Self-Targeting in U.S. Transfer Programs*

Charlie Rafkin
Adam Solomon
Evan J. Soltas

October 2023

Abstract

Transfer receipt is voluntary and costly, generating "self-targeting" through selective take-up among the eligible. How does self-targeting select on need, and what are its policy implications? We show self-targeting is advantageous in eight U.S. transfers: On average, recipients have lower consumption and lifetime incomes than eligible nonrecipients with similar current incomes. Due to self-targeting, these transfers provide 50 to 75 percent more to the consumption-poorest and lifetime-poorest than would automatic transfers that are distributionally equivalent by income. Self-targeting makes automatic transfers undesirable: We estimate the social benefits of self-targeting are approximately six cents per transfer dollar, generally exceeding the social costs of ordeals.

Keywords: transfer programs, take-up, eligibility, self-targeting, consumption inequality
JEL Codes: H23, H53, I38

*MIT Economics (crafkin@mit.edu, adamsol@mit.edu, esoltas@mit.edu). We are deeply grateful to David Autor, Amy Finkelstein, Jim Poterba, and Frank Schilbach for their guidance. We also thank Jenna Anders, Jon Cohen, Joel Flynn, Jon Gruber, Nathan Hendren, Antoine Levy, Bruce Meyer, Abby Ostriker, Dev Patel, Anna Russo, Ben Sprung-Keyser, Iván Werning, and particularly Hunt Allcott for helpful conversations. This material is based upon work supported by the Lynde and Harry Bradley Foundation and the National Science Foundation Graduate Research Fellowship under Grant No. 1122374.

1 Introduction

Millions of Americans are eligible for means-tested transfer benefits but do not claim them. These programs rarely reach more than two in three eligible households, and some serve fewer than half of the eligible. Instead of these voluntary transfers, the government could help all eligible people automatically, cutting benefits to make this switch budget-neutral. Such a reform would save the hassle costs that recipients incur to claim benefits. Yet this reform would also give up a potential advantage of voluntary transfers: self-targeting. A transfer induces self-targeting when selective take-up among the eligible implicitly reveals dimensions of need which the government cannot itself readily use in determining eligibility. This paper studies the extent of self-targeting in U.S. transfers and the resultant social trade-off between voluntary and automatic programs.

The importance of voluntary transfers is a distinguishing and controversial feature of the U.S. safety net. By comparison to the U.S., other developed countries redistribute more through automatic and near-universal programs.[1] In the Covid-19 pandemic, the U.S. made two transfers—Medicaid and school meals—essentially automatic as part of a broader, but temporary, expansion of its safety net. The post-pandemic unwinding of these expansions has renewed the debate over policy choices in U.S. transfers.[2] The issue of voluntary versus automatic redistribution also bears on a number of other longstanding policy issues, such as the simplification of transfer eligibility criteria, proposals for a universal basic income (UBI), the structure of child benefits, and the rationing of rental assistance through waitlists. Despite the broad relevance of self-targeting to transfer design, economists lack both a basic set of facts about its magnitude and a theoretical framework that maps those magnitudes into social-welfare implications.

A necessary condition to prefer voluntary over automatic transfers is that selection into take-up

[1] As a result of this difference in safety nets, a larger share of redistribution in the U.S. than in Europe is between-lifetime, rather than within-lifetime (Bartels and Neumann, 2021). Other countries also have incomplete take-up in their voluntary programs (Eurofound, 2015).

[2] In their widely-cited book, Herd and Moynihan (2019, p.24) write "[t]he most dramatic way by which the state can reduce application compliance burden is to auto-enroll eligible individuals into a program based on administrative data." About 8.2 million people are forecast to lose Medicaid benefits at the end of auto-enrollment, of which 6.8 million (83 percent) is expected from non-take-up among the eligible (U.S. Department of Health and Human Services, 2022). Considerable policy action has already occurred in the context of school meals: Five states have made the pandemic-era expansion permanent as of spring 2023 (see https://frac.org/healthy-school-meals-for-all).

is "advantageous," that is, attracting those with higher need. Without such selection, automatic transfers are superior as they avoid hassle costs. Economists have long argued that a purpose of costly "ordeals" is to induce advantageous selection (Nichols and Zeckhauser, 1982; Besley and Coate, 1992). Others have contended, however, that ordeals perversely screen out the neediest households (Currie and Gahvari, 2008; Mullainathan and Shafir, 2013). More recently, empirical studies of ordeals and information interventions have found mixed results as to their targeting properties.[3] This emerging body of research has examined selection at the margin of specific interventions, but it has not considered how take-up selects on average. Yet selection on average determines whether transfers should be voluntary or automatic, as automatic transfers would redistribute to people both near and far from the margin of taking up a voluntary transfer.

We study selection into eight means-tested transfers that constitute most of the U.S. safety net using data from the Panel Study of Income Dynamics (PSID) on households' incomes, consumption, and transfer receipt over time. Our focus is on selection on average according to household consumption and lifetime income, which are often viewed as superior measures of living standards to annual income, since households can save or dissave to smooth transient income shocks.[4] More precisely, we examine selection on consumption and lifetime income given current income. Our motivation is that the government can set taxes or transfers according to current income but cares about inequalities in consumption and in lifetime income that it cannot directly address.

We find advantageous selection on both consumption and lifetime income for all eight transfers we study, often of substantial strength. In the Supplemental Nutrition Assistance Program (SNAP), for instance, the average consumption rank of recipients is 19 percentiles lower than likely-eligible non-recipients with similar incomes, a difference of about $11,000 per person per year or 46 percent of their average per-capita consumption. These gaps are somewhat smaller for other transfers and for selection on lifetime income. Take-up appears to distinguish temporarily low-income households,

[3]Recent papers studying the targeting properties of ordeals and information include Bhargava and Manoli (2015), Ganong and Liebman (2018), Deshpande and Li (2019), Finkelstein and Notowidigdo (2019), Gray (2019), Lieber and Lockwood (2019), Homonoff and Somerville (2021), Unrath (2021), Arbogast et al. (2022), Shepard and Wagner (2022), Wu and Meyer (2022), and Ericson et al. (2023).

[4]See, e.g., Vickrey (1947) or more recently Auerbach et al. (forthcoming). Sullivan et al. (2008) find that permanent income is far more predictive of measures of material hardship than is the transient component of income.

who can smooth consumption, from persistently low-income ones who cannot.

We then examine the implications of our estimates of self-targeting for real-world policy analysis. We consider budget-neutral reforms that shift from voluntary to automatic transfers: In particular, we reallocate the value of benefits claimed by people at a given income to all people with the same income, including current non-recipients. Due to self-targeting, households at the bottom of the distributions of consumption and lifetime income currently receive 50 to 75 percent more under voluntary transfers than they would under this reform.

A second insight for policy analysis follows from contrasting the distributional incidence of transfers by current income with their incidence by consumption and lifetime income. In the context of tax policy, economists have argued that income-based analyses usually overstate the welfare-relevant notion of progressivity or regressivity due to year-to-year household income fluctuations (Poterba, 1989, 1991). We find that, in most U.S. transfers, self-targeting entirely offsets this "smoothing" effect of lifetime and consumption incidence. The other side of the coin is that the low current incomes of many eligible non-recipients overstate their need as indicated by their consumption or lifetime income. Self-targeting thus raises new concerns about income-based distributional analysis beyond those from the tax context.

Should transfers be automatic, given the extent of self-targeting when they are voluntary? We use a model of optimal redistribution to assess the trade-off between the social benefits of self-targeting and the social costs of ordeals.[5] In the model, the government sets the schedules of voluntary and automatic transfers as functions of current income alone, while marginal utilities and social welfare weights depend on consumption. Hassle costs in the voluntary transfer cause its take-up to be incomplete, which enables self-targeting but wastes real resources. We derive formulas that address the government's allocation problem between the voluntary and automatic transfers: Should it enact a small budget-neutral cut to the voluntary transfer so as to fund a small increase in the automatic one? For instance, the government could provide some SNAP benefit automatically and a top-up to only those who apply, rather than entirely on application. This

[5]This contrast also addresses the question of redistribution via taxes versus via transfers, albeit while ignoring differences in willingness to pay for in-kind goods from the government's cost to provide them.

comparison holds the ordeal fixed, thus differing from Finkelstein and Notowidigdo (2019), and has attractive theoretical properties.

Calibrating this welfare formula using our empirical estimates, we draw three conclusions about the social costs and benefits of self-targeting. First, self-targeting yields quantitatively important social benefits: In our baseline calibration, we estimate them to be to approximately six cents per transfer dollar, taking a dollar-weighted average across programs.

Second, overall across transfers, the social benefits of self-targeting likely exceed the social costs of ordeals. We provide upper-bound estimates of ordeal costs by an application of the envelope theorem, which implies that marginal ordeal costs equal fiscal costs of marginal recipients when take-up choices are made optimally. Our upper-bound estimates are large: In annual per-recipient terms, we find ordeal costs could be as high as \$240 for SNAP or \$500 for Medicaid. These ordeal costs are nevertheless often equaled or surpassed by our estimated benefits of self-targeting. Models in which incomplete take-up is, at least in part, a result of non-optimizing behavior imply smaller social costs of ordeals, as the marginal ordeal costs of non-optimizers must be less than their marginal fiscal costs. Our results establish self-targeting as an empirically credible argument for existing U.S. transfers and against reforms that move toward automatic redistribution.

The third conclusion from our welfare analysis is that there is important between-program variation in the welfare effects of moving from voluntary to automatic redistribution. This variation is explained by differences in the intensity of self-targeting. In SNAP and housing assistance, self-targeting is particularly strong and worth more than ten cents per transfer dollar to society. This amount greatly exceeds upper-bound estimates of ordeal costs, suggesting that making these transfers automatic is very unlikely to be socially desirable. By contrast, self-targeting is of less social value in Medicaid, two transfers that provide food to children (WIC and the National School Lunch Program), and one transfer for utility assistance (LIHEAP). The programs nevertheless inflict ordeal costs. Such examples show how our framework has nuanced and useful applications to real-world policy issues.

Our analysis uses PSID data from 1997 to 2019. The misreporting of transfer receipt and other key variables, namely income and consumption, poses in our view the most important challenge to our analysis (Meyer et al., 2009, 2015). We take this challenge seriously in several ways. First, on transfer misreporting, we adopt the multivariate corrections of a recent literature that estimates how misreporting probabilities vary with household characteristics by linking survey and administrative data (Davern et al., 2019; Mittag, 2019; Meyer et al., 2020). These corrections actually strengthen our results. Second, on consumption misreporting, advantageous selection holds for consumption subcategories thought to be well-measured, as well as for most measures of durable goods ownership (Meyer and Sullivan, 2023). Third, on lifetime income, we extend methods from Haider and Solon (2006) to address potential biases from incomplete income histories. The comparison to independent measures of consumption also offers a form of replication (Caspersen and Metcalf, 1994). Fourth, on transfer eligibility, our results are robust to reclassifying simulated-ineligible recipients as eligible and to comparisons of observably-similar households, which are likelier to have the same true eligibility status (Duclos, 1995). Our sensitivity analyses give us confidence that advantageous self-targeting is not an artifact of problems with survey data.

Our paper adds to two literatures in public finance. Our primary contribution is to research on the take-up of transfers (Currie, 2006; Kleven and Kopczuk, 2011). Following Nichols and Zeckhauser (1982) and Besley and Coate (1992), many theoretical analyses have considered the importance of self-targeting to the efficiency of in-kind transfers (as reviewed in Currie and Gahvari, 2008). Our key departure from the empirical literature cited above is to consider whether selection into transfers is advantageous on average, not on the margin of specific ordeals—a shift in focus that we show to be consequential. Our results suggest the mixed findings of this literature are best interpreted as faulting specific ordeals, not voluntarism in transfers generally.[6]

We also contribute to a literature on distributional incidence that assesses who pays for and benefits from government. Several recent papers examine taxes and transfers from a lifetime perspective (Blundell et al., 2015; Bengtsson et al., 2016; Roantree and Shaw, 2018; Brewer et

[6] Our test of self-targeting most resembles Alatas et al. (2016), who show that hassle costs in an Indonesian transfer program deter applications from people with unobservably high consumption, controlling for an observable proxy used in eligibility determinations.

al., 2020; Levell et al., 2021). We provide a consumption-based analysis, which is of particular relevance for policies that serve populations who are likely to be credit-constrained. Prior work has also not observed that transfers are much more progressive in consumption or lifetime terms than their current incidence would suggest, perhaps because take-up is more complete in their settings.[7] One closely related paper is Auerbach et al. (forthcoming), which studies U.S. fiscal progressivity with respect to lifetime income. We build on their analysis by identifying the central role of self-targeting in incidence and by considering its welfare implications.

2 Data and Measurement

Our main source of data is the Panel Study of Income Dynamics (PSID) in its eleven biennial survey waves from 1997 to 2019. In each PSID wave, we observe heads of household and spouses ages 18 to 65. Here we first review key aspects of the data, leaving further details to Appendix B. We then explain three imputation procedures that augment the PSID data: for cash-equivalent values of in-kind transfers, transfer eligibility, and lifetime income.

Our goal is to measure selection into transfers on consumption and lifetime income. The PSID data has several crucial features for this purpose, including its long panel dimension to estimate lifetime income, its consumption data, and its information on receipt of several transfer programs. Its major limitations are the reporting issues that we discuss in depth in Section 3. We conclude there that our results appear mostly robust to known issues with misreporting in survey data.

2.1 Income, Consumption, and Transfer Receipt

Current Income. We define household income as total annual income of the head and spouse before taxes and transfers, excluding other household members. Income includes labor, business, and capital income. Following the National Academy of Sciences (Citro and Michael, 1995), we

[7]Fullerton and Metcalf (2002) reviews the literature on the lifetime incidence of taxation. The few related U.S. studies of lifetime transfer incidence often assume away the take-up margin of transfers, treating them as equivalent to taxes (Davies et al., 1984; Fullerton and Lim Rogers, 1993; Guner et al., 2021). Of course, some research implies that self-targeting could be important for distributional analysis. For instance, Blank and Ruggles (1996) find heterogeneous income dynamics that correlate with take-up.

adjust for household size using the equivalence scale $e_h = (N_{h,\text{adult}} + 0.7N_{h,\text{child}})^{-0.7}$, where $N_{h,\text{adult}}$ and $N_{h,\text{child}}$ respectively denote the numbers of adults and of children in household h.[8] Income ranks are computed within years, pooling across birth-year cohorts.[9]

Current Consumption. The PSID has extensive coverage of consumption expenditures since 1999. Expenditure categories include food, housing, health care, transportation, education, child care, and several smaller topics. We make two adjustments to the data to better reflect consumption rather than expenditure, following Meyer and Sullivan (2023). First, for homeowners, we replace mortgage and property tax payments with equivalent rents based on reported home values. Second, for vehicle owners, we replace loan payments with estimates of lease-cost equivalents. Household consumption is then equivalized as above. Consumption ranks are also computed within years.[10]

Transfer Receipt. We observe self-reported household receipt for ten transfer programs. These are the Supplemental Assistance Nutrition Program (SNAP); Medicaid; Section 8; public housing; Temporary Assistance for Needy Families (TANF); Supplemental Security Income (SSI); Women, Infants, and Children (WIC); and the Low Income Home Energy Assistance Program (LIHEAP); and the National School Lunch Program and School Breakfast Program. We combine public housing and Section 8 into one program to which we refer as "housing assistance," and the lunch and breakfast programs into "school meals." Table 1 provides summary statistics on the transfers.

[8] Appendix A includes results not adjusted for household size and composition. These are typically quite similar to those with equivalized households. By implication, other equivalence scales are also likely to yield similar results.

[9] We do not rank households within both year and cohort because, in early adulthood, transfer receipt reaches higher up into cohort-level income distributions, as incomes are lower. Such ranks would show transfers are less progressive in current income and consumption than our within-year ranks do, only strengthening our conclusions.

[10] Consumption is inherently measured after taxes and transfers. Transfers thus elevate the consumption ranks of transfer recipients relative to their income ranks, a force that works against our conclusions.

Table 1: Means-Tested Transfer Programs in the U.S.

	SNAP	Medicaid	Housing Assistance	TANF	SSI	WIC	LIHEAP	School Meals	Any Transfer	U.S. Population
Budgetary Cost in 2019 (billions)	60.4	613.5	41.7	30.9	55.8	5.3	3.7	18.7	n.a.	n.a.
Receipt Rate, Households	14.1	20.0	6.5	1.0	6.2	7.5	5.1	18.9	33.8	n.a.
Mean Annual Benefit, Recipients	4,062	5,802	6,893	14,144	3,368	205	535	647	n.a.	n.a.
Characteristics of Households or Heads of Recipient Households										
Mean Age, Head	42.1	42.7	40.9	35.6	47.1	34.3	45.4	40.2	43.0	44.3
% Married	19.8	30.7	11.7	15.6	25.2	37.4	26.5	35.5	33.0	47.2
Mean Household Size	3.7	3.8	3.1	3.5	3.2	4.5	3.7	4.3	3.6	3.0
% Children at Home	51.6	58.9	40.3	90.9	27.7	92.0	48.2	94.1	50.5	30.9
% Nonwhite or Hispanic	64.6	62.4	73.0	74.8	60.2	70.3	58.7	69.0	62.2	40.7
% H.S. Graduate	70.6	73.4	74.3	60.0	73.8	71.4	70.6	71.9	76.2	88.6
Mean Household Income	16,263	27,992	17,829	11,735	20,501	32,091	17,223	34,860	31,741	78,506
% Employed	45.7	53.9	50.1	40.0	37.3	69.5	44.7	70.7	58.4	77.0
Mean Rank, Equivalized Households										
Current Income	16.7	22.3	19.8	12.6	18.6	24.1	17.2	25.5	25.2	50.0
Consumption	17.0	22.8	16.9	11.9	27.8	19.3	19.8	22.4	26.4	50.0
Lifetime Income	24.3	30.6	24.9	22.6	27.1	34.1	26.0	34.3	33.3	50.0

Notes: This table reports summary statistics on the eight means-tested transfer programs we study. See Appendix B for sources on budgetary costs. All other values are from the PSID. Monetary values are expressed in 2020 constant dollars.

2.2 Imputation of Other Variables

Valuing Transfers. We measure the dollar value of transfers by combining information from the PSID and the Supplemental Poverty Measure module of the U.S. Current Population Survey (CPS). For SNAP, TANF, SSI, UI, and LIHEAP, the PSID records the nominal value of transfers over various time periods, which we rebase as the per-capita annualized amount in 2020 constant dollars. The PSID does not include cash-equivalent values for in-kind transfers, namely Medicaid, Section 8, public housing, and WIC. We impute these amounts with the average values by household size and year reported in the CPS for all but WIC, where we use the national average benefit. The CPS generally values in-kind transfers dollar-for-dollar with expenditure.[11]

Lifetime Income. We construct a lifetime concept of household income from incomplete income histories. To begin, we estimate a Poisson regression model with individual fixed effects, interacted with age-specific coefficients as recommended by Haider and Solon (2006). Letting i index individuals, t index calendar years, and a index age in years, the model takes the following form:

$$\mathrm{E}[y_{it} \mid X_{it}] = \exp(\alpha_i \lambda_a + X_{it}' \beta_a), \tag{1}$$

where α_i is an individual fixed effect, X_{it} is a matrix of time-varying demographic characteristics, and λ_a and β_a are vectors of age-specific coefficients. The outcome y_{it} is individual income.

We then perform several adjustments, discussed fully in Appendix B, before using the regression results to estimate lifetime income. These adjustments shrink the fixed effects to account for sampling variation and impute demographic characteristics to balance the panel. We calculate lifetime average income from ages 18 to 65, and then we account for spousal income in a way that permits changes in household composition over time. In particular, let $j(i,t)$ indicate i's spouse in year t. Our concept of lifetime household income follows each individual through the sequence of

[11] Except for Medicaid, for which the Census produces household-level "fungible values" and individual-level "market values." We use fungible values, so as to remain at the household level.

households they experience as adults. That is, the lifetime household income of individual i is

$$\overline{y}_i^h = \sum_t e(\widehat{y}_{it}^h) = \sum_t e(\widehat{y}_{it} + \widehat{y}_{j(i,t),t}) \tag{2}$$

where t is again summed over the years in which i is between ages 18 and 65, and $e(\cdot)$ is the equivalence-scale function. If we restrict our sample to stable households (as in, e.g., Fullerton and Lim Rogers, 1993), our definition of lifetime income would coincide with the standard concept. Lifetime-income ranks are computed within birth-year cohorts.

Simulated Eligibility. Studying self-targeting among the eligible requires measures of transfer eligibility, so that we can distinguish the ineligible from eligible people who do not take up. We simulate eligibility by compiling information on program rules, mainly from primary-source documents and research databases of such rules, similar to the Urban Institute's TRIM program (Zedlewski and Giannarelli, 2015). See Appendix B for details on these eligibility simulations.

Eligibility simulations cannot perfectly capture true eligibility, as information used in actual eligibility determinations differs versus what is recorded in surveys. In validation checks in Appendix B, our simulated eligibility measure is highly predictive of benefit receipt, though misclassification is apparent. Significant fractions of recipients are simulated to be ineligible, and take-up rates are counterfactually low among the simulated-eligible. Both are routine issues in eligibility simulations (Duclos, 1995). The main threat from mismeasured eligibility in our analysis lies in understating the importance of eligibility rules relative to self-targeting among the eligible. In sensitivity analyses, we show how our results change when we reclassify simulated-ineligible recipients as eligible, which mechanically raises the influence of eligibility rules.

3 Selection into Transfer Receipt

This section quantifies advantageous selection into transfers: Holding incomes fixed, recipients have lower consumption and lifetime income than eligible non-recipients with similar incomes. By comparison, we find that eligibility rules contribute little to selection on consumption or lifetime income. Finally, we assess the sensitivity of our results to measurement issues.

3.1 Empirical Definition

We first provide a formal definition of self-targeting in transfer programs that we will take to the data. We say there is *advantageous self-targeting* on an outcome Y_i, letting i index households, if transfer recipients are negatively selected on the outcome relative to eligible nonrecipients with the same eligibility-rule observable characteristics. That is:

$$\mathrm{E}[Y_i \mid D_i = E_i = 1, X_i] \;\leq\; \mathrm{E}[Y_i \mid D_i = 0, E_i = 1, X_i], \tag{3}$$

where D_i and E_i respectively indicate transfer receipt and eligibility, and X_i contains the eligibility-rule observables. Equation 3 is motivated by the correlation test of Chiappori and Salanie (2000). Consider, in particular, the following joint model of the outcome and transfer receipt:

$$Y_i = X_i\delta + \alpha D_i + \beta E_i + \nu_i$$

$$D_i = \begin{cases} 1[X_i\gamma + \varepsilon_i \geq 0] & \text{if } E_i = 1 \\ 0 & \text{if } E_i = 0. \end{cases}$$

The analog in our context to the correlation test is $\mathrm{corr}(\nu_i, \varepsilon_i) < 0$: Households that are unobservably more likely to take up the transfer when eligible are negatively selected on the unobservable component of the outcome. If $\alpha > 0$ (transfer receipt raises the outcome), then Equation 3 is sufficient for the correlation test to be satisfied. It is stronger (i.e., not necessary) because we do not account for the direct effect of transfer receipt on the outcome.

3.2 Who Gets Transfers?

Table 2 reports the average annual per-capita value of benefits for households who fall into a given combination of income quintile and consumption quintile. Moving across the income distribution,

Table 2: Average Annual Per-Capita Total Transfer Benefits
by Quintile of Current Income, Lifetime Income, and Consumption

		Income Quintile					
		1	2	3	4	5	Avg.
Consumption Quintile	1	3,647	1,353	600	397	155	2,440
	2	1,745	719	296	134	80	666
	3	920	563	217	102	33	303
	4	572	403	168	60	33	153
	5	557	273	133	58	18	101
	Avg.	2,435	844	266	92	27	
		1	2	3	4	5	Avg.
Lifetime Income Quintile	1	3,346	1,243	498	253	28	2,208
	2	1,594	839	278	103	36	627
	3	1,272	664	230	88	36	349
	4	1,152	556	211	79	26	242
	5	1,344	522	189	66	23	239
	Avg.	2,435	844	266	92	27	

Notes: This table reports the average annual per-capita total value of transfer benefits, cash and in-kind, by quintiles of equivalized household current income, lifetime income, and consumption. Values are in constant 2020 dollars.

higher-income households unsurprisingly receive less in transfers. Yet, comparing households in the same income quintile but with different levels of consumption, lower-consumption households receive more in transfers than higher-consumption households. Thus, there is advantageous selection on consumption. For instance, among households in the bottom income quintile, those also in the bottom consumption quintile receive six times more in transfers as those also in the top consumption quintile. We find similar results for selection on lifetime income. The similarity of the results is especially helpful, in that both consumption and lifetime income require assumptions to estimate but are constructed entirely separately of one another. In particular, it provides a safeguard against the possibility that some unfortunate choice we have made in defining these variables could be responsible for the selection patterns.

These tabulations do not account for within-quintile income differences, and it is possible that apparent selection on consumption merely reflects that households with the lowest income within each quintile have the lowest consumption and also take up the transfer. A flexible rank–rank regression specification to control for income addresses this concern. We therefore estimate

$$\overline{R}_{it} = \beta D_{it} + f(R_{it}) + u_{it}, \tag{4}$$

where $\overline{R}_{it}$ is the consumption rank or lifetime-income rank for household i in year t, R_{it} is i's current-income rank, $f(R_{it})$ is a flexible function of this rank, and D_{it} indicates i's receipt status for a given transfer program. The coefficient β summarizes the extent of advantageous selection into a transfer. We parameterize $f(R_{it})$ using cubic splines with knots at the 10th, 25th, and 50th percentiles of the current-income distribution.

Panel A of Figure 1 shows the patterns in Table 2 hold up in the regression and apply generally across programs. Overall, recipients of a given transfer rank about 15 percentiles lower in the consumption distribution than non-recipients of that transfer with similar current incomes. Appendix Table A4 estimates the Poisson-regression equivalent of Equation 4, with levels of annual consumption and lifetime income per capita as outcomes. The rank differences are consistent with differences of approximately 30 to 60 percent in these outcomes, or around $7,500 to $14,000 per person per year in consumption.

The extent of selection into transfers on consumption and lifetime income varies considerably across transfer programs. For instance, SNAP and housing assistance are highly informative about consumption and lifetime income, whereas SSI is less informative about consumption given income. Some "non-differences" are also interesting: for example, it is not the case that cash-like transfers (e.g., SNAP) are systematically more selective than the transfers most unlike cash (e.g., housing or Medicaid). Results for lifetime income are similar, although there are some notable differences (Panel B of Figure 1). One is that TANF recipients are highly negatively selected on consumption but are much less selected on lifetime income.

Panel C of Figure 1 estimates the extent of advantageous selection into transfers, distinguishing by the number of transfers received. Multiple-recipient households are more advantageously selected on consumption and lifetime income than households receiving only one transfer. One explanation for such a pattern is that each decision to apply for a transfer incrementally reveals more information about living standards conditional on income.

3.3 Self-Targeting or Eligibility Rules?

Our analysis has so far pooled the ineligible and eligible. The results may thus reflect not only self-targeting but also eligibility requirements which select on correlates of consumption and lifetime income, such as asset tests or categorical eligibility for some groups (e.g., people with disabilities). We use our simulated-eligibility measures to disentangle the contributions of self-targeting and eligibility rules. We find self-targeting is the primary force, and eligibility rules a distinctly secondary one, in advantageous selection into transfers.

Table 3 shows the importance of self-targeting over eligibility rules, with SNAP as an example. Panel A shows that receipt rates of SNAP decline in both income rank and consumption rank given income rank. Panel B shows rates of simulated eligibility for SNAP by income and consumption quintile. Unsurprisingly, SNAP eligibility falls quickly in income; very few households above the second income quintile are SNAP-eligible. The rate of simulated SNAP eligibility also falls in consumption given income but less markedly than does receipt. These declines in eligibility with respect to consumption are driven by asset tests in SNAP, which existed until 2014, as well as details of the transfer's income-eligibility criteria. Take-up rates among simulated eligibles, as shown in Panel C, fall sharply in consumption given income and thus explain the difference between Panels A and B. Among eligible households in the bottom income quintile, the SNAP take-up rate among those also in the bottom consumption around 52 percent, as compared to approximately 8 percent among top-consumption-quintile households.[12]

[12]While take-up rates are sensitive in *levels* to the general expansiveness or conservativeness of any eligibility simulation, measurement issues can less easily explain the vast *differences* in take-up by consumption given income.

Figure 1: Self-Targeting in Transfer Programs

(a) Selection on Consumption

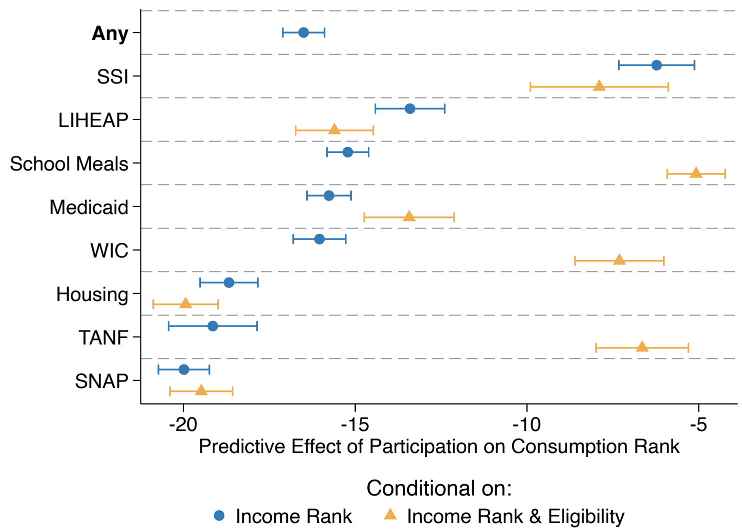

(b) Selection on Lifetime Income

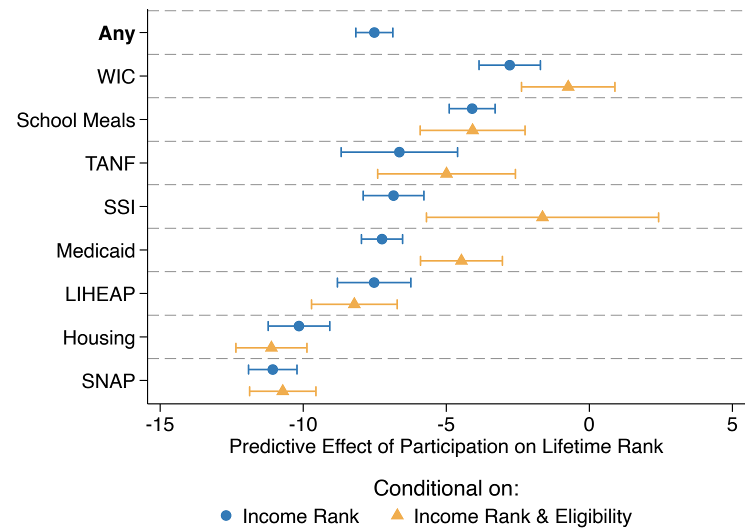

(c) Selection by Number of Transfers Received

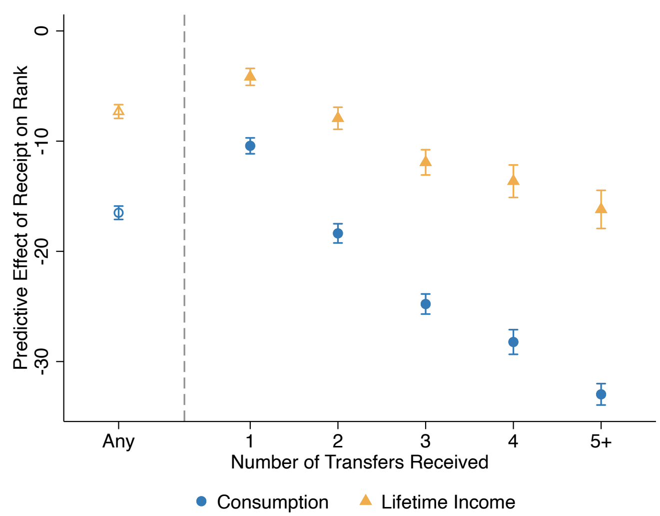

Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank (Panel A) or lifetime-income rank (Panel B), conditional on current-income rank (coefficient γ from Equation 4). For the yellow diamonds, we estimate the regression only on people whom we simulate to be eligible. The "any" row of Panel A is an indicator for receipt of at least one of the eight transfers. 95-percent confidence intervals reflect clustered standard errors by household. In Panel C, we adapt Equation 4 by replacing the transfer indicator with indicators for the number of transfers received in that year.

Table 3: SNAP Receipt, Eligibility, and Take-Up Rates by Income and Consumption Quintile

Panel A: Receipt Rate

		Income Quintile 1	2	3	4	5	Avg.
Consumption Quintile	1	51.2	22.3	7.8	4.9	4.9	35.3
	2	23.7	9.6	2.7	1.1	0.5	8.4
	3	12.3	5.9	2.3	0.5	0.3	3.3
	4	6.3	3.4	1.4	0.2	0.2	1.3
	5	5.5	2.9	1.7	0.3	0.1	0.9
	Avg.	33.6	12.2	2.8	0.6	0.2	

Panel B: Simulated Eligibility Rate

		Income Quintile 1	2	3	4	5	Avg.
Consumption Quintile	1	83.8	23.7	0.4	0.8	0.0	51.9
	2	75.5	15.8	0.4	0.2	0.0	19.0
	3	67.5	14.0	0.4	0.2	0.1	10.5
	4	61.3	13.9	0.4	0.3	0.1	7.2
	5	60.7	17.6	0.5	0.3	0.0	6.5
	Avg.	76.3	18.1	0.4	0.3	0.1	

Panel C: Take-Up Rate Among Simulated Eligibles

		Income Quintile 1	2	3	4	5	Avg.
Consumption Quintile	1	52.2	39.3	.	.	.	50.2
	2	26.8	23.7	.	.	.	25.9
	3	14.9	14.4	.	.	.	14.5
	4	8.3	8.2	.	.	.	8.1
	5	8.1	8.6	.	.	.	8.1
	Avg.	37.5	27.2	.	.	.	

Notes: This table reports the shares of households that receive SNAP (Panel A), are simulated to be eligible for SNAP (Panel B), and take up SNAP conditional on being simulated to eligible (Panel C). Households are split by quintiles of equivalized household consumption and income. Due to low rates of simulated eligibility, we do not report take-up rates for the top three income quintiles. See Appendix A for a tabulation by income and lifetime income.

We generalize this analysis in Panel A of Figure 1, where we again estimate Equation 4 but now restrict the sample to households who are simulated-eligible. For programs like SNAP or housing assistance, advantageous selection on consumption or lifetime income into transfers therefore appears to reflect self-targeting among the eligible rather than eligibility. TANF and WIC are two notable exceptions. For these programs, conditioning on eligibility reduces the predictive effect of receipt on rank from 20 percentiles to about eight. Even so, the effect of receipt among simulated-eligibles remains economically large.

Mismeasurement of eligibility mostly appears unlikely to explain our conclusions. When we reclassify any simulated-ineligible recipients of a given program to be eligible for that program, we do not see a clearer role for eligibility in concentrating incidence among the consumption-poor (Appendix Figure A4). The transfer program where our results are potentially most sensitive to mismeasurement of eligibility is SSI. Even with this adjustment to simulated eligibility, take-up remains the primary force in shaping the consumption and lifetime incidence of transfers. We further examine eligibility mismeasurement in Section 3.5.

3.4 Incidence Analysis and the Automatic Counterfactual

What are the implications of self-targeting for distributional analysis? We answer this question in two steps. First we examine the incidence of existing transfer programs by consumption and lifetime income. This is, to our knowledge, the first such non-income-based distributional analysis for these programs. Then we contrast their incidence with that of automatic transfers.

To estimate distributional incidence, we rank households by either current income, current consumption, and lifetime income.[13] We then estimate receipt rates and average dollar values of benefits per person per year as locally-linear functions of households' percentile rank in a given distribution. Panel A of Figure 2 plots the average total annual per-capita value of transfer benefits as functions of these ranks. In the left sub-panel, we combine transfers across the eight transfer programs we study. The right sub-panel excludes in-kind programs, for which we must impute cash-equivalent values of benefits. Panel B plots estimates of receipt rates as a function of ranks for the eight transfer programs we study.

Both panels also show that transfer incidence with respect to consumption, lifetime income, and current income are remarkably similar, overall and for most individual transfer programs. That is,

[13] Appendix A contains supplementary figures that probe the sensitivity of our results to definitional choices and provide detail on distributional incidence. Distributional incidence is essentially uniform with respect to current income, lifetime income, and consumption in unemployment insurance (Appendix Figure A5), whose structure is essentially opposite to a means-tested transfer. This occurs because

Figure 2: Receipt and Value of Transfer Benefits as a Function of Equivalized Household Rank
Panel A: Average Total Annual Per-Capita Benefits

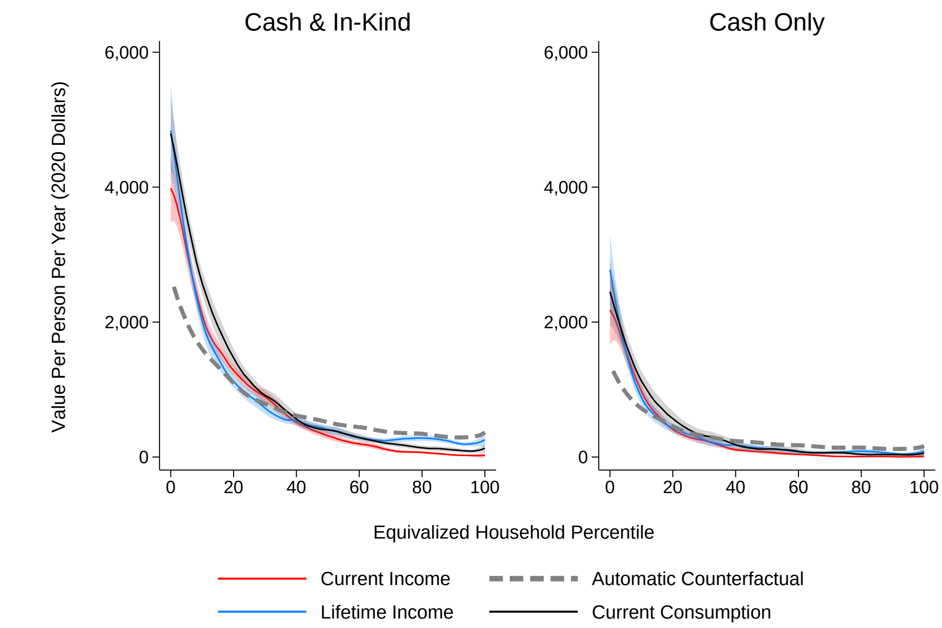

Panel B: Receipt Rate by Program

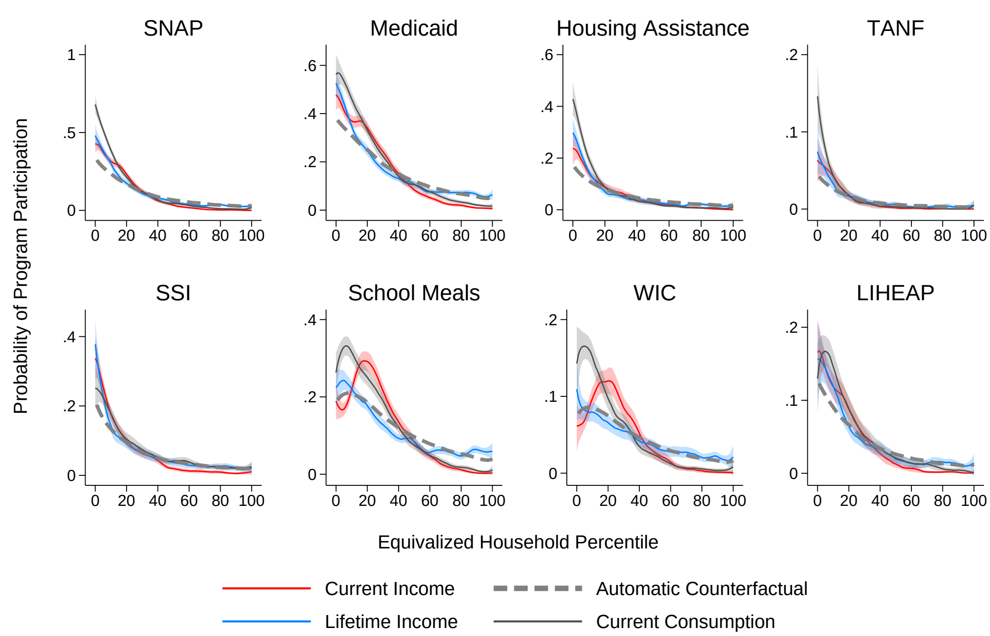

Notes: This figure displays average total annual per-capita dollar amounts and receipt rates as functions of household rank in the distributions of equivalized current income, lifetime income, and consumption. The functions are estimated by local linear regressions with bandwidths of three percentiles. Dashed lines indicate counterfactual transfer incidence with respect to consumption when the receipt rate is determined by income. Shaded regions reflect bootstrapped 95-percent simultaneous confidence bands, as in Chernozhukov et al. (2013), with clustering by household.

Figure 3: Eligible Non-Receipt Rate by Program

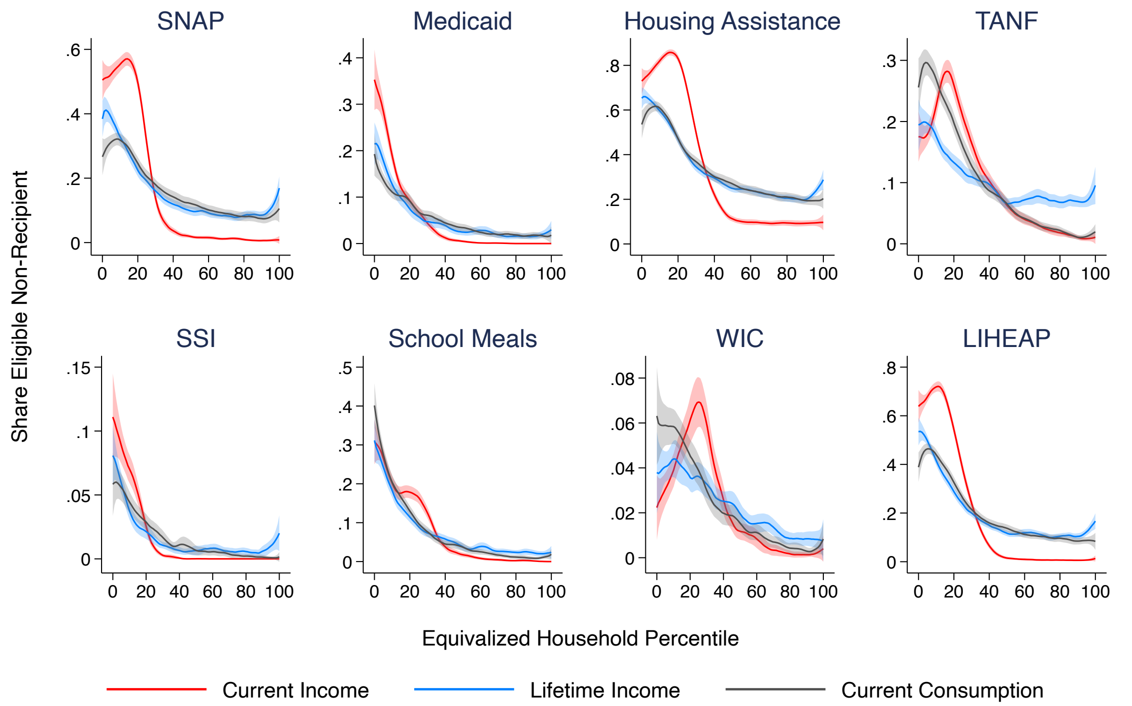

Notes: This figure displays the share of households that are eligible non-recipients of a transfer as functions of their ranks in the distributions of equivalized current income, lifetime income, and consumption. The functions are estimated by local linear regressions with bandwidths of three percentiles. Dashed lines indicate counterfactual transfer incidence with respect to consumption when the receipt rate is determined by income. Shaded regions reflect bootstrapped 95-percent simultaneous confidence bands, with clustering by household.

average annual per-capita benefits are about as high at the bottom of the distributions of consumption and lifetime income as they are at the bottom of the current-income distribution. Households in the bottom five percentiles of the income distribution, for instance, receive approximately $3,700 in transfers per person per year, as compared to $4,300 for the lowest-consumption households (Appendix Table A2). Furthermore, the similarity of consumption and lifetime incidence suggests transfer programs identify households with low consumption as a result of persistently low income, rather than a lesser ability to smooth consumption relative to income over time.

These results in Figure 2 are perhaps surprising because year-to-year income fluctuations should mechanically reduce the lifetime or consumption incidence of transfers at the bottom of the distribution. The apparent absence of this effect is not for a lack of year-to-year mobility at the bottom, as we show through a formal decomposition in Appendix A. Instead, the substantial compressive effect of income fluctuations is fully undone by selection into receipt.

Although eligible non-recipients of most transfers are poor in current-income terms, we show in Figure 3 that many have significant consumption or lifetime resources. This figure captures the two-sided implications of selective take-up. For eligible non-recipients, advantageous self-targeting compounds rather than offsets the compressive effects of consumption and lifetime distributional analysis. For instance, around one third of eligible non-recipients of SNAP and Medicaid have above-median consumption. A similar fraction has above-median lifetime income. Rates of eligible non-receipt are lower among the consumption- and lifetime-poorest than at the bottom of the current-income distribution. In Medicaid, about one fifth of the consumption- and lifetime-poorest are eligible but do not take up, as compared to one third of current-poorest.

Automatic Counterfactual. The role of self-targeting is made particularly clear by considering the consumption and lifetime incidence of an automatic transfer, in which self-targeting is inherently impossible. We consider a counterfactual in which the government automatically pays the average unconditional value of a given transfer to households at each income level. For instance, suppose households with $20,000 in income receive $1,000 if they take up a voluntary transfer, and half take up. In the counterfactual, the government would give $500 to all the households with $20,000 in income. Compared to a voluntary transfer, the automatic transfer gives less to the households who always take up and more to households who do not. We simulate this counterfactual by applying the empirical transition matrix from income ranks to consumption ranks.[14]

Relative to a voluntary transfer with the same distributional incidence, automatic transfers would give the lowest-consumption households about a third less per year on average (see the dashed lines in Figure 2). Voluntary transfers would provide the lowest-consumption households about $4,600 per person per year, compared to the $2,300 they receive with automatic transfers. These results hold for all eight programs (Panel B), and differences in incidence of the automatic counterfactual and voluntary transfers are generally large in magnitude. For instance, about 70 percent of the

[14]For the average benefit amount or receipt rate $r(y)$ at current-income rank y, this mobility counterfactual is $r(c) = \int r(y)\, dF(y|c)$, where $F(y|c)$ is the conditional distribution of income rank at consumption rank c.

consumption-poorest households receive SNAP, about double the share who would receive SNAP in a budget-neutral automatic version.

Appendix Table A2 reports point estimates and standard errors for transfer incidence in different parts of the distributions of current income and consumption, along with incidence in the automatic counterfactual. We estimate the same counterfactual for lifetime income and report similar if somewhat smaller results in Appendix Table A3.

3.5 Sensitivity to Mismeasurement

Survey data are imperfect.[15] Here we consider the potential for bias in our results due to known issues with self-reported transfer receipt, income, and consumption. We also address potential concerns with our measures of lifetime income and simulated transfer eligibility.

Transfer Receipt. Using linked survey and administrative data, Mittag (2019) and Davern et al. (2019) estimate statistical models of household survey reporting behavior for SNAP and Medicaid receipt respectively. Their models, intended for use as misreporting corrections, predict the probability of true transfer receipt given survey-reported receipt and demographic characteristics. These models allow researchers to replace assumptions of constant misreporting rates with misreporting probabilities that are functions of demographic observables.

Their corrections consistently increase our estimates of selection (Appendix Table A5). There are two reasons why. First, under constant misreporting rates, our estimates are attenuated. Consider misreporting probabilities $p_0 = \Pr(\widetilde{D}_i = 0 | D_i = 1)$ and $p_1 = \Pr(\widetilde{D}_i = 1 | D_i = 0)$, where D_i indicates true receipt and $\widetilde{D}_i$ indicates reported receipt. Comparing the feasible regression of $y_i = \widetilde{\beta}\widetilde{D}_i + u_i$ to the infeasible regression $y_i = \beta D_i + u_i$, one can show that $\beta = \widetilde{\beta}/(1 - p_0 - p_1)$.[16] Second, the parameter estimates in Mittag (2019) and Davern et al. (2019) both imply that underreporting of transfer receipt is somewhat more common among households with low consumption and lifetime

[15]Administrative data would not be a panacea, however, in a key respect: Eligibility would have to be imputed for non-recipients, likely with fewer eligibility-rule variables than are available in the PSID.

[16]Meyer et al. (2009) finds rates of under-reporting rates in the PSID in the range of 7 to 27 percent across programs we study. This suggests a presumption that our main estimates in Figure 2 are understated, even for transfer programs where heterogeneous-misreporting corrections have not yet been estimated.

income, holding income constant. Thus, their adjustments tend to amplify the increase in the selection that we would find under constant misreporting rates.

As a distinct exercise, we also calculate the rates of underreporting at the top of the consumption distribution that would be necessary to yield zero selection among the eligible (Appendix B). Overturning our conclusions requires a degree of underreporting that we view as implausible, such as "false-negative" rates of 50 percent in the top quarter of the consumption distribution. Though misreporting of transfer receipt is an important phenomenon, it is unlikely to explain our results.

Transfer Eligibility. We have already reclassified simulated-ineligible recipients as eligible to probe the sensitivity of our results to mismeasured eligibility. Another check in this spirit is to estimate the extent of selection, as in Equation 4, with rich controls for demographics and economic status so that we compare recipients and non-recipients who are observably similar and thus are likelier to have the same true eligibility status.[17] Adding these controls reduces, but does not eliminate, selection into receipt for most transfers (Appendix Figure A3). We also find similar estimates of self-targeting in demographic groups with near-certain eligibility (Appendix Figure A6). An improved measure of eligibility is thus unlikely to overturn a substantial role for take-up in selection into receipt.

Income and Consumption. Income is poorly measured at the bottom of the distribution, and while there is also mismeasurement in consumption, it appears less severe than for income (Meyer and Sullivan, 2003; Brewer et al., 2017). How is this mismeasurement likely to affect our results?

One story of intentional misreporting likely pushes in the opposite direction of our results. In this story, transfer recipients have incentives to underreport income so as to maintain eligibility, and they may do so in any quasi-official setting, including in surveys. Such incentives could apply less strongly to consumption, and to nonrecipients. All else equal, transfer recipients would thus appear *positively* selected on consumption given income. This is opposite to what we find.

We evaluate concerns about income misreporting by predicting household income using other

[17]We condition on covariates used in eligibility rules for at least one transfer program. These are the household's state of residence, household size and composition, income, earnings, ages of household members, disability status, unemployment duration and reason, and basic measures of wealth (value of any automobiles and liquid assets).

labor-market data such as weekly hours and occupation.[18] This approach provides an external check against issues in reporting that are specific to transfer recipients, assuming they report these other labor-market variables correctly. Appendix Figure A9 show that controlling for predicted income, in addition to households' reported income, has a modest impact on our results. Even this modest attenuation may reflect the informational content of occupation and education as proxies for the persistent component of income.

To examine measurement error in consumption, we show in Appendix Table A6 that similar selection patterns appear when only looking at categories of consumption deemed "well-measured" in Meyer and Sullivan (2023). For instance, on average, Medicaid recipients consume 36 percent less in housing, 25 percent less in vehicles, and 24 percent less in food at home than similar-income people who are not on Medicaid. Similar patterns also manifest in PSID data on durable-goods ownership, such as whether the household owns a home, car, computer, and the number of rooms or presence of air conditioning in the home (Meyer and Sullivan, 2012), as we show in Appendix Table A7. SNAP recipients are 16 percentage points less likely to own a home, 13 percentage points less likely to own a car, and 12 percentage points less likely to own a computer than similar-income people not on SNAP. The consistency of selection across measures, some of which seem truly unlikely to suffer from meaningful mismeasurement, bolsters our findings.

Lifetime Income. Inferring lifetime income from "snapshots" is challenging (Haider and Solon, 2006). If one thought that mismeasurement of lifetime income for people with fewer years-in-sample imparts a systematic bias in our results, then the estimated extent of selection into transfers may "drift" up or down as one examines selection for households with more or fewer years-in-sample. We re-estimate the predictive effects of transfer receipt on lifetime rank as in Equation 4, retaining only individuals with progressively more years-in-sample.

Appendix Figure A7 does find that selection on lifetime income is modestly weaker among households with many years-in-sample: results attenuate by at most two rank points, depending on

[18] Using the March Supplements to the Current Population Survey that match our PSID data years, we estimate Poisson regression models of individual income using occupation, industry, weeks per year, weekly hours, self-employment, in addition to basic demographic information. We then apply these predicted incomes to our PSID data.

the program. Yet this "drift" effect also appears as strong for consumption, suggesting the source of the phenomenon is sample attrition in the PSID and not lifetime-income estimation.[19] Altogether, measurement error in lifetime income is unlikely to drive our findings.

3.6 Extensions

Welfare Dependence. Our results so far are consistent with both transfer receipt as a "tag" of exogenous earnings ability and with endogenous dynamic responses to receipt ("welfare dependence"). Appendix Figure A8 presents suggestive evidence that advantageous selection mostly reflects tagging rather than dependence. We show that, among *current* nonrecipients of a given transfer, *future* recipients have on average a lower current-consumption rank than *future* nonrecipients with similar current incomes.[20] Such patterns hold even when one performs this comparison in the distant future, which argues against the possibility that households strategically reduce their consumption ahead of receipt. By comparison, advantageous selection in distant-future transfer receipt is quite consistent with the view of current receipt as a tag of permanent earnings ability.

Decomposition. Appendix A presents an accounting decomposition of the contributions of income mobility, eligibility, and take-up to the consumption incidence and lifetime incidence of transfers, starting from their current incidence. The decomposition builds on Brewer et al. (2020) by distinguishing between eligibility and take-up. Consistent with Figure 1, our decomposition results suggest a central role for self-targeting in concentrating incidence at the bottom of the distributions of consumption and lifetime income.

Selection Over Time. Our data span 1997 to 2019, allowing us to address how the U.S. safety net has evolved over this period. We estimate a version of Equation 4 that allows for year-specific coefficients on transfer receipt. To allow us to describe broad trends, we "stack" the data over programs and include program-specific controls for current income. Across our eight transfer programs, we see little change in selection on consumption over time, but a considerable

[19]While attrition from the PSID is another potential source of concern, it is not obvious that households that remain in the sample for 20 years are a more representative sample than those who remain for 10 years, say, before attriting.

[20]The regression of interest is $\overline{R}_{it} = \alpha_{ct} + \beta D_{i,t+k} + f(R_{it}) + u_{it}$ within the subsample such that $D_{it} = 0$, for $k > 0$.

intensification of selection on lifetime income (Appendix Figure A11).

4 Welfare Analysis of Automatic Transfers

This section estimates the social benefits of self-targeting in voluntary transfers and compares these estimates to bounds on the social costs of ordeals. We do so with a model of optimal redistribution that yields a formula for the money-metric social value of self-targeting, which we calibrate using our data and estimates above. All proofs of theoretical results are in Appendix C.

4.1 Motivating Example

Suppose there are 100 people at an income level z. Of these, 50 take up a voluntary transfer, receiving benefit $\$B$ and paying a hassle cost to do so. Of the 50, one enrolls but is exactly indifferent to not enrolling. The remaining 50 people do not take up because their ordeal cost exceeds the benefit. The government then reduces the voluntary benefit by \$1 among the 50 enrollees, which also causes the indifferent person to drop out of the transfer, saving the government an additional $\$B - 1$. The government gives all 100 people $(49 + B)/100$ dollars, which is the per-capita fiscal savings.

Let α_{AT}, α_{NT}, and α_{C} respectively be the social marginal welfare weights for the 49 inframarginal always-takers, the 50 inframarginal never-takers, and the lone complier. We focus on the case where α_{AT} exceeds α_{NT}, since we show above that always-takers have lower consumption than never-takers with the same income.

Ignoring labor-supply responses, the welfare effect of the reform is, to first order:

$$\Delta W = \underbrace{49 \times \alpha_{\text{AT}} \times \left(\frac{49 + B}{100} - 1\right)}_{\text{Welfare impact on always-takers}} + \underbrace{50 \times \alpha_{\text{NT}} \times \frac{49 + B}{100}}_{\text{Welfare impact on never-takers}} + \underbrace{1 \times \alpha_{\text{C}} \times \frac{49 + B}{100}}_{\text{Welfare impact on compliers}} . \tag{5}$$

By the envelope theorem, the only first-order welfare impact on the one complier is given by the increase in the automatic transfer, as they were indifferent prior to the reform.[21]

The key trade-off is between the fiscal savings (B) and the loss of self-targeting ($\alpha_{\text{AT}}/\alpha_{\text{NT}}$). If α_{AT} is larger than α_{NT}, then for a sufficiently small B, the welfare effect ΔW is negative, as $\frac{49+B}{100} < 1$. At a higher B, marginal fiscal savings could exceed the lost value of self-targeting. It is equivalent to describe the trade-off as between self-targeting and ordeal costs, as by the envelope theorem, marginal fiscal externalities equal the ordeal costs to marginal recipients. This logic also emerges within a model of optimal redistribution, as we now show.

4.2 Model Setup

Household's Problem. Each household has a multidimensional type $\theta = (w, \kappa)$ distributed according to p.d.f. μ. The parameter $w \in \mathbb{R}^+$ is the household's wage, encoding their productivity, and $\kappa \in \mathbb{R}^+$ is their cost of applying to the transfer. Households choose how much labor $l \in \mathbb{R}^+$ to supply to generate pre-tax income $z = wl$.

Pre-tax income is taxed according to the non-linear schedule $T(z) : \mathbb{R}^+ \to \mathbb{R}$ and the remainder is consumed. We model automatic transfers as reducing the income tax liability: An automatic transfer of \$1 reduces the income tax by \$1. There is also a (non-automatic) transfer with non-linear schedule $S(z) : \mathbb{R}^+ \to \mathbb{R}^+$. Negative taxes are possible. Households choose to apply for this benefit after observing their private cost of application κ.

To rule out bunching and ensure well-defined elasticities, we assume households choose their labor supply before they see their realization of κ.[22] As shorthand, we use $\mathbb{1}_S = 1[S(z) \geq \kappa]$ as an indicator for whether, ex post, the individual applies for and receives the transfer. To rule out income effects, we assume households have quasi-linear utility in cash and cash-equivalent transfer dollars and disutility of work hours $v(l) = v(z/w)$. We also assume no income effects are present with respect to the transfer $S(z)$. For each household, the choice of l is one-to-one with income z. Therefore, we model the household's labor supply choice as a direct choice of z. As such, each

[21]The welfare impact on compliers is generally nonzero because the reform entails both a cut to the voluntary transfer and an increase in the automatic transfer. Compliers are indifferent to the former but inframarginal to the latter.

[22]Here we abstract away from eligibility rules, which can be viewed as raising take-up costs (possibly to infinity).

household solves the program:

$$\max_{z}\left\{z-T(z)-v(z/w)+\int_{0}^{S(z)}(S(z)-\kappa)\mu(w\mid\kappa)d\kappa\right\}. \tag{6}$$

Suppressing the dependence on the wage w for clarity, the household's optimal choice $z^* = z^*(w)$ leads to ex-post consumption is $c^* = z^* - T(z^*) + \mathbb{1}_S(S(z^*) - \kappa)$. The planner maximizes the weighted sum of ex-post individual utilities, which are given by

$$V(\theta) = z^* - T(z^*) - v(z^*/w) + \mathbb{1}_S(S(z^*) - \kappa)\,. \tag{7}$$

Government's Problem. The government chooses tax and transfer schedules $T(\cdot)$ and $S(\cdot)$ to maximize utility summed across individuals according to type-specific Pareto weights ($\alpha(\theta)$):

$$\max_{T,S}\int_{\Theta}\alpha(\theta)V(\theta)d\mu(\theta),$$

subject to a balanced-budget constraint:

$$\int_{\Theta}\left[T\left(z(\theta)\right)-\mathbb{1}_S S\left(z(\theta)\right)\right]d\mu(\theta)=0 \tag{8}$$

and to household optimization. The welfare weights may capture, for instance, a higher social value of transferring to people with low consumption. As is usual in this setting, we assume that social marginal welfare weights decrease with wages, all else equal: $\frac{\partial}{\partial w}\alpha(\theta) < 0$. Because $\frac{\partial z^*(w)}{\partial w} > 0$, this assumption is tantamount to assuming that welfare weights are higher for those with lower pre-tax income z^*.

Our results are delineated by the following relationship between social welfare weights and underlying type parameters, consistent with our empirical definition of self-targeting (Equation 3).

Definition 1. *We say that take-up is* advantageously self-targeting *when* $\frac{\partial}{\partial \kappa}\alpha(\theta) < 0$.

Advantageous self-targeting means that, all else equal, households with low application costs to the transfer have higher marginal welfare weights.[23] That is, the households first to apply for transfers at any income level have high social marginal welfare weights. If recipients are negatively selected on consumption given income (as in Section 3), advantageous self-targeting holds, assuming social welfare weights that decrease in consumption. On the other hand, if ordeals perversely screen out needier people, advantageous self-targeting does not hold.

We define $M(z)$ and $m(z)$ as respectively the cumulative distribution and density functions of transfer application costs for households at income z: $M(z) = Pr(S(z) \geq \kappa)$ for all those earning income z, and similarly $m(z) = \frac{d}{dS(z)}M(z)$. We assume the distributions are continuous on $\kappa \geq 0$, that is, there are no mass points. These distributions are endogenous to tax and transfer schedules, and for brevity, we suppress the dependence on $S(z)$. Moreover, we define $h(z)$ as the mass of types that choose pre-reform income z. As noted above, this is one-to-one with primitive w.

4.3 Policy Reform

To evaluate the welfare effects of automatic transfers, we consider a budget-neutral reform that marginally shifts resources from the voluntary transfer to the automatic transfer.

The Reform. We cut the voluntary transfer and raise the automatic one (i.e., cut taxes) so that each income level is compensated on average. The transfer amount is reduced by ds at all incomes. At each income z, taxes are reduced by $\tau(z) = M(z)ds$. Marginal rates thus change by $\tau'(z) = \frac{d}{dz}M(z)ds$ at z. Fiscal savings from marginal recipients are redistributed as a lump sum $E_z\left[(S(z) + ds)m(z)\right]$. The reform does not redistribute between incomes but generically affects labor supply. It then pays for any decrease in revenue due to the labor-supply response via lump-sum taxes. Figure 4 visualizes the reform.

Of course, there are infinite possible reforms that modify both automatic and voluntary transfers while remaining budget-neutral. This reform corresponds to a level reduction in the transfer schedule, offset by the unique tax change that avoids between-income redistribution. Intuitively, each income group is "made whole" from the automatic–voluntary transfer swap. As mentioned

[23] Definition 1 follows Saez and Stantcheva (2016) in abstracting away from any specific rationale.

Figure 4: Automatic Versus Voluntary Transfers

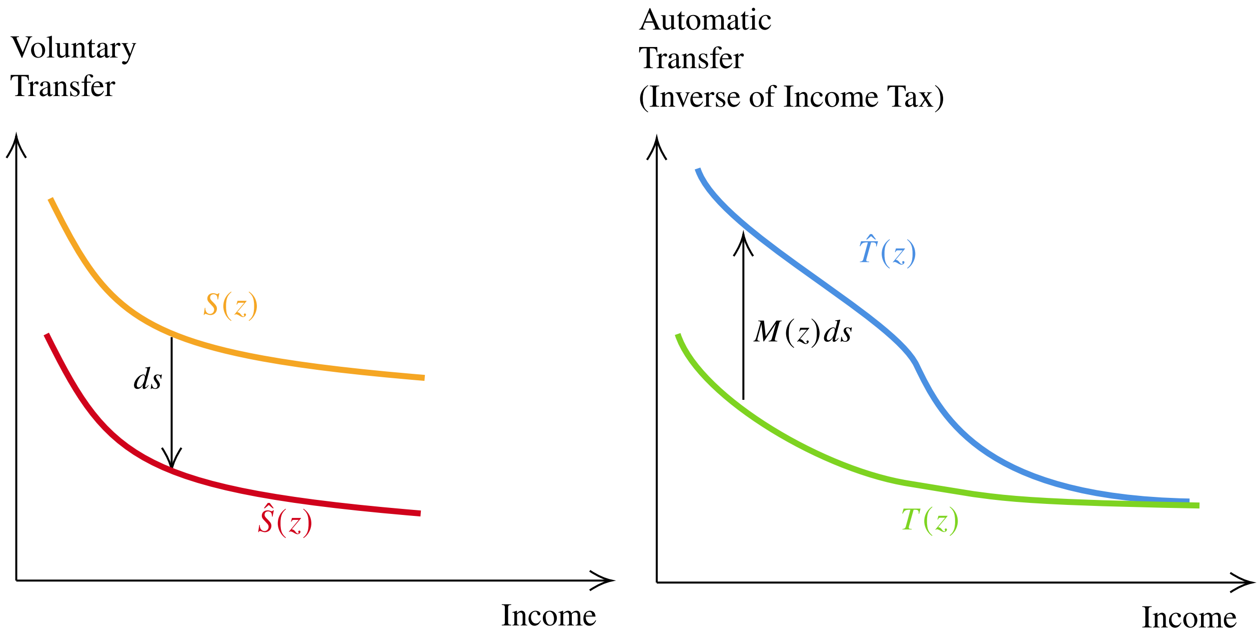

Notes: This figure shows the reform. The left diagram shows the reduction in the voluntary transfer, and the right diagram shows the corresponding increase in the automatic transfer. $S(z)$ and $T(z)$ show the pre-reform voluntary and automatic transfer schedules, whereas $\hat{S}(z)$ and $\hat{T}(z)$ show the post-reform schedules. Notably, since the automatic transfer is multiplied by $M(z)$, which is small at the top of the income distribution, the post-reform automatic transfer schedule tends toward the pre-reform schedule at the top of the distribution. Moreover, when the automatic transfer increases, note that this corresponds to a reduction in the income tax liability.

above, the reform can be envisioned as changing a fully-voluntary transfer into one with a small automatic transfer with a top-up provided upon application. The next proposition, proven using a perturbation-based argument (Jacquet and Lehmann, 2014), presents welfare formulas for this marginal shift toward automatic transfers.

Proposition 1. *The welfare effect of the reform is*

$$\begin{aligned} dW = & \underbrace{ds \int_z M(z) \left(E_\kappa \left[\alpha(z,\kappa) \right] - E_{\kappa \le S(z)} \left[\alpha(z,\kappa) \right] \right) h(z) dz}_{\textit{lost value of self-targeting}} \\ & + \underbrace{E_z[S(z)] m(z) E_{z,\kappa} \left[\alpha(z,\kappa) \right]}_{\textit{fiscal savings from marginals}} \\ & + \underbrace{ds \int_z \frac{M'(z) z \epsilon^z}{1 - T'(z)} \left(\frac{d}{dz} (S(z) M(z)) - T'(z) \right) E_{z',\kappa} \left[\alpha(z',\kappa) \right] h(z) dz}_{\textit{labor-supply effect (i)}} \\ & + \underbrace{ds \int_z \frac{M'(z) z \epsilon^z}{1 - T'(z)} S'(z) \mathrm{Cov}_k \left[\alpha(z,\kappa), \mathbb{1}(S(z) \ge \kappa) \right] h(z) dz,}_{\textit{labor-supply effect (ii)}} \end{aligned} \tag{9}$$

where $\epsilon^z = -\frac{\partial z}{\partial \tau'} \frac{1 - T'(z)}{z}$ is the elasticity of income with respect to a small change τ' in the marginal tax rate of those with initial income z.[24]

The first term in Equation 9 shows that, in moving toward automatic transfers, the government can only transfer $M(z)S(z)$ to all at income z rather than $S(z)$ to those who take up at this income. If take-up is correlated with need conditional on income, then making transfers automatic will have undesirable redistributive properties. The fiscal-savings term captures that some people no longer take up the program, since the benefit no longer exceeds their private ordeal cost. By the envelope theorem, the fiscal-savings term equals the cost of ordeals among those whose take-up decision changes when the benefit level is reduced. There are also two labor-supply terms, which appear because the reform changes marginal tax rates and hence induces a labor-supply response. The first is standard, measuring the fiscal impact of behavioral responses to changes in marginal "tax-and-transfer" rates. The second stems from our timing assumption that households choose labor supply before observing their application cost.

When self-targeting is advantageous and $S(z)$ is positive, the first term in Equation 9 is negative: the government must forgo some social benefits of self-targeting in moving toward automatic

[24] In Appendix C, we give conditions under which the elasticity of income is properly defined.

transfers. The sum of the labor supply effects in Equation 9 is also negative when the tax system is optimal and take-up decreases in income ($M'(z) < 0$). Under these assumptions, moving toward automatic transfers requires higher marginal tax rates to offset the cut to the voluntary transfer, thereby reducing labor supply.[25] See Appendix C for formal statements of these propositions.

The reform considered here differs importantly from reforms to ordeals. We instead take the ordeal as given and reallocate resources between voluntary and automatic transfers. The welfare analysis of ordeal reforms weighs the change in ordeal costs to *inframarginal* recipients against the fiscal externalities from changes in take-up (e.g., Finkelstein and Notowidigdo, 2019). By comparison, the welfare analysis of reallocating resources contrasts the value of transfers to inframarginal recipients against ordeal costs to *marginal* recipients. A virtue of our reform is that the welfare-relevant measure of ordeal costs is obtained by the envelope theorem. These are otherwise difficult to measure, as Finkelstein and Notowidigdo (2019) observe.

Using the envelope theorem to infer ordeal costs assumes that households make transfer take-up decisions optimally. However, research has found non-optimizing behavior in take-up (Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019; Anders and Rafkin, 2022). Importantly, the optimizing assumption works against our conclusion, as it yields upper bounds on ordeal costs. If households do not take-up because of mis-optimization or a lack of information, then ordeal costs would be smaller than what is implied by equating them to marginal benefits. Transfers would then achieve advantageous self-targeting in ways that do not spend real resources on ordeals.

In Appendix C, we provide a formula for the welfare effects of a more general class of transfer reforms. This formula accounts for redistribution both between and across incomes, fiscal savings from marginal recipients, and labor-supply effects. We also consider non-marginal changes to voluntary transfers, in which case we cannot apply the envelope theorem to reveal ordeal costs.

[25]These terms are broadly related to prior work. In their analysis of disability insurance, Diamond and Sheshinski (1995) observe that having the receipt rate rising in the level of disability "is sufficient to make a disability program desirable, provided the marginal utility of consumption of non-workers exceeds that of workers at the optimum without a disability program." The labor-supply terms are connected to the model of Nichols and Zeckhauser (1982). Consider, in particular, a reform that expands the voluntary transfer and cuts taxes, reversing our analysis above. Through screening, much as in Nichols and Zeckhauser (1982), this reform can both raise labor supply and redistribute toward people with higher welfare weights.

4.4 Quantification

The trade-off between self-targeting and ordeal costs means the welfare effect of shifting between voluntary and automatic transfers is ambiguous. To estimate the welfare effect, we calibrate Equation 9 using our estimates and several external inputs. We discuss where our welfare calculations are more and less sensitive to assumptions, given the room for reasonable disagreement in the calibration of the external inputs and other simplifying assumptions.

Calibration. From the PSID, we obtain receipt rates $M(z)$, average benefits $S(z)$, and the income distribution $h(z)$. To obtain the density of application costs $m(z)$, we use the result that $m(z) = \eta M(z)/S(z)$ for some elasticity η of take-up with respect to the benefit level. We set the elasticity η to 0.6, the upper end of the range in Krueger and Meyer (2002)'s review of take-up elasticities for unemployment insurance.[26] We also assume a constant elasticity of taxable income $\varepsilon = 0.3$, calibrated from Saez et al. (2012). We calibrate a piecewise-linear tax schedule $T(z)$ using effective average marginal tax rates that incorporate federal and state taxes on individual income and payroll (Congressional Budget Office, 2015).

We assume welfare weights are described by isoelastic social preferences over consumption with constant relative risk aversion parameter γ. In our primary estimates, we calibrate $\gamma = 2$ given the lack of household risk aversion in our model (Chetty and Finkelstein, 2013). A higher γ implies a stronger redistributive motive and would thus favor voluntary transfers. Fixing γ, we use the joint distribution of income, consumption, and transfer receipt to calculate the average welfare weight $E_{\kappa \leq S(z)}[\alpha(z, \kappa)]$ on recipients at income z.

Results. Panel A of Table 4 reports our primary estimates of the welfare effects of reallocating resources from a given voluntary transfer to an automatic one. Column 1 shows the social costs from giving up some self-targeting (the first term in Equation 9), while Column 2 shows the fiscal savings on marginal households who exit the transfer (the second term in Equation 9). Due to

[26]The upper end of this range is conservative for our analysis in the sense that it favors automatic transfers. McGarry (1996) likewise finds a take-up elasticity for SSI benefits of 0.5. We know of no more-recent estimates of take-up elasticities with respect to benefit level for other U.S. transfers. Additional estimates would be of clear value. We also assume the elasticity is constant with respect to income z.

the envelope argument explained above, Column 2 can also be interpreted as the social savings on ordeal costs among marginal households. Column 3 shows the labor-supply effect (the sum of the third and fourth terms in Equation 9). Column 4 shows the total effect of the reform.

To take SNAP as an example, we find that SNAP more automatic forgoes some of the social benefits of self-targeting. On the margin, self-targeting is worth about 10.5 cents per dollar of SNAP. In per-recipient terms, self-targeting yields a social benefit of about $500, an amount that we see as obviously unlikely to be offset by ordeal costs. By comparison, the government saves 5.9 cents per SNAP dollar from marginal households who exit when the voluntary benefit is cut. Finally, the automatic transfer increases marginal tax rates, which reduces labor supply and imposes a one-cent fiscal externality. Together, the net effect of making SNAP more automatic on the margin is a net social loss of 6.7 cents per SNAP dollar. The magnitudes of welfare effects seem consequential, especially for a budget-neutral reform.

Overall, we find a stark trade-off between the social benefits of self-targeting and the social costs of ordeals. Looking across transfers, social benefits are often equal to or greater than our upper-bound estimates of social costs. By consequence, the net social gains from making transfers automatic tend to be be negative or small, and they are not well-approximated by the social savings on ordeal costs alone. That result is reflected in the dollar-weighted average, which shows that on the margin the forgone social benefits of self-targeting actually exceed the social costs of ordeals. Overall, the loss of self-targeting appears to be a credible argument for the status quo of voluntary transfers and against automatic-transfer reforms.

There is also considerable heterogeneity in welfare effects across programs. Ordeals in some transfers seem ineffectual: that is, they have social costs but do not induce socially valuable self-targeting. For example, our results suggest potential welfare gains from universal free school meals or making WIC automatic. Automatic benefits, by contrast, appear most costly in housing-assistance programs. Importantly, these programs have severe ordeals: low-quality and constrained choices, as well as long waiting lists. Our framework is thus not uniformly favorable towards ordeals but makes finer distinctions according to how effective an ordeal is in causing self-targeting.

Table 4: Welfare Effects of Making Transfers Automatic (Cents per Transfer Dollar)

	(1) Self-Targeting	(2) Upper Bound on Ordeals	(3) Labor-Supply Effects	(4) Total
Panel A: Primary Estimates				
Dollar-Weighted Average	**-6.1**	**5.7**	**-0.9**	**-1.4**
SNAP	-10.5	5.9	-1.0	-5.6
Medicaid	-4.7	8.6	-1.4	2.5
Housing Assistance	-11.0	3.1	-0.5	-8.4
TANF	-1.5	0.6	-0.1	-1.0
SSI	-2.7	3.5	-0.4	0.3
School Lunch	2.5	5.7	-0.9	7.4
WIC	-0.1	2.3	-0.4	1.8
LIHEAP	-0.7	2.3	-0.3	1.3
Panel B: Sensitivity (Dollar-Weighted Average)				
SWF curvature $\gamma = \frac{1}{2} \times$ primary estimate	-2.3	5.7	-0.9	2.5
SWF curvature $\gamma = 2 \times$ primary estimate	-10.6	5.7	-0.9	-5.8
SWF over lifetime income	-5.8	5.7	-0.9	-1.0
Take-up elasticity $\eta = \frac{1}{2} \times$ primary estimate	-6.1	2.8	-0.9	-4.2
Take-up elasticity $\eta = 2 \times$ primary estimate	-6.1	11.4	-0.9	4.3
Elasticity of taxable income $\varepsilon = \frac{1}{2} \times$ primary estimate	-6.1	5.7	-0.4	-0.9
Elasticity of taxable income $\varepsilon = 2 \times$ primary estimate	-6.1	5.7	-1.8	-2.2

Notes: This table reports estimates of the welfare effects of the reform, which marginally reduces the voluntary transfer to make it automatic. We calibrate the welfare weights by assuming a CES social welfare function with curvature parameter $\gamma = 2$. We calibrate the fiscal cost of marginals by assuming the takeup elasticity is $\eta = 0.5$. We calibrate the elasticity of taxable income at $\varepsilon = 0.3$. All columns report the money-metric welfare gains in cents per transfer dollar. Columns correspond to the terms of Equation 9, where we divide each term by the average welfare weight to yield a money-metric interpretation.

Sensitivity Analysis and Discussion. Panel B examines the sensitivity of our results to the calibrated parameters. The more society cares more about redistribution, the larger the welfare losses from forgoing self-targeting in transfers. Put another way, automating transfers is likely to be socially desirable only when society cares relatively less about the poor (i.e., it has a lower γ). Redefining the social welfare function to be in terms of lifetime income rather than consumption does not much affect the welfare analysis.

The take-up elasticity is critically important to the overall fiscal costs. If take-up is more highly responsive to the benefit level than we expect, this would imply larger ordeal costs on the margin and thus could motivate automatic transfers. Results are less sensitive to the elasticity of taxable income. Across these permutations of our analysis, self-targeting remains a quantitatively important advantage of voluntary transfers. Indeed, self-targeting typically eliminates most if not all of the social savings on ordeals, even at upper-bound values for ordeal costs.

This welfare analysis has several limitations. First, it does not account for differences in the government's administrative costs between voluntary and automatic transfers. Little is known about the appropriate values for these costs (Isaacs, 2008), but it is reasonable to suspect they favor automatic transfers. Second, we ignore behavioral responses to transfers beyond take-up and labor supply, such as cross-program enrollment spillovers or dynamic incentives for human-capital investment. Third, we assume homogeneous labor supply and take-up elasticities. Heterogeneity in elasticities could shift our conclusions in either direction. For instance, non-recipients may have lower take-up elasticities than recipients, reducing the fiscal cost from marginal recipients.

5 Conclusion

A convincing body of empirical research has studied many specific ordeals in transfer programs. It finds mixed evidence that, on the margin, these ordeals have favorable selection properties. Taken as a whole, this literature would seem to have radical implications for the design of social safety nets: Why do governments hassle people by making them ask for help, if those who do not ask are no less in need? Why not just send help automatically?

What is true among the complier population for the studied ordeals does not, we find, generalize to always- and never-takers of transfer programs. We show transfer recipients are, relative to non-recipients, strongly and consistently negatively selected on consumption and lifetime income given income. This selection mostly reflects take-up among the eligible rather than eligibility rules.

Such self-targeting might rescue the case for voluntary take-up. To determine if it does, we quantify the social trade-off between self-targeting and ordeal costs. In particular, we imagine reforms that incrementally shift redistribution from voluntary to automatic transfers. Calibrating a welfare-effect formula using our empirical estimates, we find that the social benefits from self-targeting generally equal or exceed upper-bound estimtes of the social costs of ordeals. There would be, by consequence, social losses from making transfers automatic overall. However, some transfers inflict ordeal costs but achieve minimal self-targeting, and in these programs, the U.S. could indeed achieve considerable welfare gains by eliminating the need to sign up.

References

Alatas, Vivi, Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A Olken, and Rema Hanna, "Self-Targeting: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2016, *124* (2), 371–427.

Anders, Jenna and Charlie Rafkin, "The Welfare Effects of Eligibility Expansions: Theory and Evidence from SNAP," Working Paper 2022.

Arbogast, Iris, Anna Chorniy, and Janet Currie, "Administrative Burdens and Child Medicaid Enrollments," Working Paper 30580, National Bureau of Economic Research 2022.

Auerbach, Alan J, Laurence J Kotlikoff, and Darryl Koehler, "U.S. Inequality and Fiscal Progressivity: An Intragenerational Accounting," *Journal of Political Economy*, forthcoming.

Bartels, Charlotte and Dirk Neumann, "Redistribution and Insurance in Welfare States Around the World," *The Scandinavian Journal of Economics*, 2021, *123* (4), 1116–1158.

Bengtsson, Niklas, Bertil Holmlund, and Daniel Waldenström, "Lifetime Versus Annual Tax-and-Transfer Progressivity: Sweden, 1968–2009," *The Scandinavian Journal of Economics*, 2016, *118* (4), 619–645.

Besley, Timothy and Stephen Coate, "Workfare Versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs," *American Economic Review*, 1992, *82* (1), 249–261.

Bhargava, Saurabh and Dayanand Manoli, "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment," *American Economic Review*, 2015, *105* (11), 3489–3529.

Blank, Rebecca M and Patricia Ruggles, "When Do Women Use Aid to Families with Dependent Children and Food Stamps? The Dynamics of Eligibility Versus Participation," *Journal of Human Resources*, 1996, pp. 57–89.

Blundell, Richard, Michael Graber, and Magne Mogstad, "Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family," *Journal of Public Economics*, 2015, *127*, 58–73.

Brewer, Mike, Ben Etheridge, and Cormac O'Dea, "Why Are Households that Report the Lowest Incomes So Well-Off?," *Economic Journal*, 2017, *127* (605).

_ , Robert Joyce, Tom Waters, and Joseph Woods, "A Method for Decomposing the Impact of Reforms on the Long-Run Income Distribution, with an Application to Universal Credit," *Economics Letters*, 2020, *192*, 109230.

Caspersen, E and Gilbert Metcalf, "Is a Value Added Tax Regressive? Annual Versus Lifetime Incidence Measures," *National Tax Journal*, 1994, *47* (4), 731–746.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly, "Inference on Counterfactual Distributions," *Econometrica*, 2013, *81* (6), 2205–2268.

Chetty, Raj and Amy Finkelstein, "Social Insurance: Connecting Theory to Data," in "Handbook of Public Economics," Vol. 5, Elsevier, 2013, pp. 111–193.

Chiappori, Pierre-André and Bernard Salanie, "Testing for asymmetric information in insurance markets," *Journal of political Economy*, 2000, *108* (1), 56–78.

Citro, Constance Forbes and Robert T Michael, *Measuring Poverty: A New Approach*, National Academy Press, 1995.

Congressional Budget Office, "Effective Marginal Tax Rates for Low- and Moderate-Income Workers in 2016," Report 2015.

Currie, Janet, "The Take-Up of Social Benefits," in Alan J Auerbach, David Card, and John M Quigley, eds., *Public Policy and the Income Distribution*, Russell Sage Foundation, 2006.

_ and Firouz Gahvari, "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature*, 2008, *46* (2), 333–83.

Davern, Michael E, Bruce D Meyer, and Nikolas K Mittag, "Creating Improved Survey Data Products Using Linked Administrative-Survey Data," *Journal of Survey Statistics and Methodology*, 2019, *7* (3), 440–463.

Davies, James, France St-Hilaire, and John Whalley, "Some Calculations of Lifetime Tax Incidence," *American Economic Review*, 1984, *74* (4), 633–649.

Deshpande, Manasi and Yue Li, "Who is Screened Out? Application Costs and the Targeting of Disability Programs," *American Economic Journal: Economic Policy*, 2019, *11* (4), 213–48.

Diamond, Peter and Eytan Sheshinski, "Economic Aspects of Optimal Disability Benefits," *Journal of Public Economics*, 1995, *57* (1), 1–23.

Duclos, Jean-Yves, "Modelling the Take-up of State Support," *Journal of Public Economics*, 1995, *58* (3), 391–415.

Ericson, Keith Marzilli, Timothy J Layton, Adrianna McIntyre, and Adam Sacarny, "Reducing Administrative Barriers Increases Take-up of Subsidized Health Insurance Coverage: Evidence from a Field Experiment," Working Paper 30885, National Bureau of Economic Research 2023.

Eurofound, "Access to Social Benefits: Reducing Non-Take-Up," 2015.

Finkelstein, Amy and Matthew J Notowidigdo, "Take-Up and Targeting: Experimental Evidence from SNAP," *Quarterly Journal of Economics*, 2019, *134* (3), 1505–1556.

Fullerton, Don and Diane Lim Rogers, *Who Bears the Lifetime Tax Burden?*, Brookings Institution Washington, DC, 1993.

— **and Gilbert E Metcalf**, "Tax Incidence," *Handbook of Public Economics*, 2002, *4*, 1787–1872.

Ganong, Peter and Jeffrey B Liebman, "The Decline, Rebound, and Further Rise in SNAP Enrollment: Disentangling Business Cycle Fluctuations and Policy Changes," *American Economic Journal: Economic Policy*, 2018, *10* (4), 153–76.

Gray, Colin, "Leaving Benefits on the Table: Evidence from SNAP," *Journal of Public Economics*, 2019, *179*, 104054.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura, "Rethinking the Welfare State," Discussion Paper DP16275, CEPR 2021.

Haider, Steven and Gary Solon, "Life-Cycle Variation in the Association Between Current and Lifetime Earnings," *American Economic Review*, 2006, *96* (4), 1308–1320.

Herd, Pamela and Donald P Moynihan, *Administrative Burden: Policymaking by Other Means*, Russell Sage Foundation, 2019.

Homonoff, Tatiana and Jason Somerville, "Program Recertification Costs: Evidence from SNAP," *American Economic Journal: Economic Policy*, 2021, *13* (4), 271–98.

Isaacs, Julia, "The Costs of Benefit Delivery in the Food Stamp Program: Lessons From a Cross-Program Analysis," 2008.

Jacquet, Laurence and Etienne Lehmann, "Optimal Nonlinear Income Taxation with Multidimensional Types: The Case with Heterogeneous Behavioral Responses," THEMA Working Papers 2014-01, Université de Cergy-Pontoise 2014.

Kleven, Henrik Jacobsen and Wojciech Kopczuk, "Transfer Program Complexity and the Take-Up of Social Benefits," *American Economic Journal: Economic Policy*, 2011, *3* (1), 54–90.

Krueger, Alan B and Bruce D Meyer, "Labor Supply Effects of Social Insurance," *Handbook of Public Economics*, 2002, *4*, 2327–2392.

Levell, Peter, Barra Roantree, and Jonathan Shaw, "Mobility and the Lifetime Distributional Impact of Tax and Transfer Reforms," *International Tax and Public Finance*, 2021, *28* (4), 751–793.

Lieber, Ethan MJ and Lee M Lockwood, "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care," *American Economic Review*, 2019, *109* (4), 1461–85.

McGarry, Kathleen, "Factors Determining Participation of the Elderly in Supplemental Security Income," *Journal of Human Resources*, 1996, pp. 331–358.

Meyer, Bruce D and James X Sullivan, "Measuring the Well-Being of the Poor Using Income and Consumption," *Journal of Human Resources*, 2003, pp. 1180–1220.

__ **and** __ , "Identifying the Disadvantaged: Official Poverty, Consumption Poverty, and the New Supplemental Poverty Measure," *Journal of Economic Perspectives*, 2012, *26* (3), 111–36.

__ **and** __ , "Consumption and Income Inequality in the US Since the 1960s," *Journal of Political Economy*, 2023, *131* (2).

__ **, Nikolas Mittag, and Robert M Goerge**, "Errors in Survey Reporting and Imputation and Their Effects on Estimates of Food Stamp Program Participation," *Journal of Human Resources*, 2020, pp. 0818–9704R2.

__ **, Wallace KC Mok, and James X Sullivan**, "The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences," Working Paper 15181, National Bureau of Economic Research 2009.

__ **,** __ **, and** __ , "Household Surveys in Crisis," *Journal of Economic Perspectives*, 2015, *29* (4), 199–226.

Mittag, Nikolas, "Correcting for Misreporting of Government Benefits," *American Economic Journal: Economic Policy*, 2019, *11* (2), 142–64.

Mullainathan, Sendhil and Eldar Shafir, *Scarcity: Why Having Too Little Means So Much*, Macmillan, 2013.

Nichols, Albert L and Richard J Zeckhauser, "Targeting Transfers Through Restrictions on Recipients," *American Economic Review*, 1982, *72* (2), 372–377.

Poterba, James M, "Lifetime Incidence and the Distributional Burden of Excise Taxes," *American Economic Review*, 1989, *79* (2), 325–330.

__ , "Is the Gasoline Tax Regressive?," *Tax Policy and the Economy*, 1991, *5*, 145–164.

Roantree, Barra and Jonathan Shaw, "What a Difference a Day Makes: Inequality and the Tax and Benefit System from a Long-Run Perspective," *Journal of Economic Inequality*, 2018, *16* (1), 23–40.

Saez, Emmanuel and Stefanie Stantcheva, "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 2016, *106* (1), 24–45.

— **, Joel Slemrod, and Seth H Giertz**, "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, 2012, *50* (1), 3–50.

Shepard, Mark and Myles Wagner, "Reducing Ordeals through Automatic Enrollment: Evidence from a Subsidized Health Insurance Exchange," Working Paper 30781, National Bureau of Economic Research 2022.

Sullivan, James X, Lesley Turner, and Sheldon Danziger, "The Relationship Between Income and Material Hardship," *Journal of Policy Analysis and Management*, 2008, *27* (1), 63–81.

Unrath, Matthew, "Targeting, Screening, and Retention: Evidence from California's Food Stamps Program," Working Paper 2021.

U.S. Department of Health and Human Services, "Unwinding the Medicaid Continuous Enrollment Provision: Projected Enrollment Effects and Policy Approaches," Issue Brief HP-2022-20, Office of Health Policy, Assistant Secretary for Planning and Evaluation 2022.

Vickrey, William S, *Agenda for Progressive Taxation*, Ronald Press Company, 1947.

Wu, Derek and Bruce D Meyer, "Certification and Recertification in Welfare Programs: What Happens When Automation Goes Wrong," Working Paper, University of Chicago 2022.

Zedlewski, Sheila and Linda Giannarelli, "TRIM: A Tool for Social Policy Analysis," Research Report, Urban Institute 2015.

Appendices for Online Publication

A Additional Tables and Figures

Figure A1: Receipt and Value of Transfer Benefits as a Function of Household Rank

Panel A: Average Total Annual Per-Capita Transfer

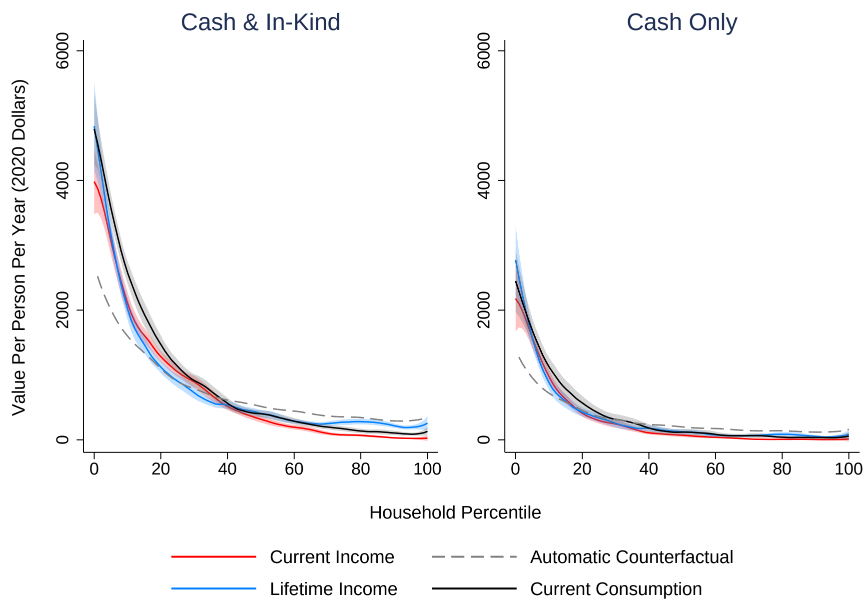

Panel B: Receipt Rate by Program

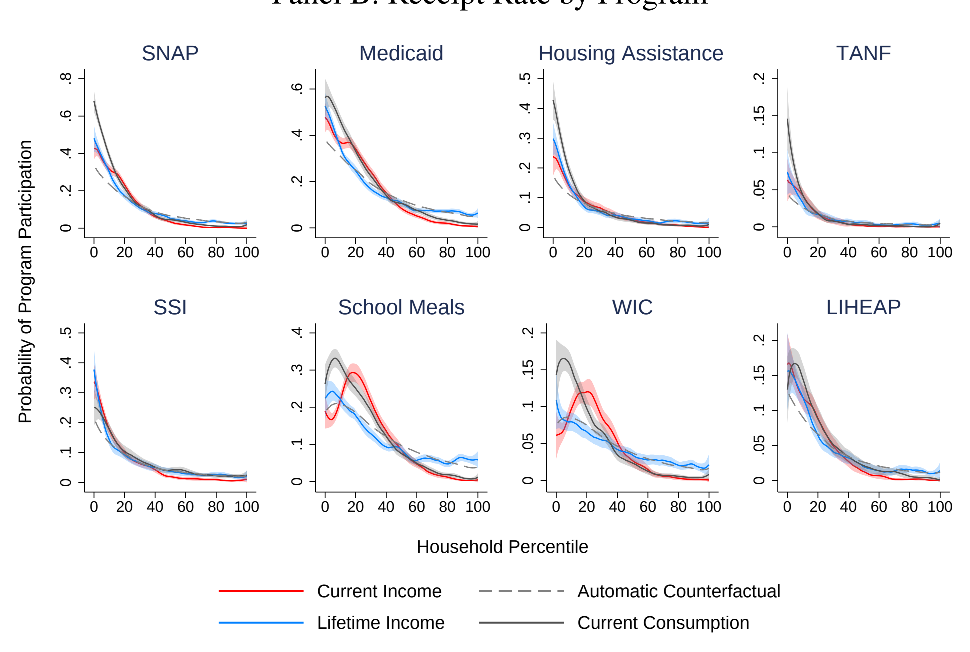

Notes: This figure displays average total annual per-capita values of benefits and receipt rates by program as functions of household ranks in the distributions of household current income, lifetime income, and current consumption. There is no equivalence scale applied to household income. The functions are estimated by local linear regressions with bandwidths of three percentiles. Shaded regions reflect bootstrapped 95-percent simultaneous confidence bands, as in Chernozhukov et al. (2013), with clustering by household.

Figure A2: Value of Transfer Benefits, by Program, as a Function of Rank

Panel A: Equivalized Household Rank

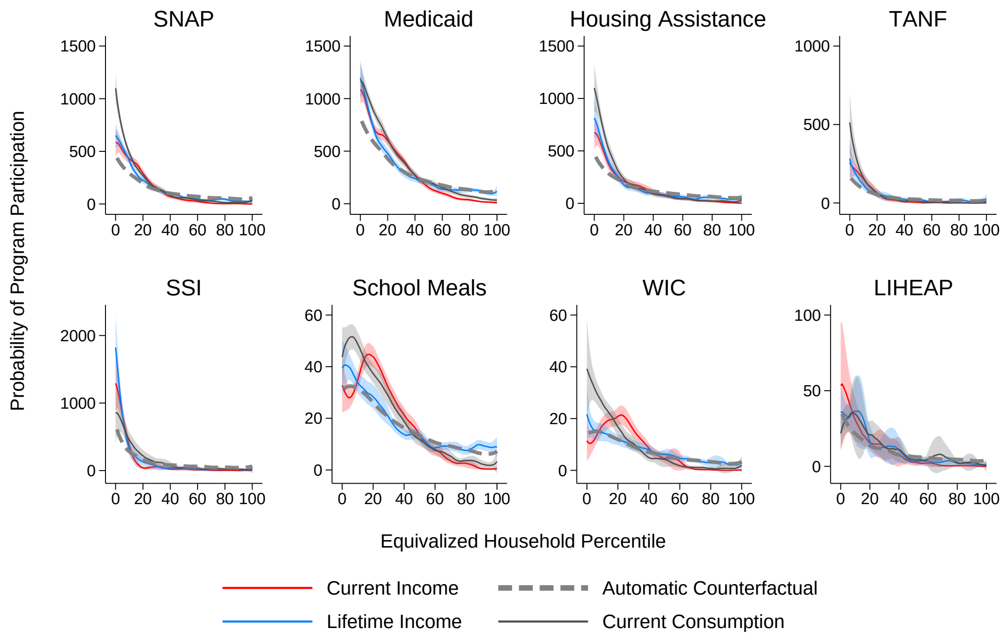

Panel B: Household Rank

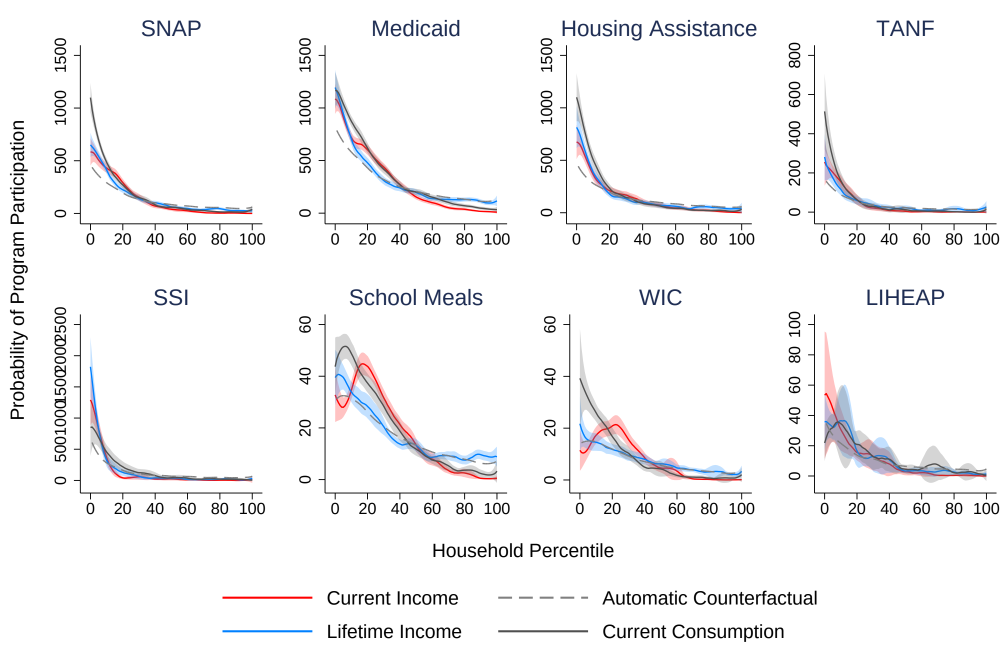

Notes: This figure displays the average total annual per-capita value of transfer benefits by program as functions of ranks in the distributions of current income, lifetime income, and current consumption. Panel A reports results for equivalized income or consumption, whereas Panel B does not adjust for differences in household size and composition. The functions are estimated by local linear regressions with bandwidths of three percentiles. Shaded regions reflect bootstrapped 95-percent simultaneous confidence bands, as in Chernozhukov et al. (2013), with clustering by household.

Figure A3: What Explains Selection into Transfer Receipt? With Controls

Panel A: Selection on Consumption

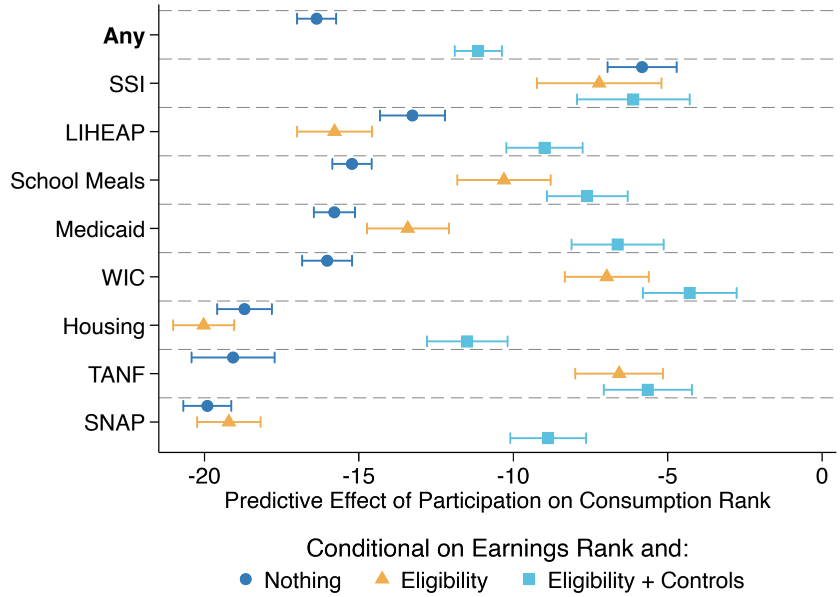

Panel B: Selection on Lifetime Income

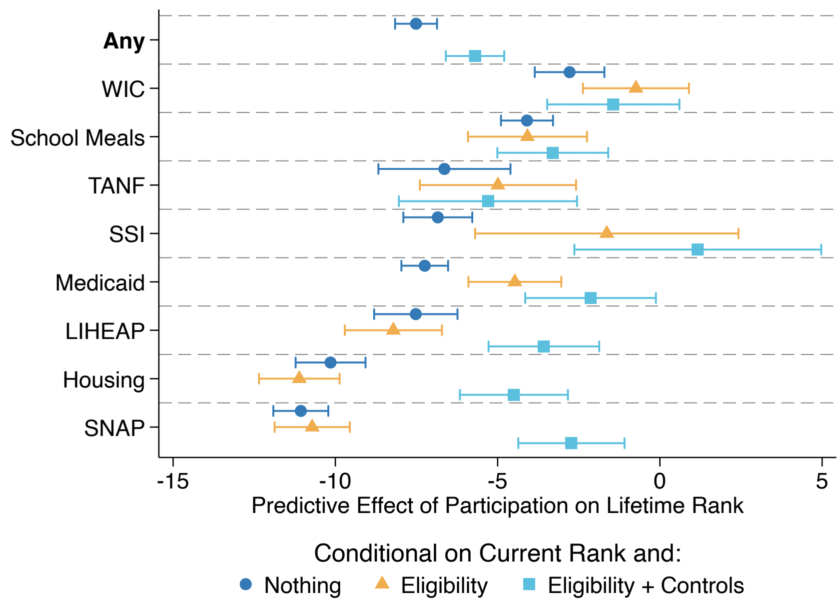

Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank (Panel A) or lifetime-income rank (Panel B), conditional on current-income rank (coefficient γ from Equation 4). For estimates represented by blue circles, we do not add additional control variables to the specification, whereas for the yellow diamonds, we estimate the regression only on people whom we simulate to be eligible. For estimates represented by the teal squares, we condition on eligibility and household demographic characteristics. The "any" row is an indicator for receipt of at least one of the eight transfers. The confidence intervals are for the 95-percent level and reflect clustered standard errors by household.

Figure A4: What Explains Selection into Transfer Receipt?: Reclassifying Simulated Eligibles

Panel A: Selection on Consumption

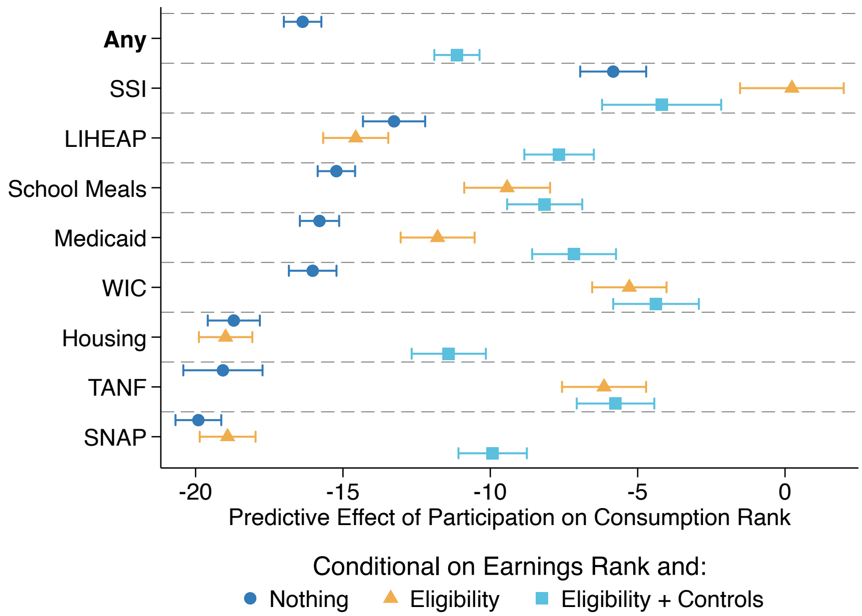

Panel B: Selection on Lifetime Income

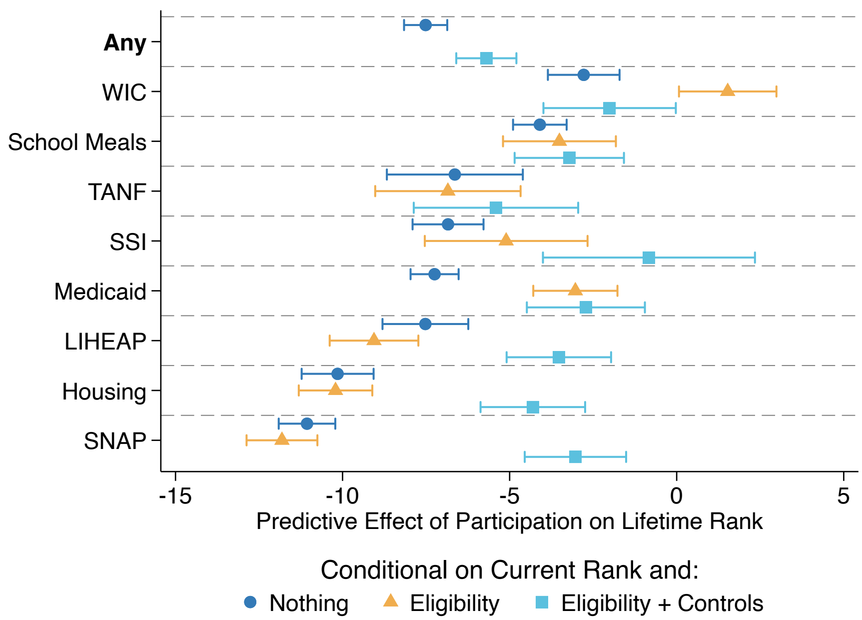

Notes: This figure displays estimates of the predictive effect of transfer benefit receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient γ from Equation 4). For estimates represented by blue circles, we add no additional control variables to the specification, whereas for the yellow diamonds, we add program eligibility. The confidence intervals are for the 95-percent level and clustering by household. In Panel B, we adapt Equation 4 by replacing the transfer indicator with indicators for the number of unique transfers received.

Figure A5: Receipt and Value of Unemployment Insurance
as a Function of Equivalized Household Rank

Panel A: Receipt Rate

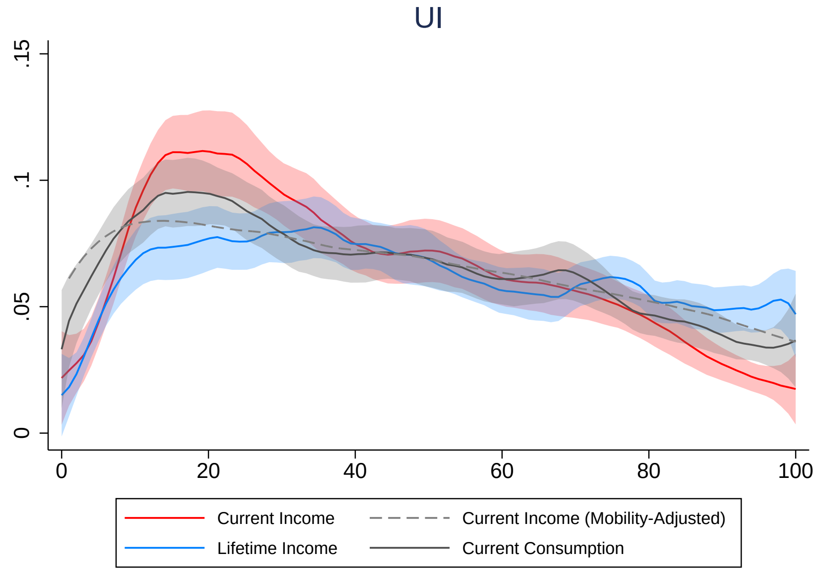

Panel B: Average Annual Per-Capita Transfer

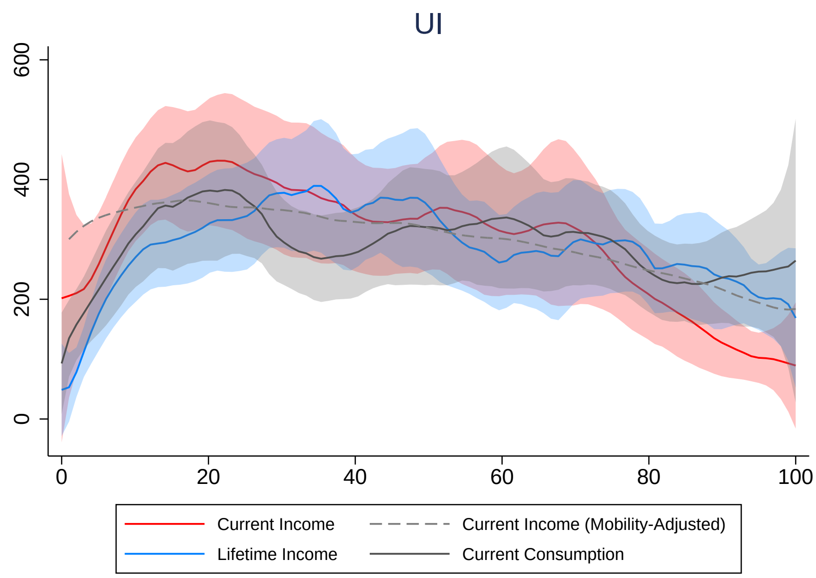

Notes: This figure displays average total annual per-capita values of benefits and receipt rates for unemployment insurance as functions of household ranks in the distributions of equivalized current income, lifetime income, and current consumption. The functions are estimated by local linear regressions with bandwidths of three percentiles. Shaded regions reflect bootstrapped 95-percent simultaneous confidence bands, as in Chernozhukov et al. (2013), with clustering by household.

Figure A6: Selection into Transfer Receipt: Very Likely Eligible Subsample

Panel A: Selection on Consumption

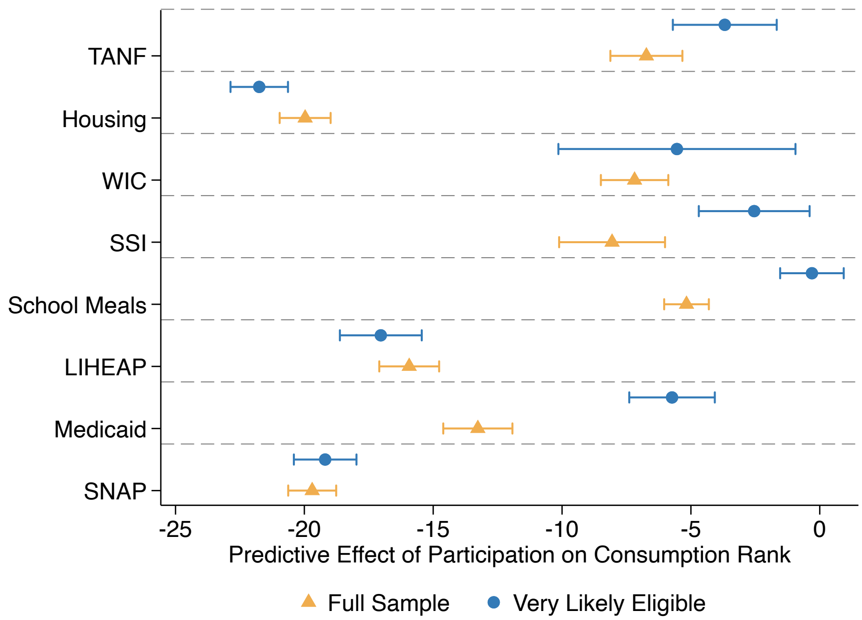

Panel B: Selection on Lifetime Income

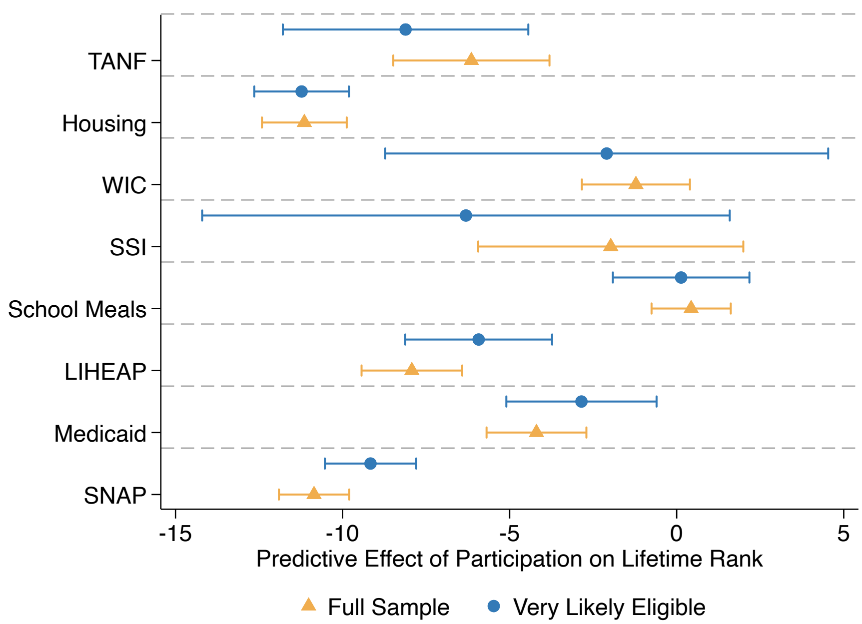

Notes: This figure displays estimates of the predictive effect of transfer benefit receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient γ from Equation 4). Both yellow diamonds and blue circles restrict the sample to simulated eligibles. For estimates represented by blue circles, we further limit the sample to people who, in a logistic regression of simulated eligibility status on demographic observables, have a predicted probability of eligibility above 0.8. The eligibility logit uses the following demographic variables: age (in ten bins), sex, marital status, race/ethnicity (white, black, Hispanic, other), education (less than high school, high school, some college, BA, more than BA), household size, homeownership, disability, and rank-transformed current income, lifetime income, and consumption. The confidence intervals are for the 95-percent level and clustering by household. In Panel B, we adapt Equation 4 by replacing the transfer indicator with indicators for the number of unique transfers received.

Figure A7: Transfer Receipt as a Function of Rank

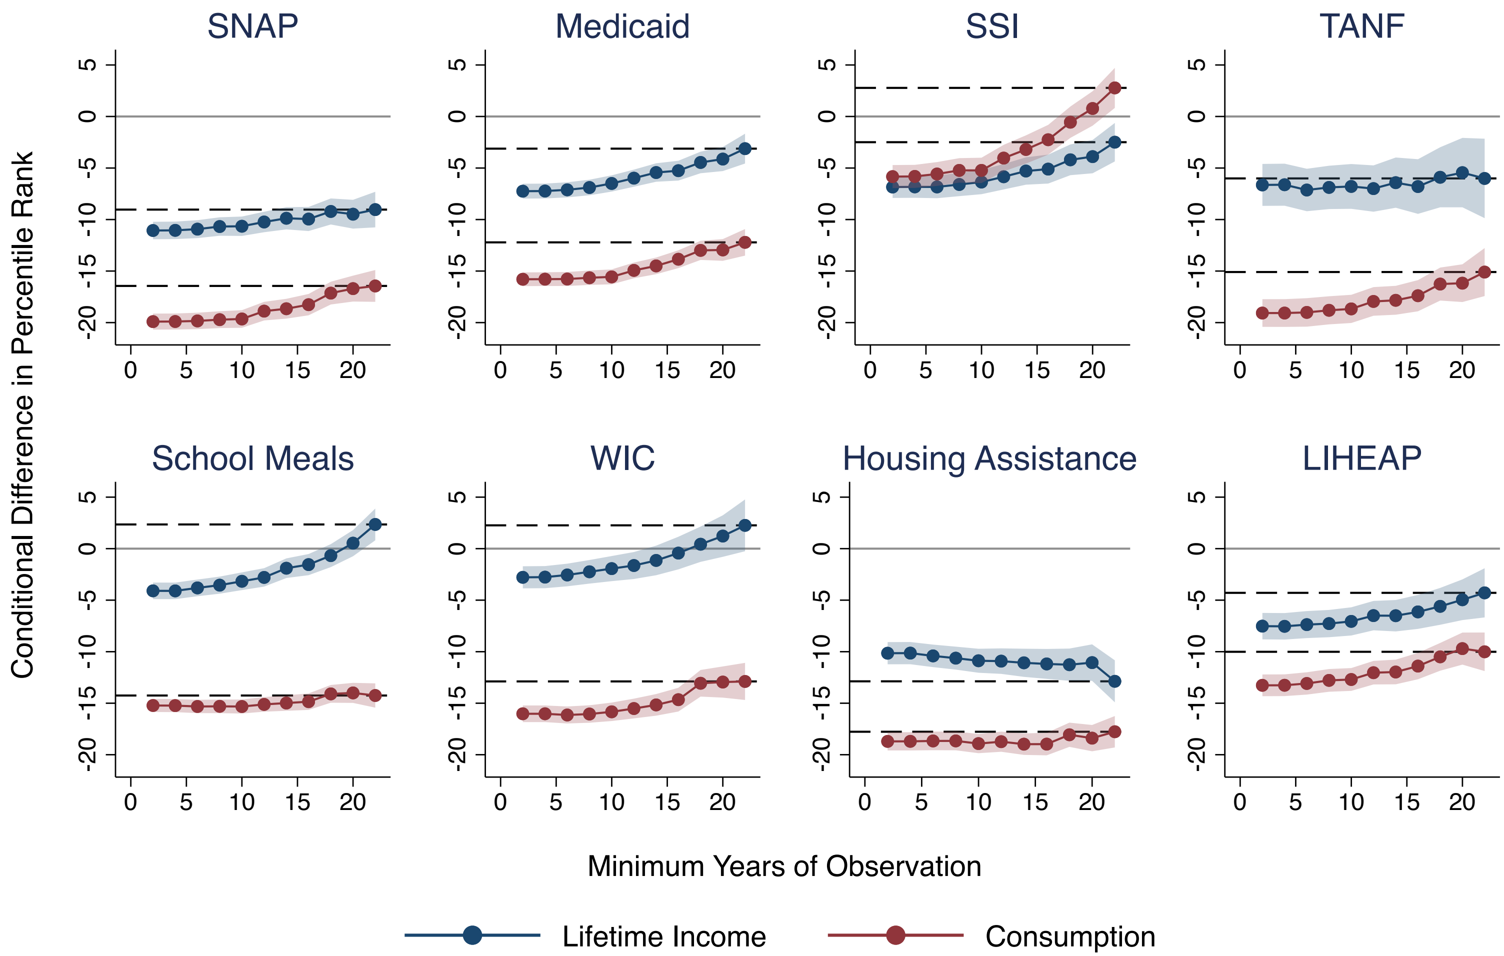

Notes: This figure displays program receipt rates as functions of income percentile ranks, variously for households or individuals, and for current or lifetime income. Section 2 introduces data sources and the measurement of lifetime income. The functions are estimated using a local linear regression with a bandwidth of three percentiles. Shaded regions reflect bootstrapped 95-percent simultaneous confidence bands, as in Chernozhukov et al. (2013), with clustering by household.

Figure A8: Advantageous Selection on Transfer Receipt in the Distant Future

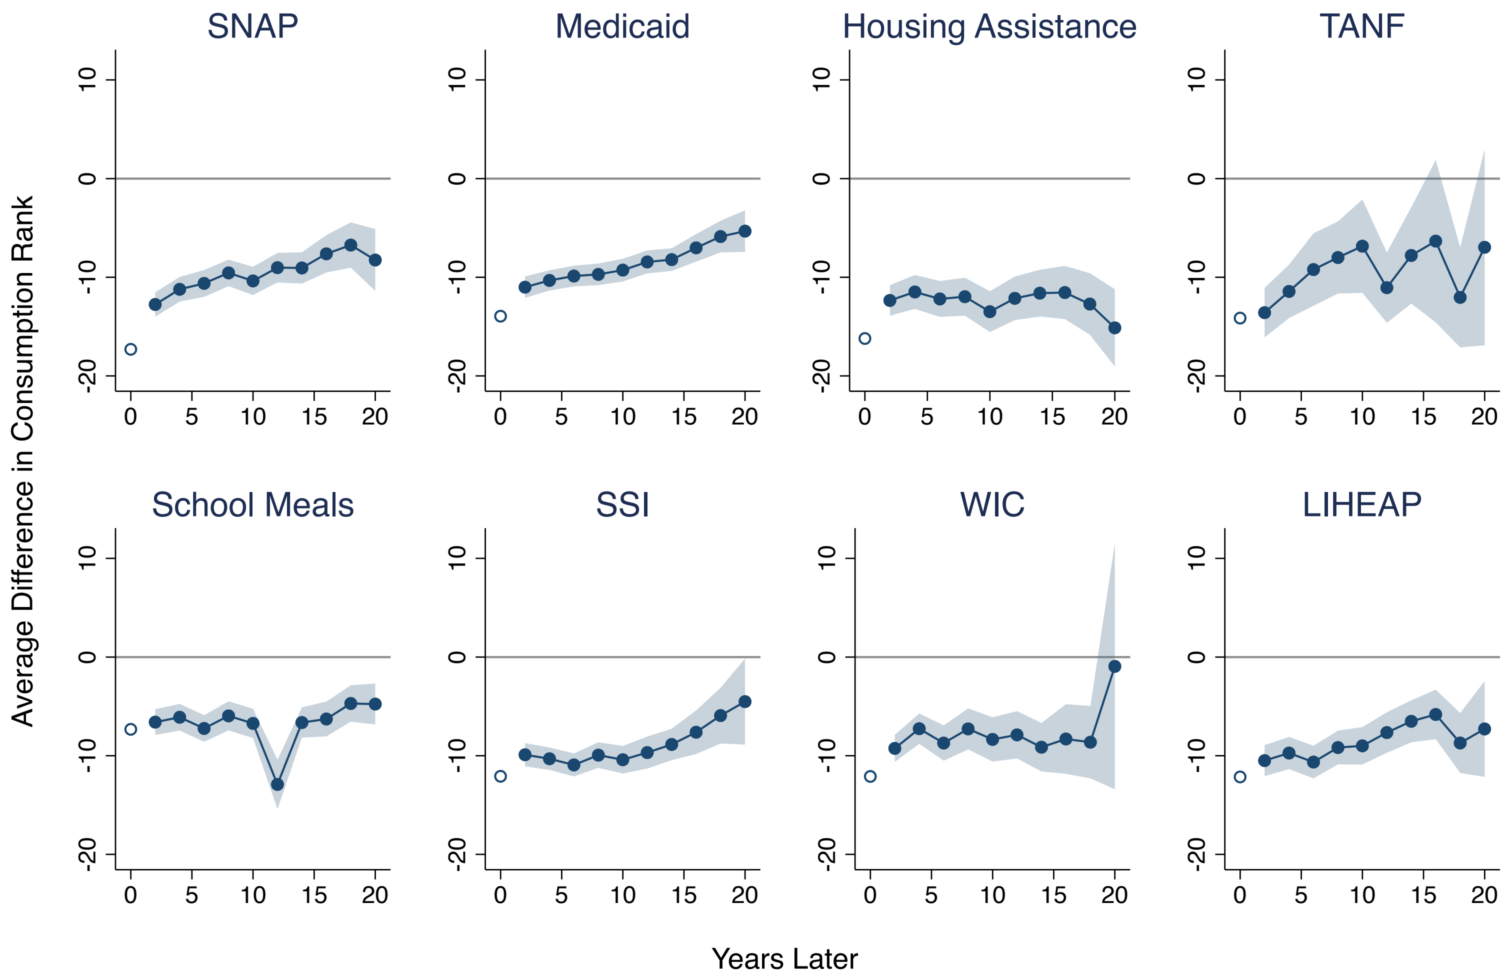

Notes: This figure displays the predictive effect of transfer receipt k years ahead on consumption rank this year conditional on current income rank. The regression equation is $\overline{R}_{it} = \alpha_{ct} + \beta D_{i,t+k} + f(R_{it}) + u_{it}$, where we plot β for each horizon k. The estimation sample is always restricted to current non-recipients, $D_{it} = 0$. Shaded regions reflect bootstrapped 95-percent pointwise confidence intervals, with clustering by household.

Figure A9: Selection into Transfers, with Predicted-Income Control

Panel A: Consumption

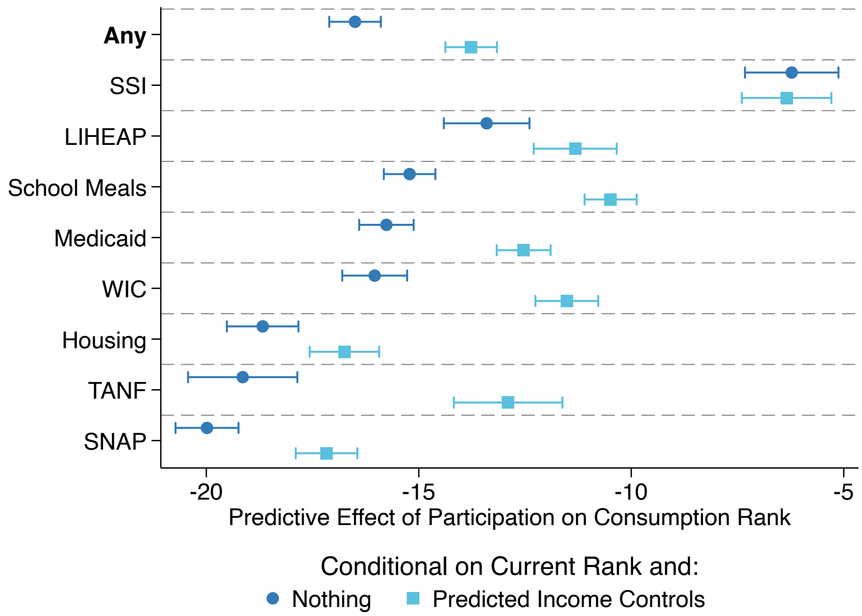

Panel B: Lifetime Income

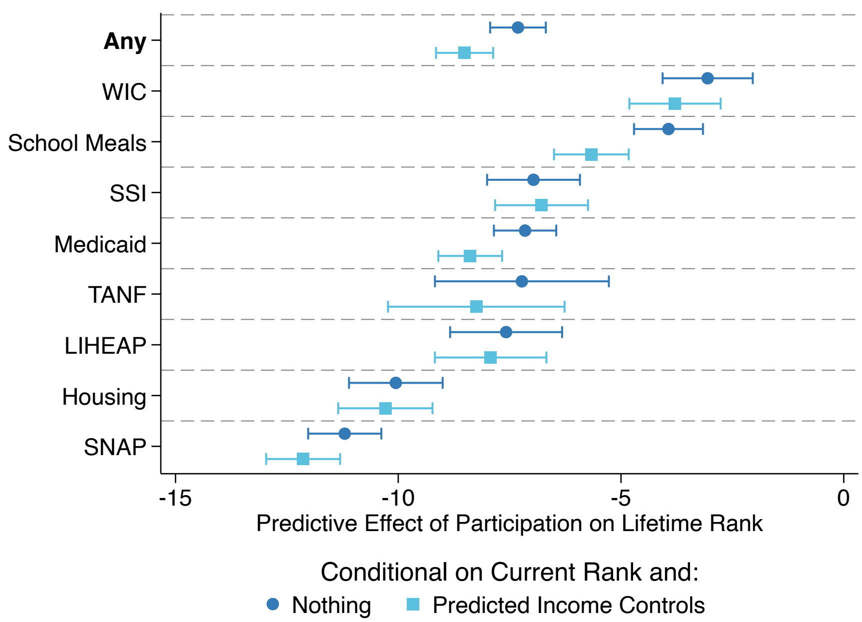

Notes: This figure displays estimates of the predictive effect of transfer receipt on equivalized household consumption rank, conditional on current-income rank as well as predicted-income rank. Income prediction uses a Poisson regression as explained in Section 3. For estimates represented by blue circles, we do not add additional control variables to the specification, whereas for the teal squares, we estimate the regression only on people whom we simulate to be eligible. The "any" row of Panel A is an indicator for receipt of at least one of the eight transfers. The confidence intervals are for the 95-percent level and reflect clustered standard errors by household.

Figure A10: Measurement Error Simulations

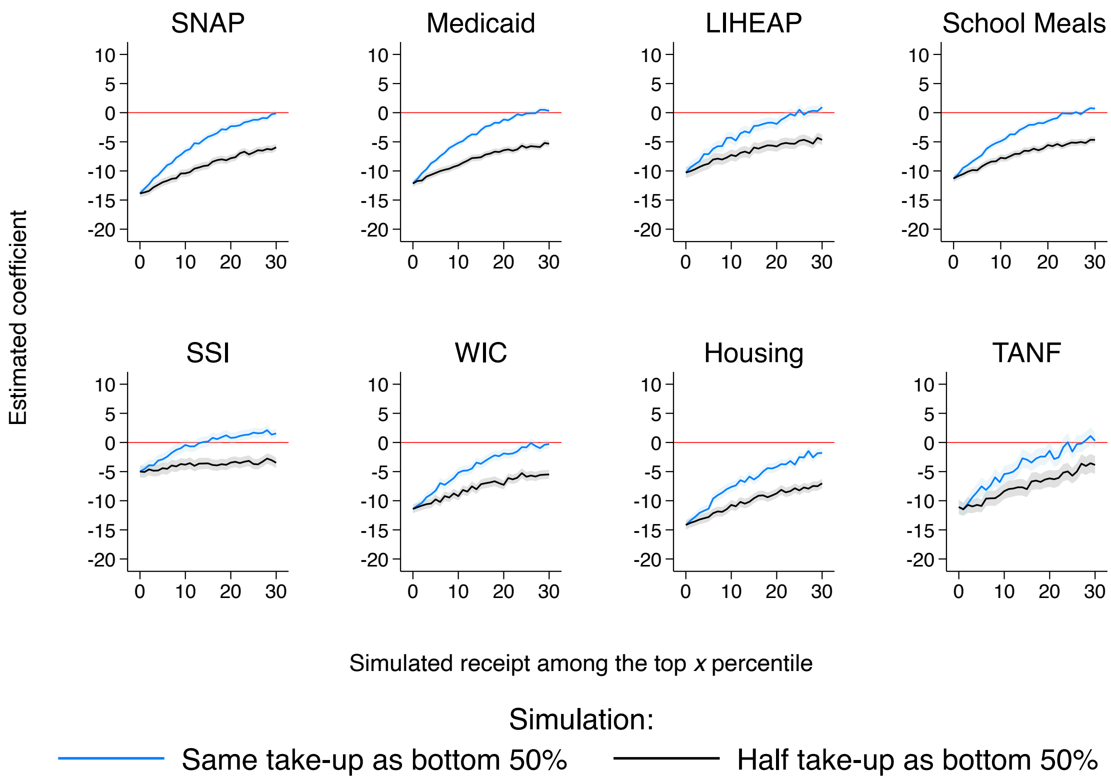

Notes: This figure displays the predictive effect of transfer receipt on consumption rank given income rank (Equation 4) when we assume that take-up is underreported for the top x consumption percentiles. In blue, we assume that the top x percentiles actually have the same take-up rate as the bottom half of the consumption distribution. In black we assume that top take-up rate is half that of the bottom half of the consumption distribution. Shaded regions reflect 95-percent pointwise confidence intervals, with standard errors clustered by household.

Figure A11: Selection into Transfer Receipt Over Time

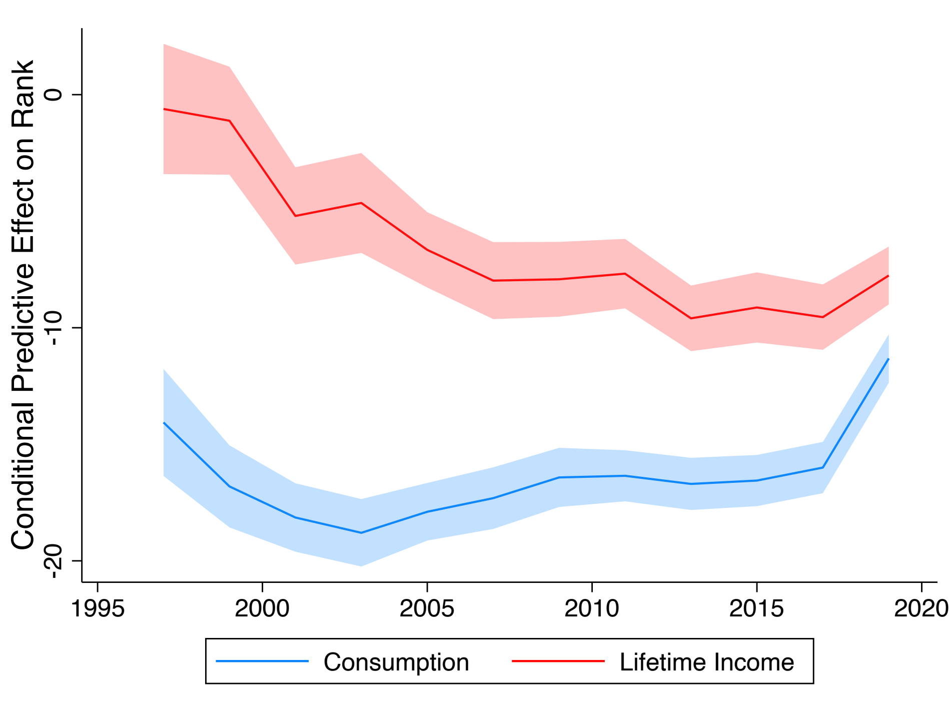

Notes: This figure displays coefficients from the following regression specification:

$$\overline{R}_{its} = \alpha_{cts} + \beta_t D_{its} + f_s(R_{it}) + u_{its},$$

where i denotes households, t denotes years, and s denotes transfer programs. The outcome $\overline{R}_{its}$ is equivalized household consumption rank in the blue line and equivalized household lifetime income rank in red. The data are stacked across programs, so that each individual–year appears eight times, once for each transfer program s. We thus include cohort-year effects α_{cts} specific to each transfer, as well as transfer-specific spline controls $f_s(\cdot)$ for current-income rank R_{it}. The coefficients β_t thus report an average selection effect across transfer programs in a given year t. Shaded regions reflect 95-percent pointwise confidence intervals, with standard errors clustered by household. All regressions use PSID sample weights but are not otherwise adjust to account for variation in transfer program size.

Figure A12: Decomposition of Difference Between Consumption and Current Incidence

Panel A: Average Total Annual Per-Capita Benefits

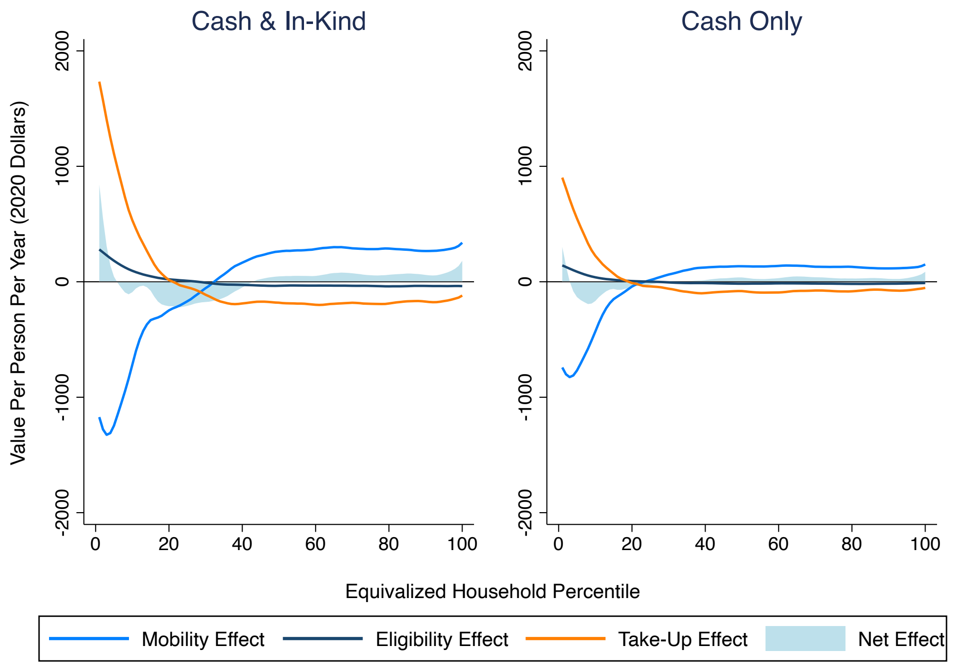

Panel B: Receipt Rate by Program

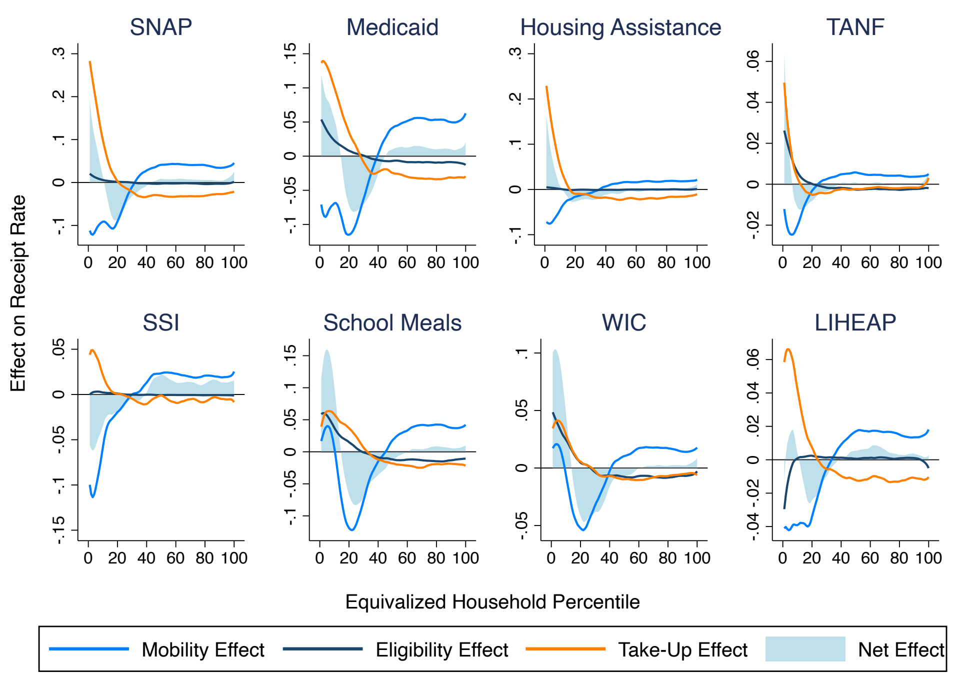

Notes: This figure displays decompositions of differences between current-income and consumption incidence into mobility, eligibility, and take-up effects. The underlying functions are estimated using local linear regressions with bandwidths of three percentiles.

Figure A13: Decomposition of Difference Between Lifetime and Current Incidence

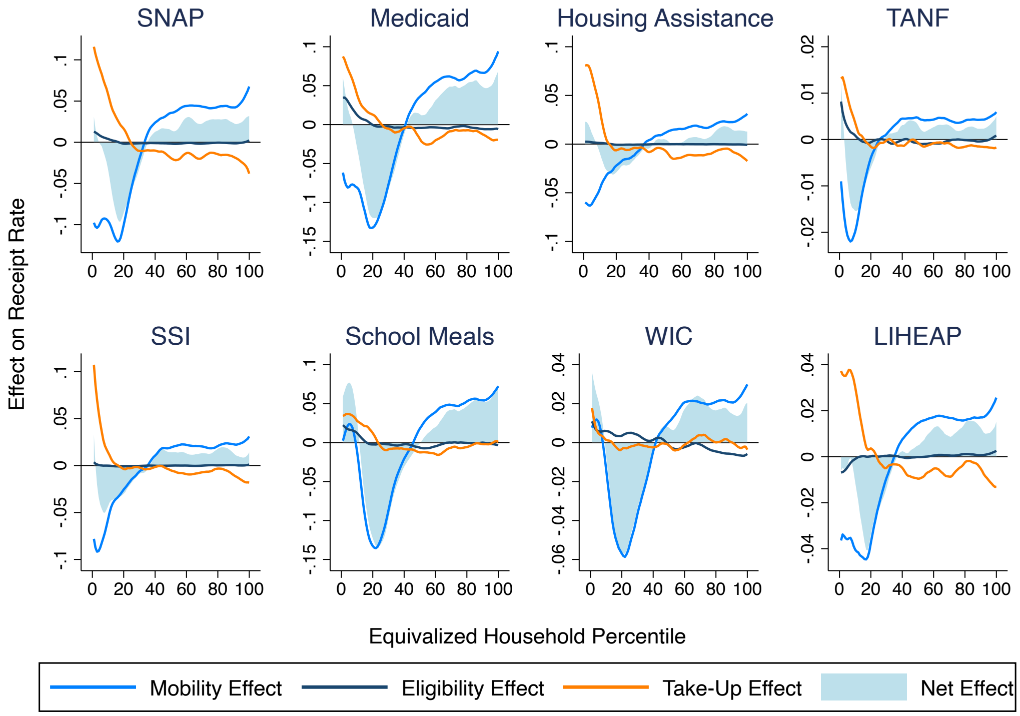

Notes: This figure displays a decomposition of the difference between current and lifetime incidence into mobility and take-up effects. The original functions are estimated using a local linear regression with a bandwidth of three percentiles.

Table A1: SNAP Receipt, Eligibility, and Take-Up Rates by Income and Lifetime Income Quintile

Panel A: Receipt Rate

		Income Quintile					
		1	2	3	4	5	Avg.
Lifetime Income Quintile	1	43.2	19.9	6.4	1.7	0.1	35.3
	2	26.1	11.9	3.0	1.0	0.1	8.4
	3	20.5	9.5	2.0	0.4	0.2	3.3
	4	19.7	6.9	2.1	0.3	0.2	1.3
	5	20.2	4.6	1.9	0.4	0.2	0.9
	Avg.	33.6	12.2	2.8	0.6	0.2	

Panel B: Simulated Eligibility Rate

		Income Quintile					
		1	2	3	4	5	Avg.
Lifetime Income Quintile	1	80.5	24.0	0.0	0.0	0.0	49.7
	2	71.0	16.0	0.2	0.0	0.0	17.5
	3	70.8	13.7	0.0	0.0	0.0	10.0
	4	68.3	15.5	0.1	0.1	0.0	7.9
	5	76.3	19.0	2.3	1.0	0.1	10.0
	Avg.	76.3	18.1	0.4	0.3	0.1	

Panel C: Take-Up Rate Among Simulated Eligibles

		Income Quintile					
		1	2	3	4	5	Avg.
Lifetime Income Quintile	1	46.2	36.3	.	.	.	44.9
	2	30.1	27.5	.	.	.	29.4
	3	25.0	25.1	.	.	.	24.9
	4	23.9	17.8	.	.	.	22.3
	5	23.0	10.3	.	.	.	18.9
	Avg.	37.5	27.2	.	.	.	

Notes: This table reports the shares of households that receive SNAP (Panel A), are simulated to be eligible for SNAP (Panel B), and take up SNAP conditional on being simulated eligible (Panel C). Households are split by quintiles of equivalized household current and lifetime income. Due to low rates of simulated eligibility, we do not report take-up rates for the top three income quintiles.

Table A2: Transfer Amounts and Receipt Rates at the Bottom of the Distributions of Income and Consumption

	Income				Consumption				Counterfactual for Consumption			
Percentiles	0–5	0–10	0–25	0–50	0–5	0–10	0–25	0–50	0–5	0–10	0–25	0–50
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: Average Dollars Per Person Per Year												
Total	3,634	3,216	2,184	1,380	3,973	3,294	2,136	1,316	2,271	2,015	1,530	1,099
	(97)	(67)	(33)	(18)	(88)	(58)	(31)	(18)	(25)	(18)	(11)	(8)
Cash	1,928	1,676	1,014	577	1,917	1,545	927	545	1,086	931	663	454
	(72)	(49)	(23)	(12)	(62)	(41)	(21)	(12)	(16)	(11)	(7)	(4)
Panel B: Receipt Rates												
SNAP	43.3	38.2	30.9	19.1	56.2	48.2	30.7	18.3	30.3	27.5	21.1	15.0
	(1.1)	(0.7)	(0.4)	(0.2)	(1.0)	(0.7)	(0.4)	(0.2)	(0.3)	(0.2)	(0.1)	(0.1)
Medicaid	44.8	42.0	37.1	26.5	54.8	50.2	38.4	25.3	35.7	33.6	28.0	21.2
	(1.1)	(0.8)	(0.4)	(0.3)	(1.0)	(0.7)	(0.4)	(0.3)	(0.2)	(0.2)	(0.1)	(0.1)
Housing Assistance	23.0	20.3	14.2	9.4	34.4	27.1	15.9	9.4	15.0	13.5	10.4	7.6
	(1.0)	(0.6)	(0.3)	(0.2)	(1.0)	(0.6)	(0.3)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)
TANF	5.9	5.6	3.7	2.1	9.5	7.0	3.7	2.0	3.8	3.3	2.4	1.6
	(0.4)	(0.3)	(0.2)	(0.1)	(0.5)	(0.3)	(0.2)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)
SSI	29.6	25.5	17.2	10.7	23.1	20.2	13.9	9.2	17.9	15.7	11.8	8.5
	(1.0)	(0.7)	(0.4)	(0.2)	(0.9)	(0.6)	(0.3)	(0.2)	(0.2)	(0.2)	(0.1)	(0.1)
School Meals	17.7	17.0	23.0	18.9	32.1	32.3	27.3	18.6	20.4	20.7	19.1	15.1
	(0.7)	(0.5)	(0.4)	(0.2)	(0.9)	(0.6)	(0.4)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)
WIC	6.5	6.8	9.8	7.7	17.2	16.3	12.0	7.6	8.5	8.7	7.8	6.1
	(0.4)	(0.3)	(0.2)	(0.1)	(0.7)	(0.5)	(0.3)	(0.2)	(0.1)	(0.1)	(0.0)	(0.0)
LIHEAP	15.2	14.7	11.9	7.4	16.5	16.0	11.3	7.0	11.4	10.5	8.2	5.8
	(0.8)	(0.6)	(0.3)	(0.2)	(0.8)	(0.5)	(0.3)	(0.2)	(0.1)	(0.1)	(0.1)	(0.0)

Notes: This table reports sample means of average total transfer payments per person per year (in 2020 constant dollars) as well as receipt rates by transfer program. Each column is for a different range of percentiles in a distribution. Columns 1–4 are with respect to household equivalized current income, and Columns 5–8 are with respect to household equivalized consumption. Columns 9–12 calculate the consumption incidence under the counterfactual in which receipt is a function of income rank. Parentheses report standard errors clustered by household.

Table A3: Transfer Amounts and Receipt Rates at the Bottom of the Distributions of Current and Lifetime Income

	Current Income				Lifetime Income				Counterfactual for Lifetime Income			
Percentiles	0–5 (1)	0–10 (2)	0–25 (3)	0–50 (4)	0–5 (5)	0–10 (6)	0–25 (7)	0–50 (8)	0–5 (9)	0–10 (10)	0–25 (11)	0–50 (12)
Panel A: Average Dollars Per Person Per Year												
Total	3,634	3,216	2,184	1,380	3,815	3,099	1,927	1,220	2,581	2,196	1,551	1,080
	(97)	(67)	(33)	(18)	(102)	(63)	(32)	(18)	(25)	(18)	(12)	(8)
Cash	1,928	1,676	1,014	577	2,011	1,544	881	521	1,293	1,056	693	456
	(72)	(49)	(23)	(12)	(77)	(46)	(22)	(12)	(16)	(11)	(7)	(4)
Panel B: Receipt Rates												
SNAP	43.3	38.2	30.9	19.1	43.2	38.2	26.2	16.6	32.5	28.7	21.0	14.6
	(1.1)	(0.7)	(0.4)	(0.2)	(1.0)	(0.7)	(0.4)	(0.2)	(0.3)	(0.2)	(0.1)	(0.1)
Medicaid	44.8	42.0	37.1	26.5	48.5	43.4	31.8	22.1	37.1	34.1	27.1	20.2
	(1.1)	(0.8)	(0.4)	(0.3)	(1.0)	(0.7)	(0.4)	(0.3)	(0.2)	(0.2)	(0.1)	(0.1)
Housing Assistance	23.0	20.3	14.2	9.4	25.1	21.5	13.2	8.6	16.6	14.3	10.4	7.4
	(1.0)	(0.6)	(0.3)	(0.2)	(1.0)	(0.6)	(0.3)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)
TANF	5.9	5.6	3.7	2.1	6.1	4.8	2.9	1.8	4.4	3.6	2.5	1.6
	(0.4)	(0.3)	(0.2)	(0.1)	(0.4)	(0.3)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)
SSI	29.6	25.5	17.2	10.7	27.3	22.1	13.9	9.2	20.5	17.3	12.1	8.4
	(1.0)	(0.7)	(0.4)	(0.2)	(1.0)	(0.6)	(0.3)	(0.2)	(0.2)	(0.2)	(0.1)	(0.1)
School Meals	17.7	17.0	23.0	18.9	24.7	24.3	20.5	15.0	18.9	19.2	17.5	14.0
	(0.7)	(0.5)	(0.4)	(0.2)	(0.8)	(0.6)	(0.3)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)
WIC	6.5	6.8	9.8	7.7	9.6	9.0	7.7	6.0	7.7	7.8	7.1	5.6
	(0.4)	(0.3)	(0.2)	(0.1)	(0.5)	(0.3)	(0.2)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)
LIHEAP	15.2	14.7	11.9	7.4	14.8	14.1	9.9	6.3	12.3	11.0	8.1	5.7
	(0.8)	(0.6)	(0.3)	(0.2)	(0.8)	(0.5)	(0.3)	(0.2)	(0.1)	(0.1)	(0.1)	(0.0)

Notes: This table reports sample means of average total transfer payments per person per year (in 2020 constant dollars) as well as receipt rates by transfer program. Each column is for a different range of percentiles in a distribution. Columns 1–4 are with respect to household equivalized current income, and Columns 5–8 are with respect to household equivalized lifetime income. Columns 9–12 calculate the lifetime incidence under the counterfactual in which receipt is a function of current-income rank. Parentheses report standard errors clustered by household.

Table A4: Dollars and Percentage Differences Between Recipients and Similar-Income Non-Recipients

	Proportion Difference		Difference in 2020 Constant Dollars	
	Consumption (1)	Lifetime Income (2)	Consumption (3)	Lifetime Income (4)
SNAP	-0.461***	-0.440***	-10,843***	-23,648***
	(0.017)	(0.077)	(413)	(4,262)
Medicaid	-0.405***	-0.495***	-9,516***	-26,614***
	(0.013)	(0.053)	(300)	(3,058)
Housing Assistance	-0.367***	-0.205**	-8,626***	-11,029**
	(0.022)	(0.099)	(516)	(5,359)
TANF	-0.589***	-0.839***	-13,848***	-47,394***
	(0.079)	(0.168)	(1,848)	(9,855)
SSI	-0.084***	-0.353***	-1,976***	-18,979***
	(0.018)	(0.081)	(413)	(4,409)
School Meals	-0.487***	-0.420***	-12,080***	-22,385***
	(0.012)	(0.038)	(310)	(2,113)
WIC	-0.506***	-0.481***	-12,572***	-25,721***
	(0.023)	(0.053)	(585)	(2,939)
LIHEAP	-0.321***	-0.405***	-7,541***	-21,810***
	(0.021)	(0.075)	(501)	(4,152)

Notes: This table reports estimates of differences in consumption and lifetime income between transfer receipients and nonrecipients, conditional on current income. All columns report estimates obtained via Poisson regression. Columns 1 and 2 report exponentiated coefficients ($\exp(\beta) - 1$) from these regressions. Columns 3 and 4 report the dollar effects. Each cell is its own regression. All specifications control flexibly for the logarithm of equivalized current household income using cubic basis splines. Standard errors are clustered by household. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A5: Selection into Transfers, Adjusted for Misreporting of Transfer Receipt

	Baseline		Adjusted for Misreporting	
	Consumption (1)	Lifetime Income (2)	Consumption (3)	Lifetime Income (4)
Panel A: SNAP				
Receives Transfer	-17.6***	-11.1***	-26.4***	-14.3***
	(0.6)	(0.6)	(0.8)	(0.9)
Panel B: Medicaid				
Receives Transfer	-14.4***	-7.0***	-23.4***	-12.2***
	(0.5)	(0.5)	(0.7)	(0.8)

Notes: This table examines the effect of corrections for misreporting of transfer receipt on estimates of selection into transfers by consumption rank and lifetime-income rank, conditional on current-income rank. The estimating equation is Equation 4. In Columns 3 and 4, we replace reported receipt with the adjusted measures from Mittag (2019) for SNAP and Davern et al. (2019) for Medicaid. All specifications control flexibly for current-income rank using cubic basis splines. Standard errors are clustered by household. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A6: Well-Measured Consumption and Transfer Receipt

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)	School Meals (6)	WIC (7)	LIHEAP (8)
Panel A: Rent and Owner's Equivalent Rent (1997–2019)								
Receives Transfer	-0.431***	-0.333***	-0.670***	-0.615***	-0.137***	-0.374***	-0.370***	-0.320***
	(0.014)	(0.011)	(0.021)	(0.046)	(0.018)	(0.011)	(0.016)	(0.021)
Panel B: Vehicle Lease Cost and Equivalent Lease Cost (1999–2019)								
Receives Transfer	-0.216***	-0.242***	-0.113***	-0.372***	-0.026*	-0.296***	-0.291***	-0.167***
	(0.010)	(0.007)	(0.012)	(0.023)	(0.014)	(0.008)	(0.010)	(0.014)
Panel C: Food at Home Expenditure (1999–2019)								
Receives Transfer	-0.531***	-0.245***	-0.243***	-0.507***	-0.153***	-0.160***	-0.261***	-0.328***
	(0.014)	(0.011)	(0.017)	(0.040)	(0.018)	(0.011)	(0.016)	(0.021)
Panel D: Utility Expenditure (1999–2019)								
Receives Transfer	-0.045***	-0.095***	-0.266***	-0.215***	0.016	-0.113***	-0.124***	0.027
	(0.014)	(0.010)	(0.021)	(0.036)	(0.019)	(0.011)	(0.014)	(0.018)
Panel E: Gasoline Expenditure (1999–2019)								
Receives Transfer	-0.155***	-0.127***	-0.129***	-0.278***	-0.115***	-0.130***	-0.143***	-0.179***
	(0.016)	(0.013)	(0.022)	(0.045)	(0.023)	(0.013)	(0.017)	(0.024)

Notes: This table reports estimates of the predictive effect of transfer receipt on (log) levels of reported consumption, conditional on current-income rank. Each panel row is for a different consumption outcome, and each column is for a different transfer. The year ranges in parentheses indicate data coverage for the outcome of interest. All specifications control flexibly for current-income rank using cubic basis splines. Standard errors are clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A7: Durable-Goods Ownership and Transfer Receipt

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)	School Meals (6)	WIC (7)	LIHEAP (8)
Panel A: HH Owns Primary Residence (1997–2019)								
Receives Transfer	-0.161***	-0.070***	-0.422***	-0.138***	-0.074***	0.001	0.032***	-0.064***
	(0.008)	(0.007)	(0.006)	(0.015)	(0.011)	(0.007)	(0.009)	(0.012)
Panel B: Number of Rooms in Home (1997–2019)								
Receives Transfer	-0.522***	-0.554***	-0.668***	-0.603***	-0.080*	-0.763***	-0.667***	0.000
	(0.032)	(0.029)	(0.034)	(0.072)	(0.048)	(0.031)	(0.042)	(0.046)
Panel C: Central Air Conditioning at Home (1997–2009)								
Receives Transfer	-0.044***	-0.020**	-0.002	-0.144***	0.025	-0.014	-0.033**	-0.068***
	(0.012)	(0.010)	(0.015)	(0.025)	(0.020)	(0.010)	(0.014)	(0.018)
Panel D: HH Owns a Car (1999–2019)								
Receives Transfer	-0.121***	-0.038***	-0.224***	-0.113***	-0.091***	0.035***	0.050***	-0.038***
	(0.008)	(0.006)	(0.011)	(0.022)	(0.011)	(0.006)	(0.009)	(0.012)
Panel D: HH Owns a Computer (2003–2019)								
Receives Transfer	-0.115***	-0.042***	-0.130***	-0.082***	-0.033***	-0.011	-0.056***	-0.083***
	(0.009)	(0.007)	(0.013)	(0.027)	(0.012)	(0.008)	(0.011)	(0.014)
Panel E: HH Owns a Smartphone (2015–2019)								
Receives Transfer	-0.004	0.006	-0.052***	0.056	0.007	0.043***	0.056***	-0.040*
	(0.012)	(0.009)	(0.017)	(0.035)	(0.012)	(0.008)	(0.011)	(0.021)

Notes: This table reports estimates of the predictive effect of transfer receipt on measures of household durable-goods ownership, conditional on current-income rank. Each panel row is for a different consumption outcome, and each column is for a different transfer. The year ranges in parentheses indicate data coverage for the outcome of interest. All specifications control flexibly for current-income rank using cubic basis splines. Standard errors are clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A8: Simulated Eligibility and Receipt Rates

Program	(1) P(Receive | Simulated Eligible)	(2) P(Receive | Not Sim. Elig.)	(3) P(Sim. Elig. | Receive)	(4) P(Sim. Elig. | Not Receive)
Housing Assistance	0.10	0.02	0.82	0.39
LIHEAP	0.11	0.01	0.97	0.72
Medicaid	0.54	0.09	0.49	0.08
SNAP	0.35	0.04	0.69	0.15
SSI	0.28	0.05	0.12	0.02
TANF	0.07	0.00	0.78	0.11
UI	0.33	0.05	0.25	0.03
WIC	0.44	0.02	0.50	0.03

Notes: This table presents the share of households who do or do not receive transfers, conditional on our simulated eligibility measures. See Appendix B for details on forming the measures of simulated eligibility.

B Data Appendix

This appendix first explains measurement details for consumption, lifetime income, and simulated eligibility. Next, it presents the extensions we mention in Section 3. These are the decomposition of differences in incidence into mobility, eligibility, and take-up effects, as well as the inverse approach of finding the income-tax reform that would replicate the consumption or lifetime incidence of transfers.

B.1 Consumption Ranks

We compute households' equivalized consumption ranks using the expenditure data available in the PSID in a given year. Not all consumption categories are available in each year. In particular, we observe expenditures on clothing, furniture, travel, and recreation starting in 2005 and computer expenditures starting in 2017. We observe housing rents (actual and imputed) starting in 1997, and all other expenditures starting in 1999. These expenditure categories are childcare, education, food, health, transportation, and utilities (energy and water starting in 1999, phone/cable/internet starting in 2005).

As noted in Section 2, we follow Meyer and Sullivan (2023) in making two adjustments so that we more closely measure consumption rather than expenditure. Broadly, these adjustments estimate consumption flows from households' two key durable goods, homes and vehicles.

For renters, we take their paid rents as their housing consumption. For owner-occupiers, we obtain imputed rents in several steps. In 2017 and 2019, owner-occupier households were asked "If someone were to rent this (apartment/mobile home/home) today, how much do you think it would rent for per month, unfurnished and without utilities?" We take these values as housing consumption for such households. For all years in our sample period, households who report that their housing is free are asked "How much would it rent for if it were rented?", which we use as their housing consumption. Finally, we construct a mapping from home values to owners' equivalent rents using the cross-sectional relationship in 2017 and 2019 between households' estimates of their home's value and its equivalent rent.[27]

Transportation consumption is constructed as follows. We count any expenditures on gasoline, parking, public transportation, taxis, other transportation toward the household's transportation consumption. Due to PSID data limitations, we also count as consumption any expenditures on vehicles other than the household's three reported primary vehicles. For households that lease any of their three primary vehicles, we count their lease costs toward transportation consumption. For

[27]For households that do not report an exact home value, we use the midpoint of the elicited range. For households who say their homes are worth more than $400,000, but do not report an exact value, we impute it as the sample mean conditional on exceeding $400,000 among households who report exact home values.

households that own any of their three primary vehicles, we impute the equivalent lease cost from a hedonic regression.

To estimate this hedonic regression, we restructure our data into a vehicle-level dataset. Households that lease or own a vehicle report the vehicle's manufacturer (e.g., Toyota), its make (e.g., Lexus), its age at acquisition (year of purchase or lease minus model year), and its "type" (car, pickup/truck, van, utility, or motor home). We estimate Poisson regression models of all two-way interactions of these variables, along with indicator variables for calendar year and the rank (1/2/3) of the vehicle in the household's list, The outcomes are purchase price or lease cost, winsorizing values at the first and 99th percentiles. We then collapse these predicted values for purchase price and lease cost to the level of manufacturer, make, age, and type. This procedure yields an estimated lease cost equivalent for owned vehicles.[28]

B.2 Lifetime Income Ranks

Step 1: Estimate lifecycle regression parameters. Letting i index individuals, t index calendar years, and a index age in years, we estimate Poisson regression models of the following form:

$$\mathrm{E}[y_{it} \mid X_{it}] = \exp(\alpha_i \lambda_a + X'_{it}\beta_a), \tag{10}$$

where α_i is an individual fixed effect, α_t is a calendar-year fixed effect, X_{it} is a matrix of time-varying demographic characteristics, and λ_a and β_a are vectors of age-specific coefficients. The outcome y_{it} is individual income. For individuals with zero income in all observed years, we impute a constant annual income of \$100.[29] The age-specific coefficients are initialized to $\lambda_a = 1$ for all a but will be estimated in an outer loop discussed below. We make several adjustments before using the regression results to estimate lifetime-income ranks.

Step 2: Shrink fixed effects. First, we apply the empirical Bayes methods in Morris (1983) to shrink the estimated individual fixed effects $\widehat{\alpha_i}$ toward a conditional expectation fit from several time-invariant individual characteristics.[30] These methods accommodate both unequal individual means and unequal sampling variances in the fixed effects by iteratively re-estimating the extent of true heterogeneity among individuals and the conditional expectation function using weighted least

[28] For missing values, we impute using the cross-sectional relationship between fitted purchase prices and fitted lease values from these Poisson regressions.

[29] In an unadjusted Poisson regression, estimates of the individual fixed effects α_i diverge to negative infinity for any individual i who earns $y_{it} = 0$ for all observed periods t. By setting their y_{it} to a very low positive value, we obtain convergent fixed effects and rank these individuals at the bottom of the lifetime-income distribution. Importantly, this procedure does affect our estimates of β, as the fixed effects α_i perfectly explain the income of these individuals.

[30] Recent applications of these methods in economics include Chandra et al. (2016) and Sorkin (2018). We refer interested readers to their appendices for detailed expositions. One key modification we make to their approach is to use a within-individual Bayesian bootstrap (Rubin, 1981) instead of actual resampling.

squares. Our baseline specification uses sex, race, and ethnicity to fit this conditional expectation.

Step 3: Outer loop. Haider and Solon (2006) emphasize that the "error-in-variables" model of lifetime income is misspecified, as the predictive effect of individual fixed effects grows over the lifecycle. To account for this, we estimate the λ_a terms in Equation 1 through the following outer loop. Consider the first loop, in which we have initialized $\lambda_a = 1$ and have shrunken estimates of α_i. We can estimate the following Poisson model:

$$\mathrm{E}[y_{it} \mid X_{it}] = \exp(\widehat{\alpha}_i \lambda_a + X_{it}' \beta_a), \tag{11}$$

importantly treating $\{\widehat{\alpha}_i\}$ as data rather than as parameters. We then obtain coefficient estimates $\{\widehat{\lambda}_a\}$, and with these in hand, we return to step 1 and iterate until convergence of $\{\widehat{\alpha}_i, \widehat{\lambda}_a, \widehat{\beta}_a\}$. In practice, we find that convergence is fast; three runs of the outer loop are sufficient.

Step 4: Balance the panel. Having estimated the model in Equation 1, we use it to predict income from ages 18 to 65, irrespective of the years in which we observe an individual's actual income. An individual's predicted income in year t is $\widehat{y}_{it} = \exp(\widehat{\alpha}_i^* \widehat{\lambda}_a + X_{it}' \widehat{\beta}_a)$, where $\widehat{\alpha}_i^*$ are the shrunken estimates of the individual fixed effects. Lifetime income are then

$$\overline{y}_i = \sum_t \widehat{y}_{it}, \tag{12}$$

where the summation over t is for the years $\{\underline{T}_i, \ldots, \overline{T}_i\}$ in which individual i is between the ages of 18 and 65. Importantly, however, we do not observe individual characteristics X_{it} in all years and therefore must impute them. In our baseline specification, we assume these characteristics are unchanged from the nearest period of observation, except for age.

Step 5: Construct ranks. We define an individual's lifetime income percentile rank as $\Pr(y \leq \overline{y}_i \mid c_i = c)$, where $\overline{y}_i$ is their estimated lifetime income and c_i is their birth-year cohort. We define an individual's current income percentile rank as $\Pr(y \leq y_{it} \mid c_i = c)$, again ranking individuals each year within their birth cohorts. Appendix A presents figures of our main results when we do not rank current income within cohorts.

We construct current and lifetime household income percentile ranks as follows. Let $j(i,t)$ indicate i's spouse in year t, and let $h(i,t)$ indicate the household of which i is a part at t. As explained above, current household income is the sum of the head's and spouse's individual current income: $y_{i,t}^h = y_{it} + y_{j(i,t),t}$. Our lifetime concept of household income follows each individual through the sequence of households during their adult life, again using individuals' income fitted from Equation 1 and the subsequent adjustments. That is, the lifetime household income of individual i is

$$\overline{y}_i^h = \sum_t e(\widehat{y}_{it}^h) = \sum_t e(\widehat{y}_{it} + \widehat{y}_{j(i,t),t}) \tag{13}$$

where t is again summed over the years in which i is between ages 18 and 65. The function $e(\cdot)$ equivalizes household income for differences in household size in each year. If we were to restrict our sample to stable households over time (as in, e.g., Fullerton and Lim Rogers, 1993), our definition of household income would coincide exactly with the natural concept. However, it accommodates unstable households in a way that is meaningful as a measure of living standards.

B.3 Simulated Eligibility

Supplemental Nutrition Assistance Program (SNAP). SNAP eligibility is determined on the basis of three tests: (1) a gross-income test, (2) a net-income test, and (3) an asset test. Recipients of TANF and SSI are always categorically SNAP-eligible.

We use state-level gross-income tests from 1996 to 2016 from SNAP Policy Database, maintained by the Economic Research Service of the U.S. Department of Agriculture.[31] We assume these thresholds are unchanged from 2016 through 2019. Until 2000, all U.S. states had a SNAP gross-income test at 130 percent of the Federal Poverty Level (FPL). Under "broad-based categorical eligibility" (BBCE), states raised gross-income limits.

The net-income test requires that income net of specific deductions is less than 100 percent of the FPL. Starting from gross income, all households take a standard deduction as a function of their household size; they also deduct 20 percent of household earnings from gross income. There are four further deductions that may be applied to gross income. We focus on the most important, the "excess shelter deduction." This deduction subtracts housing costs, inclusive of utilities, that exceed half of net income after accounting for all other deductions. The excess-shelter deduction is capped at a level that depends on household size. Standard deductions and excess-shelter deduction caps vary by year but are different for Alaska and Hawaii; we collected these policy parameters from Federal Register notices. The three other deductions—for child support, medical expenses, and dependent care—appear rarely used in eligibility determinations, and we ignore them.[32]

We use asset-test thresholds from the SNAP Policy Database. We apply the asset test rules to household liquid savings, due to the exemption of most relevant other categories of wealth. The asset limit for nonelderly households was $2,000 from the 1980s until 2014, when it was raised to $2,250. The asset test is eliminated under BBCE.

There are special eligibility rules covering households with elderly or disabled adults. In particular, these households are only subject to the net-income test (no gross-income test). They

[31] See https://www.ers.usda.gov/data-products/snap-policy-data-sets/.

[32] For further details, see Center on Budget and Policy Priorities, "A Quick Guide to SNAP Eligibility and Benefits."

also face higher asset-test threshold of $4,250, unless the threshold has been raised under BBCE. We assume the asset-test threshold for such households is the maximum of $4,250 and their BBCE asset-test threshold for all other households.

Medicaid. Medicaid eligibility is determined by income and asset tests that vary by state and with household characteristics. In most states, SSI recipients are categorically Medicaid-eligible; we apply this to states which, under the "209(b)" rules, in principle have some Medicaid eligibility rules that are more stringent than for SSI.

Income eligibility thresholds come primarily from the Kaiser Family Foundation (KFF), with our supplementation to fill gaps in the data. We imputed that thresholds did not change when there are data gaps but we know thresholds on both ends of the gap were the same. Different income tests apply to non-disabled adults, parents, and pregnant women. Income eligibility is most complicated for disabled adults, who may become eligible under a number of pathways, including Medicaid buy-in and being "medically needy." We determine whether a household qualifies as medically needy using reported health expenditures.

We hand-collected Medicaid asset-test thresholds from state-agency websites and policy reports that will be included in our replication files. The thresholds vary for singles and couples, and for the Medicaid buy-in and medically-needy pathways. When we were unable to find state asset-test thresholds in a given year, we imputed it from surrounding years or used the federal thresholds.

Housing Assistance. Eligibility for housing assistance (Section 8 and public housing) is determined by income. Income is measured relative to Area Median Income (AMI) at the level of metropolitan area or non-metropolitan county. As we do not have sub-state geographic identifiers, we use state-level AMIs by household size. Public housing authorities may set their income thresholds between 50 percent and 80 percent of AMI. We assume an eligibility threshold of 50 percent of AMI, as large-city public housing authorities typically impose this threshold at voucher take-up or occupancy of the public-housing unit.

There is no asset test for housing assistance. Until 2014, however, households with no actual asset income but significant wealth could be excluded from housing assistance on the basis of imputed asset income. This imputation used a "passbook savings rate" of two percent until 2014. In 2014, HUD Notice H 2014-15 set this rate to almost zero, essentially eliminating the treatment of assets as income.

Supplemental Security Income (SSI). SSI eligibility is determined by disability of an adult or child member of the household, an income test, and an asset test.

Households are ineligible if their income exceeds a federal "substantial gainful activity" (SGA) threshold. This SGA threshold rose gradually from $500 per month in 1997 to $1,220 in 2019. We also label households ineligible if their countable income exceeds the Federal Benefit Rate (FBR), which implies they would not be eligible for a positive SSI benefit amount. Monthly countable income for SSI is defined by the following formula:

$$y_{\text{countable}} = \max\{0, y_{\text{earned}} + y_{\text{unearned}} - 0.5 \cdot \max\{0, y_{\text{earned}} - 65\} - 20\},$$

where y_{earned} and y_{unearned} are monthly earned and monthly unearned income respectively.

Single-adult households are ineligible for SSI if they possess more than $2,000 in countable assets. The asset threshold is $3,000 for couples. Countable assets are financial assets only after 2005 and financial assets plus the excess of vehicle wealth above $4,500 before 2005.

Women, Infants, and Children (WIC). WIC eligibility is determined by the presence of a child under age five in the household and an income test. The income test is that their income is no greater than 185 percent of the FPL. Households are also categorically WIC-eligible if they have such a child and receive SNAP, TANF, or Medicaid.

Low-Income Heating and Energy Assistance Program (LIHEAP). A household is LIHEAP-eligible if they pay utilities, satisfy an income test, and satisfy an asset test. We determine whether a household pays utilities based on reported utility expenditures.

States set their own income-test thresholds, and these differ by LIHEAP sub-program. Our eligibility simulation focuses on non-crisis heating assistance, the largest sub-program. For 1997–2007 and 2015–2019, we obtain these from the LIHEAP Clearinghouse website, using Internet Archive to obtain the first interval. We obtained the intermediate years from LIHEAP Reports to Congress.

Information was more limited on LIHEAP asset tests. From the Clearinghouse, Reports to Congress, and state-agency websites, we were able to determine whether states had asset tests for all years. The levels of the asset threshold, however, we have only beginning in 2015. We assume these thresholds were unchanged from 1997 to 2015 if the state always had an asset test. For states that had an asset test but eliminated it before 2015, we impute a limit of $5,000. We assume the assets covered by the test are liquid savings, although definitions appear to vary somewhat by state.

States may also make SNAP, SSI, and TANF recipients categorically eligible for LIHEAP. We obtained states' categorical-eligibility rules for fiscal year 2019 from the "Detailed Model Plan" submissions included in their SF-424 grant applications for federal LIHEAP funds. We assume that categorical-eligibility rules are unchanged over the entire period.

School Meals. A household is eligible for the National School Lunch Program and the School Breakfast Program if they have a school-age child (ages 5 to 18) and have an income less than 185 percent of the FPL. We use the threshold to qualify for reduced-price meals. The threshold is 150 percent of the FPL for free meals. Households can also be categorically eligible if they receive

SNAP, TANF, or other means-tested transfers.[33]

The Healthy Hunger-Free Kids Act of 2010 established the Community Eligibility Provision (CEP), which offers free school meals universally in high-poverty areas. We do not account for school-meals eligibility via the CEP, as we lack sub-state geographic identifiers.

Temporary Aid for Needy Families (TANF). We heavily rely on data and eligibility simulations from the Urban Institute's TRIM3 model. We use the following variables from TRIM3: state-year-household size gross income eligibility thresholds; state-year "standard of need" data, which scale the income eligibility thresholds; state-year "adjustment variables," which adjust the standard of need; state-year TANF asset thresholds; state-year earnings disregards (fixed levels and shares of income). Some states use a net asset test; in those cases, we impute their gross threshold as the median non-missing gross standard of need.

A household's income is below the TANF eligibility threshold if their income, less the disregard, is less than the standard of need times the eligibility threshold. A household is eligible for TANF if they have a child, are below the eligibility threshold, and are below the asset threshold.

This eligibility simulation neglects the following forces. First, we do not incorporate the TANF net income thresholds. Second, we assume that the vehicle asset test deducts the full value of the vehicle, which occurs in 39 states of 50 states. Third, some states do not require children in the household to be eligible.

B.4 Data Sources on Budgetary Cost

- *SNAP*: Laura Tiehen, "The Food Assistance Landscape: Fiscal Year 2019 Annual Report," Economic Research Service, U.S. Department of Agriculture, July 2020.
- *Medicaid*: U.S. Centers for Medicare & Medicaid Services, "CMS Office of the Actuary Releases 2019 National Health Expenditures," 16 December 2020.
- *Housing Assistance*: Donna Kimura, "Fiscal 2019 HUD Budget Approved," *Affordable Housing Finance*, 20 February 2019.
- *SSI*: Office of Research, Evaluation, and Statistics and Office of Retirement and Disability Policy, Social Security Administration, "SSI Annual Statistical Report, 2019," SSA Publication No. 13-11827, August 2020.
- *TANF*: Office of Family Assistance, Administration for Children & Families, U.S. Department of Health and Human Services, "TANF and MOE Spending and Transfers by Activity, FY 2019," 22 October 2020.

[33] See Rebecca R. Skinner and Randy Alison Aussenberg, "Overview of ESEA Title I-A and the School Meals' Community Eligibility Provision," Congressional Research Service Report R44568, 2016.

- *WIC*: Food and Nutrition Service, U.S. Department of Agriculture, "WIC Program Participation and Costs," 10 February 2023.
- *LIHEAP*: Office of Community Services, Administration for Children & Families, U.S. Department of Health and Human Services, "LIHEAP DCL Funding Release FY 2019," 26 October 2018.
- *School Meals*: Economic Research Service, U.S. Department of Agriculture, "National School Lunch Program" and "School Breakfast Program," 3 August 2020.

B.5 A Decomposition of Distributional Incidence

We have argued informally through regressions that take-up, more so than eligibility, explains why transfer receipt identifies those with low consumption given income. Here we extend the decomposition approach of Brewer et al. (2020) to carefully distinguish between selection via eligibility rules and selective take-up among the eligible.

Mobility Effect. The mobility effect is the component of the difference between incidence with respect to consumption and incidence with respect to income that results from year-to-year income mobility. For instance, because college graduates have significant debt early in their careers but later greatly out-earn non-graduates, the lifetime incidence of student loan forgiveness is less progressive than its incidence with respect to current income.

To measure the mobility effect, we estimate the share $q(r)$ of people with consumption rank $\overline{R}_i = r$ who would have received a given transfer in a given year if, counterfactually, transfer receipt were only a function of current-income rank. That is, our mobility-only counterfactual is

$$q(r) = \int \Pr(D_{it} = 1 \mid R_{it})\, dF(R_{it} \mid \overline{R}_i = r), \tag{14}$$

where D_{it} indicates transfer receipt and $F(R_{it} \mid \overline{R}_i = r)$ is the conditional distribution of current-income ranks R_{it} at consumption rank r. The mobility effect at rank r is the difference between $q(r)$ and the empirical receipt rate at a current-income rank r, $p(r) = \Pr(D_{it} = 1 \mid R_{it} = 1)$. Intuitively, the mobility effect applies the consumption–income "transition matrix" to the receipt rate by income rank, yielding the counterfactual receipt rates at each consumption rank.

Eligibility Effect. The eligibility effect is the component of the difference in incidence due to eligibility rules that, among low-income households, tag those with low consumption and low lifetime income. For instance, asset tests and categorical eligibility for single parents, people with disabilities, and similar groups tend to target benefits to the persistently poor, not people with merely low current income. By contrast, eligibility rules for contributory social insurance programs such as unemployment insurance operate in the opposite direction, restricting transfers to the persistently poor.

To measure the eligibility effect, we define the receipt rate at consumption rank r under a second counterfactual, $s(r)$. Here the probability of take-up conditional on eligibility is a function of current income alone, whereas we let eligibility depend on both current income and consumption (or current income and lifetime income). This counterfactual receipt rate $s(r)$ thus depends only on eligibility and current-income rank. The difference between these two counterfactuals, $q(r)$ and $s(r)$, is the eligibility effect.

When eligibility can be measured perfectly—that is, $E_i = 0$ implies $D_i = 0$—our second counterfactual is defined as

$$s(r) = \int \Pr(D_{it} = 1 \mid R_{it},\, E_{it} = 1) \Pr(E_{it} = 1 \mid R_{it},\, \overline{R}_i = r)\, dF(R_{it} \mid \overline{R}_i = r),$$

where E_{it} indicates eligibility. This expression emerges from the following reasoning. The probability $\Pr(D_{it} = 1 \mid R_{it},\, E_{it} = 1)$ is the transfer take-up rate among the eligible at current-income rank R_{it}. The probability $\Pr(E_{it} = 1 \mid R_{it},\, \overline{R}_i = r)$ is the eligibility rate at current-income rank R_{it} and consumption rank r. Multiplying these two probabilities yields a predicted receipt rate for people with current-income rank R_{it} that embeds the desired independence assumption: Given current income, take-up among the eligible is uninformative about consumption or lifetime income.[34] Integrating over the conditional distribution of current-income ranks given consumption, $F(R_{it} \mid \overline{R}_i = r)$, we have a counterfactual receipt rate $s(r)$ at consumption rank r that permits a role for eligibility rules while shutting down selective take-up.

Imperfect measurement of eligibility requires a more complex expression for the same counterfactual concept. However, the counterfactual answers the same question: What would the distributional impact of a transfer program be if take-up, conditional on eligibility and income, was uninformative about consumption? This counterfactual is

$$\begin{aligned} s(r) = \int \Big[&\Pr(D_{it} = 1 \mid R_{it},\, E_{it} = 0) \Pr(E_{it} = 0 \mid R_{it},\, \overline{R}_i = r) \\ &+ \Pr(D_{it} = 1 \mid R_{it},\, E_{it} = 1) \Pr(E_{it} = 1 \mid R_{it},\, \overline{R}_i = r) \Big]\, dF(R_{it} \mid \overline{R}_i = r), \end{aligned}$$

where all terms are as defined above. We use this counterfactual in our analysis. As a robustness check, we also use the earlier counterfactual that requires perfect measurement of eligibility and

[34] To better understand this counterfactual, note that the law of conditional probability implies that: $\Pr(D_{it} = 1 \mid \overline{R}_i = r) = \int \Pr(D_{it} = 1 \mid R_{it},\, \overline{R}_i = r)\, dF(R_{it} \mid \overline{R}_i = r)$. By consequence, the counterfactual receipt rate $s(r)$ equals the empirical receipt rate at a given consumption rank if and only if take-up among eligibles is conditionally independent of consumption given current income at all consumption ranks.

reclassify simulated-ineligible recipients as eligible.

Take-Up Effect. The take-up effect is the component of the difference in incidence that results from the influence of consumption and lifetime income on take-up rates among the eligible. For instance, procedural complexity (Deshpande and Li, 2019; Gray, 2019), "ordeal" costs (Nichols and Zeckhauser, 1982), and information frictions (Finkelstein and Notowidigdo, 2019) could select for or against low lifetime income and consumption among eligible households with similar current incomes. We define the take-up effect as the residual difference between the counterfactual receipt rate $s(r)$ and the empirical receipt rate at the same consumption rank, $\overline{p}(r) = \Pr(D_{it} = 1 \mid \overline{R}_i = 1)$.

B.6 Results of Decomposition

This decomposition enables us to interpret differences between the consumption incidence and the income incidence of transfer programs. For instance, why do a greater or smaller share of people at consumption rank r receive a transfer than do people at current-income rank r? Our decomposition is

$$\overline{p}(r) - p(r) = \underbrace{[\overline{p}(r) - s(r)]}_{\text{take-up effect}} + \underbrace{[s(r) - q(r)]}_{\text{eligibility effect}} + \underbrace{[q(r) - p(r)]}_{\text{mobility effect}}, \tag{15}$$

using the objects defined above. Figure A12 displays decomposition results for the difference between current-income and consumption incidence. Panel A displays results for the average total annual per-capita value of benefits, and Panel B displays results for receipt rates by program. The blue, black, and yellow lines plot the mobility, eligibility, and take-up effects, respectively; each represents the difference in consumption incidence with respect to current-income incidence if only this effect were present. The blue shaded region indicates the net effect, equal to the difference between the consumption and current-income incidence at the indicated rank; this region integrates to zero.

Absent eligibility and take-up effects, the incidence of transfer programs would be markedly less progressive at the bottom of the distribution. Consistent with our prior results, we find a central role for selective take-up among eligible households, and a modest role for eligibility rules, in shaping the consumption incidence of transfer programs.

In total over transfer programs, selective take-up increases the average annual value of benefits received by the consumption-poorest people by about \$750 per person relative to a counterfactual in which transfer receipt is a function of current income and eligibility alone. Eligibility rules contribute about \$250 per person at the bottom of the consumption distribution. Together these amounts represent about one quarter of the average transfer per capita per year to consumption-poorest people. Mobility tends to shift the incidence of transfers up the consumption distribution, but bograms, the eligibility and take-up effects largely offset the mobility effect.

The primacy of take-up applies broadly across the eight transfer programs we study. Eligibility rules appear most important in Medicaid, TANF, and WIC, whereas they appear essentially irrelevant or operating in the opposite direction for SNAP, LIHEAP, and housing assistance. While mobility alone would reduce the receipt rate among the consumption poor by about 10 percentage points for SNAP, Medicaid, and housing assistance, selective take-up increases receipt rates in these programs at the bottom of the consumption distribution by 10 to 20 percentage points. These are economically large differences relative to underlying rates of program participation as well as relative to the effects of mobility.

B.7 Measurement Error: Simulation

Method. How much measurement error is required to overturn our results? We conduct an adversarial simulation to probe the robustness of Figure 1 to extreme amounts of measurement error. We consider the coefficient γ in Equation 4, which represents the marginal effect of take up of a given transfer program on lifetime rank, controlling for current rank. We simulate measurement error as follows:

1. Obtain the take-up rate among the bottom 50% of households ranked in consumption terms, $\hat{D} \in [0, 1]$.

2. Assign the top x% to have some constant $c \in [0, 1]$ times the take-up rate of the bottom 50%: $c\hat{D}$.

3. Estimate Equation 4 using the simulated data.

This exercise generates a large amount of measurement error at the top of the distribution. The take-up rate in the bottom 50% of the current consumption distribution is a natural bound on the take-up rate of the top x% of the consumption distribution, unless the programs' targeting properties are very perverse. The parameter c governs whether the measurement error is as severe as possible ($c = 1$).

Results. Measurement error would need to be very severe to overturn our results. Figure A10 shows the estimated coefficient $\hat{\gamma}$ as a function of the share of the top of the consumption distribution that has severe measurement error. In black, we present the estimates if the true take-up rate is half the bottom 50%'s take-up rate. The blue lines show the estimates if the true take-up rate is the same as the bottom 50%'s take-up rate. When $x = 0$, the estimates coincide with Figure 1. As long as true take-up at the top is half the poorest's take-up rate, we continue to reject $\gamma = 0$. If true take-up is equal, then we can no longer reject $\gamma = 0$ for $x \geq 15$ or so. These results are logical: if we impute take-up rates that are the same as at the bottom of the distribution for much of the top of the distribution, we no longer find evidence of selection. But as long as measurement error does not exceed half the take-up rate, we decisively reject the null.

C Theory Appendix

C.1 Alternative Reforms

In this section we consider alternate reforms and analyze their welfare effects.

Alternate Reform. We marginally reduce the voluntary transfer schedule by ds at all incomes. This finances an expansion of a universal basic income, i.e. a lump sum tax decrease of $\tau = \int_z [dsM(z) + S(z)m(z)]\, h(z)dz$. The decrease offsets the static cost of the decreased benefit to the inframarginals and fiscal cost from marginal recipients. Because this reform imposes lump-sum changes to both the transfer and income tax, it has no efficiency effects via labor supply.

$$
\begin{aligned}
dW_{\text{UBI}} = & \underbrace{ds \int_z M(z) \left(E_{\kappa \leq S(z)} [\alpha(z,\kappa)] - E_\kappa [\alpha(z,\kappa)]\right) h(z)dz}_{\text{reduced redistribution within } z} \\
& + \underbrace{ds \int_z E_\kappa [\alpha(z,\kappa)] \left(M(z) - E_{z'} [M(z')]\right) h(z)dz}_{\text{reduced redistribution between } z} \\
& - \underbrace{E_z [(S(z))m(z)]\, E_{z,\kappa} [\alpha(z,\kappa)]}_{\text{fiscal saving from marginals}} .
\end{aligned} \tag{16}
$$

The only new term in Equation 16 is the redistribution between z term. To speak to this, we define an alternative notion of advantageous selection.

Definition 2. *We say households are* advantageously selected between income *when the distribution of costs κ conditional on income z increases in z in the sense of first-order stochastic dominance.*

Between-income advantageous selection evokes Nichols and Zeckhauser's (1982) targeting argument for transfers. If take-up costs increase in z, then take-up rates are declining in income, and hence transfers will have positive redistributive effects between incomes. With this definition, we can sign the redistribution between term.

Proposition 2. *Assuming the transfer $S(z)$ is positive and weakly decreasing, the welfare effect of between-income redistribution is negative if there exists advantageous selection as well as between-income advantageous selection.*

This proposition shows that voluntary transfer program has an advantage over a UBI. A dollar of voluntary transfer flows in expectation to income levels with lower application costs. Assuming advantageous selection between incomes, the voluntary transfer is taken-up relatively more often by lower incomes. The UBI reform therefore redistributes regressively: A dollar is taken disproportionately from lower incomes and given to the average individual.

This suggests an efficiency motive for voluntary transfers. For any tax change that is redistributed as an automatic transfer, one could instead redistribute as a conditional transfer and the recipients would be strictly better targeted. The price of targeting precision is the real take-up cost that must be incurred to enroll. The following proposition shows that as long as some people have zero cost, voluntary transfers should exist in the planner's optimum.

Proposition 3. *Suppose there is a mass point of individual's with cost zero at each income z. If we have advantageous selection and between advantageous selection starting at $S(z) = 0$, then optimally there is a non-zero transfer system.*

This proposition echoes Nichols and Zeckhauser's (1982) rationale. So long as the first dollar of voluntary transfer has infinitismally improved targeting properties compared to costs, then the planner should utilize this tool to relax screening constraints that frustrate redistribution.

UBI Reform We remove the entire voluntary transfer program $S(z)$ and expand the automatic transfer program by $M(z)S(z)$ in its place.

We calculate the welfare impacts of this reforms below, as before assuming no income effects,

$$\begin{aligned}\Delta W_{\text{UBI}} = &\underbrace{\int_z S(z)M(z)\left(E_\kappa\left[g(z,\kappa)\right] - E_{\kappa\leq S(z)}\left[g(z,\kappa)\right]\right)h(z)dz}_{\text{reduced redistribution within } z}\\ &+ \underbrace{E_{z,\kappa\leq S(z)}\left[\kappa g(z,\kappa)\right]}_{\text{fiscal savings/ordeal costs}}\\ &+ \mathscr{L},\end{aligned} \tag{17}$$

where $\mathscr{L}$ contains all labor-supply effects, which are additively separable. The perturbation approach does not admit a simple closed-form solution for $\mathscr{L}$ since changes in the slopes of the voluntary transfer and tax schedules are not infinitesimal.

If the voluntary transfer scheme is entirely replaced, the primary benefit is that ordeal costs paid by everyone in the program are removed. The main cost is that there is a redistribution from always-takers (i.e. those inframarginal to the old voluntary transfer schedule) to never-takers. When the always-takers have higher social marginal welfare weights than never-takers — as our empirical analysis demonstrates is the case — this redistributive effect is bad for welfare. As noted in the text, the ordeal costs term requires estimating the *average* ordeal cost among inframarginal always-takers, not just the ordeal cost among compliers. This cost is difficult to estimate without auxiliary assumptions, since we lack data on inframarginals' willingness to pay to avoid ordeals.

C.2 Formal Analysis of the Consumer's Problem

To ensure that there is no bunching and that labor supply elasticities are well defined, we assume every consumer of type w chooses their pre-tax income z before their realization of κ is drawn. With the quasi-linearity assumptions, this means that labor supply will be chosen according to expected amount of social program dollars, $M(z)S(z)$ that the consumer expects to accrue.

Each individual maximizes

$$\max_{z} z - T(z) + \int_{0}^{S(z)} (S(z) - \kappa)\mu(w \mid \kappa)d\kappa - v(z/w). \tag{18}$$

Hence the consumer's FOC is

$$0 = 1 - T'(z) + S'(z)M(z) - v'(z/w). \tag{19}$$

We use the notation $M(z)$ as take-up of the transfer, post income choice z is of primary interest. But for an individual with type (w), there is no causal effect of z on $M(z)$ beyond $S(z)$ changing. The distriubition of that κ is fixed by w, and hence even if type w expands their labor supply choice, absent any $S(z)$ change, there will be no change to the probability of $\kappa \geq S(z)$.

Next, we calculate individual elasticities derivatives of labor supply with respect to tax and SNAP changes in terms of primitives. Following Jacquet and Lehmann (2014), we apply perturbations to the tax and snap system about z_0 of the form $\hat{T} = T + \tau(z - z_0) - \nu$ and $\hat{S} = S + \varsigma(s - s_0) - \vartheta$. In each case, the marginal tax (transfer program) rate has increased by τ (ς) but the level has decreased by ν (ϑ).

Now the consumer's problem is

$$\max_{z} z - T(z) - \tau(z - z_o) - \nu + \int_{0}^{S(z)+\varsigma(z-z_0)-\vartheta} (S(z) + \varsigma(z - z_0) - \vartheta) - \kappa)\mu(w \mid \kappa)d\kappa - v(z/w). \tag{20}$$

The new FOC is

$$\mathcal{F} = 1 - T'(z) - \tau + \int_0^{S(z)+\varsigma(z-z_0)-\vartheta} (S'(z) + \varsigma)\mu(w \mid \kappa)d\kappa - v'(z/w) \tag{21}$$

$$= 1 - T'(z) - \tau + (S'(z) + \varsigma)Pr\,(z + \varsigma(z - z_0) - \vartheta \leq \kappa \mid w) - v'(z/w). \tag{22}$$

To use the implicit function theorem, we calculate the derivatives:

$$\mathcal{F}_z \mid_{\tau=\varsigma=0} = -T''(z) - \frac{v''(z/w_i)}{w_i^2} + S''(z)M(z) + S'(z)^2 m(z) \tag{23}$$

$$\mathcal{F}_\tau \mid_{\tau=\varsigma=\vartheta=0} = -1 \tag{24}$$

$$\mathcal{F}_\varsigma \mid_{\tau=\varsigma=\vartheta=0} = M(z), \tag{25}$$

noting that $z \longrightarrow z_0$ as $\tau, \varsigma, \vartheta \longrightarrow 0$. Hence by the implicit function theorem we have that

$$\frac{\partial z}{\partial \tau} = \frac{-1}{\mathcal{F}_z} \tag{26}$$

$$\frac{\partial z}{\partial \varsigma} = \frac{M(z)}{\mathcal{F}_z}, \tag{27}$$

where we have assumed $S(z)$ income effects are small: $\frac{\partial z}{\partial \vartheta} = 0$.

And the (compensated) elasticities needed are then defined as:

$$\epsilon^z = -\frac{\partial z}{\partial \tau} \mid_{\tau=0} \frac{1 - T'(z)}{z} \tag{28}$$

$$\epsilon^s = \frac{\partial z}{\partial \varsigma} \mid_{=0} \frac{S'(z)}{z}. \tag{29}$$

C.3 Proofs

C.3.1 Proof of Proposition 1

Proof. As in Section C.2 the reform is composed of perturbations to the tax and snap system about z_0 of the form $\hat{T} = T + \tau(z - z_0) - \nu$ and $\hat{S} = S + \varsigma(s - s_0) - \vartheta$. In particular, the reform is composed of a level shift in $S(z)$ of $\vartheta = ds$, a change in the marginal tax rate of $\tau = \frac{d}{dz}M(z)ds$, [35] and decreases in everyone's level of taxes $\nu = E_z\left[S(z)m(z)\right] + M(z)ds$ due to the reduced voluntary transfer

[35]Note that the derivative $\frac{d}{dz}M(z)$ is a total derivative from the planner's perspective, i.e. shifting between people at different z's. It includes changes the $M(z)$ in z due to both the the schedule $S(z)$ varying in z and the distribution of $\kappa \mid w$ varying in z.

expenditure. Finally, any changes in revenue due to changes in labor supply are redistributed lump sum. We analyse these in turn. Throughout, for convenience, we integrate over z instead of w, where z is the pre-reform income that is one to one with w, per the consumer's FOC in Section C.2.

Writing $V^* = V^*(z^*, w, \kappa)$ for the consumers pre-reform optimized utility function, since $W = \int_z \int_\kappa \alpha(w,\kappa) V^*(z^*, w, \kappa) d\mu$, it remains $\frac{\partial W}{\partial \tau} = \int_z \int_\kappa \alpha(w,\kappa) \frac{\partial}{\partial \tau} V^*(z^*, w) d\mu$ and similarly for ς and ϑ.

This total derivative is generally of the form

$$\frac{dV^*}{d\tau} = \frac{\partial V^*}{\partial \tau} + \frac{\partial V^*}{\partial z}\frac{\partial z}{\partial \tau}.$$

For ν we have

$$\frac{dV^*}{d\nu} = \frac{\partial V^*}{\partial \nu} + \frac{\partial V^*}{\partial z}\frac{\partial z}{\partial \nu} = \frac{\partial V^*}{\partial \nu}$$

since $\frac{\partial z}{\partial \nu} = 0$ by the assumption of no income effects. Hence we need only calculate the direct effect. Identically for $\frac{dV^*}{d\vartheta} = \frac{\partial V^*}{\partial \vartheta}$.

However, in the case of τ, we have $\frac{\partial V^*}{\partial \tau} = 0$ (slope changes do not have a direct utility effect, yet $\frac{\partial V^*}{\partial z}\frac{\partial z}{\partial \tau} \neq 0$ as there are substitution effects, and the model does not admit an envelope theorem.

In particular, note that $\frac{\partial V^*}{\partial z} = 1 - T'(z) - v'(z/w) + \mathbb{1}(\kappa \leq S(z)) \cdot S'(z)$ whereas $\frac{\partial U^*}{\partial z} = 1 - T'(z) - v'(z/w) + M(z)S'(z)$. No envelope theorem applies since when the individual averages over κ they do so without weights, whereas the planner averages over κ according to weights $\alpha(w,\kappa)$ To evaluate the change in utility, we subtract off

$$\frac{\partial U^*}{\partial z} = 1 - T'(z) - v'(z/w) + M(z)S'(z) = 0$$

by the individual's FOC in Section C.2. Hence we have:

$$\begin{aligned}
\int_\kappa \alpha(w,\kappa)\frac{\partial}{\partial \tau} V^*(z^*, w) d\mu &= \int_\kappa \alpha(w,\kappa)\frac{\partial z}{\partial \tau}\left(\frac{\partial V^*}{\partial z}\right) d\mu \\
&= \frac{\partial z}{\partial \tau}\int_\kappa \left(\alpha(w,\kappa)\frac{\partial V^*}{\partial z} - E_\kappa\left[\alpha(w,\kappa)\right]\frac{\partial U^*}{\partial z}\right) d\mu \\
&= \frac{\partial z}{\partial \tau}\int_\kappa \left(\alpha(w,\kappa)\mathbb{1}(\kappa \leq S(z)) \cdot S'(z) - E_\kappa\left[\alpha(w,\kappa)\right] M(z)S'(z)\right) d\mu \\
&= \frac{\partial z}{\partial \tau} S'(z) E_\kappa\left[\left(\alpha(w,\kappa)\mathbb{1}(\kappa \leq S(z)) - E_\kappa\left[\alpha(w,\kappa)\right] M(z)\right)\right] \\
&= \frac{\partial z}{\partial \tau} S'(z) Cov_\kappa\left[\alpha(w,\kappa), \mathbb{1}(\kappa \leq S(z))\right].
\end{aligned}$$

The last line follows by noting the $E_k\left[\mathbb{1}(\kappa \leq S(z))\right] = M(z)$.[36] Integrating over z yields the final term in the welfare formula.

It remains to count the direct effects of changes T and S levels, and any fiscal consequences of labor supply changes.

The total direct tax changes that accrue to pre-reform income level z are a reduction of tax of $M(z)ds$. This has a welfare effect of $E_{w,\kappa}\left[g(w,\kappa)M(z)\right]ds$ where implicitly the FOC defines a unique $z = z(w)$.

This is counterbalanced with an reduction of $S(z)$ by ds at all incomes. Welfare can be written as

$$\int_z \left(\int_0^{S(z)} \alpha(w,\kappa)(S(z)-\kappa)\mu d\kappa + \int_\kappa \alpha(w,\kappa)\left(z - T(z) - v(z/w)\right)\mu d\kappa \right) dw.$$

Hence, an expansion of (z) by ds has a direct effect of $\int_z S'(z) \int_0^{S(z)} \alpha(w,\kappa) d\mu = \int_z S'(z) E_{\kappa \leq S(z)}\left[\alpha(w,\kappa)\right] M(z)$. Again since w and z are in one to one correspondence, we switch the integration label and arrive at the other half of the first term in the propositon.

In addition there is a lump-sum transfer to all z due to the fiscal savings from marginal recipients. As $S(z)$ declines by ds, the fiscal saving is $E_z\left[S(z)m(z)\right]$ which accrues to everyone, at welfare gain of

$$E_z\left[S(z)m(z)\right] \int_z \int_\kappa g(w,\alpha) d\mu$$

which is precisely the second term in the proposition.

Finally, since marginal tax rates have risen (as $M'(z) < 0$), labor supply contracts by $dsM'(z)\frac{\partial z}{\partial \tau}$ at each income z. The fiscal cost per $d\tau$ is $T'(z) - \frac{d}{dz}\left(S(z)M(z)\right)$. Plugging in the definition of the elasticity: $\epsilon^z = -\frac{\partial z}{\partial \tau}\frac{1-T'(z)}{z}$ and by assumption all fiscal costs are paid lump sum, i.e. by the average welfare weight $E_{w,\kappa}\left(\alpha(g,w)\right)$, yields the final term in proposition 1. This completes the proof.

□

C.3.2 Proof of Proposition 4

Proposition 4. *Assuming the transfer $S(z)$ is positive, the first term in Equation 9 is negative if there exists advantageous selection, implying lost value from self-targeting.*

Proof. Directly from the definition of advantageous selection we get that $E_{z,\kappa \leq S(z)}\left[\alpha(z,\kappa)\right] > E_{z,\kappa)}\left[\alpha(w,\kappa)\right]$. This implies the integrand is positive for all values of z, and hence positive overall. □

[36]Since each component is a decreasing function of k, we conclude by Schmidt (2003) that the covariance is positive.

C.3.3 Proof of Proposition 5

Proposition 5. *Suppose (1) the tax system is optimal and (2) take-up decreases in income (i.e., $M'(z) < 0$). Then the sum of the labor supply effects (i) and (ii) in Equation 9 is negative.*

Proof. We derive a neccesary condition from the optimal tax schedule that ensures the labor supply effect is signed as proposed. Suppose the planner increases the tax rate at income z by $d\tau$, and that the net fiscal gain/loss from this change is redistributed as a lump sum transfer/tax.

Breaking down the effect into fiscal and behavioural responses, we have the following changes to welfare:

1. Direct effect (fiscal and welfare):

$$d\tau \int_{x\geq z} (E_{z,\kappa}(\alpha(z,\kappa)) - E_{\kappa}(\alpha(x,\kappa))))dH(x) = \mathbb{E}_{x\geq z}\left[E_{z,\kappa}(\alpha(z,\kappa) - E_{\kappa}(g(x,\kappa))\right](1-H(z)). \tag{30}$$

2. Compensated price effect (effect on tax receipts):

$$d\tau \frac{\partial z}{\partial \tau} \cdot [T'(z)] \cdot h(z) E_{z,\kappa}(\alpha(z,\kappa) = d\tau \left(-\epsilon^{z} \cdot z \cdot \frac{T'(z)}{1-T'(z)} h(z) E_{z,\kappa}(\alpha(z,\kappa))\right). \tag{31}$$

3. Compensated price effect (effect on social program payments):

$$d\tau \frac{\partial z}{\partial \tau} \frac{d}{dz}\left[-S(z)M(z)\right] h(z) E_{z,\kappa}(\alpha(z,\kappa)) = d\tau \frac{\epsilon^{z} \cdot z}{1-T'(z)} \frac{d}{dz}\left[S(z)M(z)\right] h(z) E_{z,\kappa}(\alpha(z,\kappa)). \tag{32}$$

4. Non-envelope effect (welfare):

$$d\tau \int_{\kappa} \alpha(w,\kappa) \frac{\partial z}{\partial \tau}\left(\frac{\partial V^{*}}{\partial z}\right) d\mu = -d\tau \frac{\epsilon^{z} \cdot z}{1-T'(z)} S'(z) Cov_{k}\left[\alpha(w,k), \mathbb{1}(\kappa \leq S(z))\right]. \tag{33}$$

The equality of term 33 follows from the working in the proof of proposition 1.

A necessary condition for the optimality of the tax system is that the sum of these welfare effects weakly negative. In particular, to convert to utility units, suppose the net fiscal gain/loss from this MTR change was redistributed as a lump sum. This need not be the optimal way to redistribute, but for optimality it cannot deliver a positive welfare benefit.

$$\mathbb{E}_{x\geq z}\left[E_{z,\kappa}(\alpha(z,\kappa)) - E_{\kappa}(g(x,\kappa))\right](1-H(z)) \leq d\tau\Big[E_{z,\kappa}(\alpha(z,\kappa))\left(T'(z) - \frac{d}{dz}\left[S(z)M(z)\right]\right) \tag{34}$$

$$+ S'(z)Cov_k\alpha(w,k), \mathbb{1}(\kappa \leq S(z))\Big]\left(\frac{h(z)\epsilon^z z}{1-T'(z)}\right) \tag{35}$$

The left hand side is positive, and therefore so must the right hand side be positive. Immediately we have then that the sum of the labor supply terms is negative, as hypothesized.

□

C.3.4 Proof of Proposition 2

Proof. The positivity of the redistribution within z term follows directly from the definition of within-advantageous selection (definition 1).

It remains to show the positivity of the redistribution between z term under within- and between-advantageous selection. Write

$$f_1(z) = -E_\kappa(\alpha(z,\kappa)) \tag{36}$$

and

$$f_2(z) = -\left(M(z) - E_{z'}(M(z'))\right) \tag{37}$$

such that the redistribution between z term can be written as

$$\int_z E_\kappa\left[\alpha(z,\kappa)\right]\left\{M(z) - E_{z'}\left[M(z')\right]\right\}h(z)dz = \int_z f_1(z)f_2(z)h(z)dz = E_z\left[f_1(z)f_2(z)\right]. \tag{38}$$

By the assumption of between-adverse selection, we have that the distribution of $\kappa \mid z$ increases in z in the FOSD sense. Since $M(z)$ is precisely the CDF of $\kappa \mid z$ evaluated at $S(z)$, and $S(z)$ is weakly decreasing, immediately we have that $M(z)$ decreases in z and hence $f_2(z)$ increases in z.

By within-adverse selection, $\alpha(z,\kappa)$ decreases in κ, hence $-\alpha(z,\kappa)$ increases in κ. Then again by FOSD, we have that $E_\kappa(-\alpha(z,\kappa)) = f_1(z)$ increases in z.

Since $f_1(z)$ and $f_2(z)$ are both increasing functions of z, it follows from Schmidt (2003) that $Cov(f_1(z), f_2(z)) \geq 0$. Noting that $E_z f_2(z) = 0$ we have

$$E_z\left[f_1(z)f_2(z)\right] = E_z\left[f_1\right]E_z\left[f_2(z)\right] + \mathrm{Cov}(f_1, f_2) = \mathrm{Cov}(f_1, f_2) \geq 0 \tag{39}$$

as required. □

C.3.5 Proof of Proposition 3

Proof. Suppose, for a contradiction, that optimally $S(z) = 0$ for all z. We implement the opposite of the reform described in the main text. That is, we increase the transfer to $S(z) = ds$ for all z, and taxes rise by $\tau(z) = M(z)ds$ at income z.

For now ignore labor supply changes. By assumption, there are no take-up costs for this mass of individuals. Thus each $\kappa = 0$ individual receives a dollar, paid for by everyone at their income level z. The net welfare effect of this is $\int_z M(z) \{\alpha(z, \kappa = 0) - E_\kappa [\alpha(z, \kappa)]\} h(z)dz$, the redistribution within term. Assuming advantageous selection within income, this term is positive.

Next, consider the effects of the altered labor supply choices in response to the tax changes. The marginal rates have changed by $dsM'(z)$ which is negative due to advatageous selection between incomes. Hence labor supply everywhere increases. By the envelope theorem, there are no first order utility effects to this. There is a positive fiscal externality, which further increases the welfare effect of this reform.

In sum, the reform has a strictly positive effect. Hence $S(z) = 0$ cannot have been optimal. □

References for Appendices

Brewer, Mike, Robert Joyce, Tom Waters, and Joseph Woods, "A Method for Decomposing the Impact of Reforms on the Long-Run Income Distribution, with an Application to Universal Credit," *Economics Letters*, 2020, *192*, 109230.

Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson, "Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector," *American Economic Review*, 2016, *106* (8), 2110–44.

Davern, Michael E, Bruce D Meyer, and Nikolas K Mittag, "Creating Improved Survey Data Products Using Linked Administrative-Survey Data," *Journal of Survey Statistics and Methodology*, 2019, *7* (3), 440–463.

Deshpande, Manasi and Yue Li, "Who is Screened Out? Application Costs and the Targeting of Disability Programs," *American Economic Journal: Economic Policy*, 2019, *11* (4), 213–48.

Finkelstein, Amy and Matthew J Notowidigdo, "Take-Up and Targeting: Experimental Evidence from SNAP," *Quarterly Journal of Economics*, 2019, *134* (3), 1505–1556.

Fullerton, Don and Diane Lim Rogers, *Who Bears the Lifetime Tax Burden?*, Brookings Institution Washington, DC, 1993.

Gray, Colin, "Leaving Benefits on the Table: Evidence from SNAP," *Journal of Public Economics*, 2019, *179*, 104054.

Haider, Steven and Gary Solon, "Life-Cycle Variation in the Association Between Current and Lifetime Earnings," *American Economic Review*, 2006, *96* (4), 1308–1320.

Jacquet, Laurence and Etienne Lehmann, "Optimal Nonlinear Income Taxation with Multidimensional Types: The Case with Heterogeneous Behavioral Responses," THEMA Working Papers 2014-01, Université de Cergy-Pontoise 2014.

Meyer, Bruce D and James X Sullivan, "Consumption and Income Inequality in the US Since the 1960s," *Journal of Political Economy*, 2023, *131* (2).

Mittag, Nikolas, "Correcting for Misreporting of Government Benefits," *American Economic Journal: Economic Policy*, 2019, *11* (2), 142–64.

Morris, Carl N, "Parametric Empirical Bayes Inference: Theory and Applications," *Journal of the American Statistical Association*, 1983, *78* (381), 47–55.

Nichols, Albert L and Richard J Zeckhauser, "Targeting Transfers Through Restrictions on Recipients," *American Economic Review*, 1982, *72* (2), 372–377.

Rubin, Donald B, "The Bayesian Bootstrap," *The Annals of Statistics*, 1981, pp. 130–134.

Schmidt, Klaus D, *On the covariance of monotone functions of a random variable*, Professoren des Inst. für Math. Stochastik, 2003.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," *Quarterly journal of Economics*, 2018, *133* (3), 1331–1393.