frac{\partial h^e}{\partial y} \right) \right]$$

A further transformation into elasticities yields the result in the main text.

C.4 Imputation of Rent Subsidies

Since the PSID and HUD data cannot be merged, I must impute the subsidy amount that PSID households receive from rental assistance. To do so, I predict household subsidies using similar measured observables between PSID households and the HUD data, focusing on those inputs that are key to rental assitance. In the HUD data, I estimate the following model of rent subsidies

$$B_{it} = \alpha + \beta p_{\ell t}^h + \gamma \text{TTP} + \delta_{\text{hhsize}} + \varepsilon_{it}$$

where B_{it} is the reported subsidy payment to the landlord, TTP is the total tenant payment (i.e., the out-of-pocket rent in levels), and δ_{hhsize} represents dummies for household size. This model captures the important determinants of the rent subsidy, especially the private market rent facing the household and how much the household pays of that rent. The results of this regression are presented in Table 14.

	Dependent variable:
	subs_amt
CZ Rent	0.071***
	(0.00001)
TTP	-0.856***
	(0.0002)
HH Size 2	139.999***
	(0.119)
HH Size 3	273.319***
	(0.130)
HH Size 4	369.714***
	(0.150)
HH Size 5	469.776***
	(0.195)
HH Size 6	534.723***
	(0.287)
HH Size 7	637.606***
	(0.441)
HH Size 8	704.383***
	(0.666)
HH Size 9	806.012***
	(0.955)
HH Size 10	921.011***
	(0.985)
Constant	-85.998***
	(0.156)
Observations	25,540,311
R ²	0.679
Note:	*p<0.1; **p<0.05; ***p<0.

Table 14: Subsidy Imputation Regression, HUD 2012-2022.

62

C.5 Robustness of Sufficient Statistics Results

I present robustness of the sufficient statistic estimates to the various assumptions provided.

C.5.1 Coefficient of Relative Risk Aversion

The coefficient of relative risk aversion only affects the insurance value of the program, and subsequently the MVPF through insurance value. In Table 15, I present the sensitivity of the insurance value to the risk aversion parameter γ . The MVPF remains greater than 1 except in the extreme case of $\gamma=0.5$. If comparing these insurance value results to other papers, it's useful to benchmark using $\gamma=2$, given that the vast majority use this magnitude of risk aversion.

Risk Aversion (γ)	EAWTP	Ins Val	MVPF
0.5	340.30	82.03	0.97
1.0	429.46	171.18	1.24
1.5	518.33	260.06	1.51
2.0	600.31	342.04	1.77
2.5	671.49	413.21	1.99
3.0	730.79	472.51	2.18
3.5	779.01	520.74	2.34

Table 15: Sensitivity of insurance value and MVPF to risk aversion. Baseline results in main text are $\gamma = 1.5$. Baseline net government costs are \$349.68.

C.5.2 Normalization of Marginal Utility

An alternative normalization of marginal utility is to normalize across all states rather than the states where the household is not receiving rental assistance. Theoretically, this is a tax that a household must pay in *every* state of the world rather than just states of the world when they do not receive. If rental assistance does target high marginal utility states, then an all-states normalization will lower the insurance value of the program.

In Table 16, I present the alternative normalization results. There is a 13% decrease in the insurance value relative to the baseline results.

Risk Type	Prob Assist	EAWTP	Ins Val	Gov Cost	MVPF
Less Than HS	0.12	785.30	310.49	611.93	1.28
HS Grad	0.06	394.00	200.90	270.76	1.46
Average	0.07	484.52	226.25	349.68	1.42

Table 16: Sensitivity of Results to normalization of marginal utility across all states. Compare with Table 4, which normalizes marginal utility relative to non-assisted states.

C.5.3 The Deadweight Loss of Landlord Transfers

The transfers to landlords may not be valued one-to-one with government revenues. I examine the sensitivity of the results to this assumption by varying how much of the landlord transfer is a deadweight loss and hence a fiscal externality. To do so, I interpolate between the net and gross government costs, including a share of the landlord transfer as a government cost.

I present the results of this exercise in Table 17. The government costs exceed the ex ante benefits only if all of the landlord transfer is counted as a cost. However, even small deadweight losses have moderate impacts on the MVPF.

Landlord DWL	Gov Cost	MVPF
0.0	349.68	1.51
0.2	389.42	1.36
0.4	429.17	1.23
0.6	468.91	1.13
0.8	508.66	1.04
1.0	548.40	0.96

Table 17: Sensitivity of government costs to the portion of landlord transfers counted as a government cost. Baseline EAWTP is \$518.33.

C.5.4 Sensitivity of Insurance Value Decomposition

To be completed later...

D Model Details

D.1 Preferences

While labor supply affects the choice over total consumption (i.e., both numeraire and housing consumption). Once again, I use the Stone-Geary utility function

$$U_{i\ell t}(c,h,P) = \frac{\left(c^{\alpha}(h - \bar{h}_{i\ell t})^{1-\alpha} \exp(\phi P)\right)^{1-\gamma}}{1-\gamma}$$

Similar to Low and Pistaferri (2015), the complementarity or substitutability between consumption and working depends on the sign of ϕ . They show that Frisch complementarity is related to the sign of ϕ :

$$sign\left(\frac{dc}{dP}\right) = sign\left(U_{cP}\right) = -sign\left(\phi\right)$$

Therefore, numeraire consumption and labor supply are Frisch complements because I assume $\phi < 0$. Simlarly, housing consumption and labor supply are Frisch complements.

The division of consumption between numeraire and housing is independent of labor supply. At an interior optimum, $U_c = U_h/p$. Therefore,

$$\alpha c^{\alpha(1-\gamma)-1} \left(h - \bar{h}\right)^{(1-\alpha)(1-\gamma)} \exp\left(\phi \left(1 - \gamma\right) P\right) = \left(1 - \alpha\right) c^{\alpha(1-\gamma)} \left(h - \bar{h}\right)^{(1-\alpha)(1-\gamma)-1} \exp\left(\phi \left(1 - \gamma\right) P\right) / p$$

$$\frac{c}{h - \bar{h}} = \frac{p\alpha}{1 - \alpha}$$

which does not depend on the labor supply. Intuitively, households have a commodity consumption subutility, where the choice of total expenditures depends on working but preferences over which commodity are independent of working.

The minimum housing constraint \bar{h} plays an important role in the marginal utility of consumption. To see this,

$$U_{c\bar{h}} = -\alpha \left(1 - \alpha\right) \left(1 - \gamma\right) c^{\alpha(1 - \gamma) - 1} \left(h - \bar{h}\right)^{(1 - \alpha)(1 - \gamma) - 1} \exp\left(\phi \left(1 - \gamma\right) P\right)$$

For $\gamma > 1$, this expression is positive, and raising the minimum housing consumption increases the marginal utility of consumption. Given CRRA preferences, this also raises the marginal utility of income in low consumption states relative to high consumption states. Hence, demand for insurance is higher when \bar{h} is higher.

D.1.1 Marginal Utility Expressions with Stone Geary Preferences

[To be completed later...]

D.2 Wage-Rent Process Estimation Details

D.2.1 Location Wage-Rent Process

To estimate the location wage-rent process, I follow the same process as Kueng et al. (2023), which involves estimating a panel autoregressive process. I first estimate CZ wages and rents following Appendix B.1. To maintain a consistent time period in the panel, I restrict to decadal years between 1970 and 2020. The estimating equation is

$$w_{\ell t} = \alpha w_{\ell t-1} + \varepsilon_{\ell t}$$

$$p_{\ell t}^h = \beta p_{\ell t-1}^h + \eta_{\ell t}$$

where ℓ is the location (i.e., CZ), t is the decade, $w_{\ell t}$ is location log wages, $p_{\ell t}^h$ are location rent prices, and $(\varepsilon_{\ell t}, \eta_{\ell t} \sim N(0, \Sigma))$ is a shock error term. I also test robustness to this specification by adding the cross prices as covariates to each autoregressive process.

To estimate this process, I regress each price on its lagged decadal price, collecting the autoregressive coefficients and residuals from the regression. I estimate Σ using the variance-covariance matrix of the residuals.

As demonstrated in Kueng et al. (2023), this process can be reduced to smaller time steps using the eigenvectors and eigenvalues of the autoregressive coefficients. In Table 19, I present the estimation of both the multivariate and univariate wage-rent processes. The univariate processes capture almost all of the variation that is found in the multivariate process, and therefore I use the univariate processes in my baseline estimation. I find that while the location wage process is stationary, the location rent process is non-stationary. However, I cannot reject that the rent process follows a random walk and assume so for the baseline estimation.

D.2.2 Idiosyncratic Wage Process

The estimation of the idiosyncratic wage process involves using the PSID to estimate selection-corrected wages, following Low and Pistaferri (2015). In a future draft, this will also involve wage estimation using the HUD quarterly wages.

The model for (log) wages is

$$\ln w_{it} = \ln \mu_{\ell t} + X_{it}\beta + f_i + \omega_{it} + \varepsilon_{it}$$

where w_{it} is the total wage, $\mu_{\ell t}$ is the location component, X_{it} are household characteristics, f_i is fixed idiosyncratic productivity, ω_{it} is measurement error, and $\varepsilon_{it} = \xi_{it} + \varepsilon_{it-1}$ are productivity shocks that follow a random walk. In household characteristics, I include whether the head is married, quadratic age, and disability status.

The first differences equation is

$$\Delta \ln w_{it} = \Delta \ln \mu_{\ell t} + \Delta X_{it} \beta + \Delta \omega_{it} + \Delta \varepsilon_{it}$$

Estimation this equation via OLS is biased for two reasons. First, measurement error in the PSID earnings $\Delta \omega_{it}$ will bias the variance of productivity shocks upward. Second, wages are not observed for those who did not work.

Table 18: Panel Results: pVAR Estimates of CZ Wage-Rent Process (Annualized). Annualized standard errors estimated via Delta Method.

Panel A: Baseline Specification

(a) pVAR Coefficient

Variable	Rent	Income
Lag Rent	1.001 (0.001)	
Lag Income	(====)	0.989 (0.001)
R^2	0.87	0.86

(b) Covariance Matrix

	Rent	Income	
Rent	.0013		
Income	.0004	.0003	
Correlation: 0.55			

Panel B: Alternative Specification: Multivariate pVAR

(c) pVAR Coefficients

Variable	Rent	Income
Lag Rent	0.998	0.008
	(0.002)	(0.001)
Lag Income	0.008	0.975
· ·	(0.004)	(0.001)
R^2	0.88	0.86
Eigenvalues	1.0005	0.972

(d) Covariance Matrix and Eigenvalues

	Rent	Income	
Rent	.0012		
Income	.0003	.0003	
Correlation: 0.51			

Panel C: Initial Wage-Rent Covariance

(e) Covariance Matrix and Eigenvalues

	Rent	Income
Rent Income	.0414 .027	.026
Correlation: 0.85		

Table 19: Annualized estimates of the location wage-rent process using Census data. Estimates transformed from decadal to annual following Kueng et al. (2023).

To correct for employment selection, I follow Low and Pistaferri (2015) and use welfare benefits as a simulated instrument. Households who are eligible for higher welfare benefits when not working are more likely to select out of employment. This leads to a first-stage to estimate the selection into employment. The correction for measurement error comes from the autocovariance in differenced wages over time, since measurement error shocks are mean reverting. I estimate the resulting wages in a GMM framework similar to Low and Pistaferri (2015).

D.3 Details on Model Solution

A major challenge of applying a discrete-time lifecycle model to rental assistance is the dimensionality of the problem. Adding location processes, rationing mechanisms, and heterogeneity in assistance preferences greatly expands the dimensionality over other lifecycle models in the literature. Due to the curse of dimensionality, the model solution quickly becomes intractable to carry out indirect inference, let alone solve the model once.

Adaptive sparse grids overcome many of the primary dimensionality issues that these models face. The sparse component builds a sparse basis grid that has mathematical guarantees of how the error rate decreases with the resolution of the grid. Since these guarantees only work for smooth functions, the adaptive component allows the grid to refine in parts of the state space that are non-smooth, such as near eligibility cutoffs.

In addition, I implement the model solution on the GPU to greatly speed up the computation of the model. Due to recent advances in GPUs, a single GPU may outperform thousands of CPUs for the same task. However, the problem needs to be implemented to run efficiently on GPU, often re-casting computation as an embarrassingly parallel problem. In the case of this lifecycle model, this involves batching value function evaluations in ways that can be run embarrassingly parallel.

I implement the solution across a wide variety of GPUs. Running on a single H100 (with 8 supporting CPU cores), the model solution takes on average 24 minutes to solve. Other GPUs, such as Tesla V100 or L40S, take around 65 minutes to solve. Model solution times can be greatly improved by more efficiently allocating memory than the current adaptive sparse grid implementation and by parallelizing the computation across GPUs.

D.4 Homeownership Transition

To reduce the complexity of the model, homeownership transitions follow an exogenous probabilistic process. Households begin renting and will transition into homeownership when wages are high and/or housing costs are low. I assume household transitions take place before retirement.

Variable	Coefficient
Budget Variables	
Log Wage	0.184
Log Rent Price	-0.461
Log Asset	0.096
Age Bands	
24-25	0.222
26	0.477
28	0.584
30	0.609
32	0.643
34	0.496
36	0.330
38	0.311
40-44	0.309
45	0.233
50	0.212
55	-0.084
60	0.007
Const	-2.788
N	16,446
R^2	0.0669

Table 20: Coefficients of probit model

Using all households that are still renting, I estimate the follow probit model:

homeown_{it} =
$$\beta_0 + \beta_1 \ln(w_{it} + w_{\ell t}) + \beta_2 \ln p_t^h + \beta_3 \ln A_{it} + \beta_4 b_{it} + \delta_t + \varepsilon_{it}$$

where b_{it} and δ_t are indicators for head age bin and year, respectively. I group households into age bins for power, using two-year increments before 40 and five-year increments after.

I present the resulting estimates in Table 20. I find that all budget variables (i.e., wages, rents, and assets) significantly impact the probability of purchasing a home. I also find that the probability of purchasing a home peaks between the ages of 28-33. In Figure D.1, I show how model simulation of the homeownership transition match the data. Because households begin the lifecycle as renters, the simulations underestimate the number of homeowners compared to the data. However, what is important for eligibility is that the homeownership rates match those of older ages who may be less likely to work and therefore eligible for rental assistance.

D.5 Baseline Exit Estimation

[TODO]

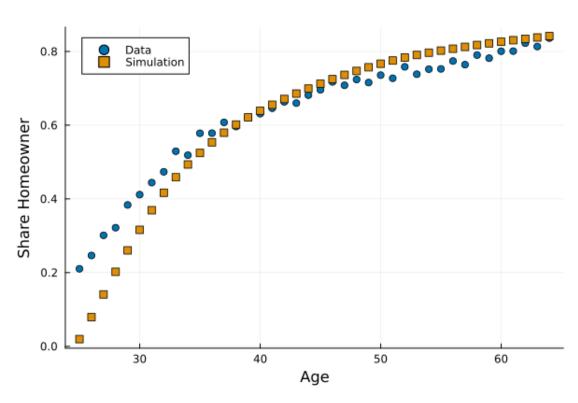


Figure D.1: Share of households who are homeowners by age, comparing the PSID to the simulations.