

Self-Targeting in U.S. Transfer Programs*

Charlie Rafkin
Adam Solomon
Evan J. Soltas

September 2025

Abstract

This paper explores a classic rationale for why some transfer programs require take-up rather than enroll people automatically: Take-up may be advantageously “self-targeted” on characteristics that cannot be used as eligibility criteria. We find self-targeting on consumption and lifetime income in eight U.S. transfers, as recipients are needier on average than observably-similar eligible nonrecipients. Compared to automatic receipt, self-targeting focuses redistribution toward the lifetime poor, and it also modestly raises the within-lifetime insurance value of transfer dollars. In several transfers, these social benefits of self-targeting appear to offset the social costs of take-up, challenging some economic arguments against ordeals.

Keywords: transfer programs, take-up, eligibility, self-targeting

JEL Codes: H23, H53, I38

*Rafkin: UC Berkeley (craffkin@berkeley.edu). Solomon: MIT (adamsol@mit.edu). Soltas: Princeton University (esoltas@princeton.edu). Our editor Rachel Griffith and three anonymous referees provided exceptional comments. We are very grateful to David Autor, Amy Finkelstein, Jim Poterba, and Frank Schilbach for guidance. We also thank Jenna Anders, Jon Cohen, Stefano DellaVigna, Joel Flynn, Jon Gruber, Nathan Hendren, Antoine Levy, Bruce Meyer, Abby Ostriker, Dev Patel, Anna Russo, Ben Sprung-Keyser, Michel Strawczynski, Iván Werning, and particularly Hunt Allcott for helpful conversations. Rohan Jha and Allan Lee provided excellent research assistance. 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

The take-up of many means-tested transfers is voluntary, and in the U.S., a substantial share of transfer-eligible people does not receive benefits (Currie, 2006; Ko and Moffitt, 2024). Instead of relying on voluntary take-up, the government could provide benefits automatically to all eligible people, adjusting benefit levels to hold spending constant. Automatic provision would eliminate the costs recipients incur to claim benefits, but it would also forgo a potential advantage of voluntary transfers: self-targeting. Self-targeting occurs when selective take-up among the eligible implicitly reveals dimensions of need that the government cannot incorporate into eligibility rules. How much self-targeting occurs in U.S. transfer programs, and can it justify why they remain voluntary?

Answers to these questions would inform contentious debates over the role of voluntary take-up in the U.S. social safety net. Critics of “administrative burdens” promote reforms to raise take-up, including automatic enrollment, over alternatives that would increase benefits or expand eligibility (e.g., Herd and Moynihan, 2019). Some of their proposed reforms would build upon U.S. policy experimentation in the Covid-19 pandemic, when two programs—Medicaid and school meals—became partly automatic amid a temporary expansion of the safety net.¹ Across countries, safety nets vary in their breadth of eligibility and their use of burdens.² The issue of voluntary versus automatic redistribution also bears on other longstanding issues, such as the complexity of eligibility rules or fundamental reforms like a negative income tax or a basic income.

The classic theoretical rationale for voluntary transfers is that selection into take-up may be “advantageous.” That is, a household’s choice to take up a voluntary transfer in the face of costs or “ordeals” may reveal that it has a higher level of unobservable need (Nichols and Zeckhauser, 1982; Besley and Coate, 1992). However, there are also prominent arguments against voluntary transfers. Ordeals may instead perversely screen out the neediest households, who may face greater take-up costs or behavioral frictions (Currie and Gahvari, 2008; Mullainathan and Shafir, 2013; Herd et al., 2023). Ordeals may also reduce the insurance value of transfers if they deter take-up among households hit by adverse shocks (see, e.g., Blundell et al., 2008). And ordeals indisputably impose social costs on inframarginal recipients. Whether, in actual transfers, the benefits of self-targeting offset ordeal costs thus requires both new theory and careful measurement.

This paper studies self-targeting in U.S. transfer programs. We first measure the extent of self-targeting in eight transfers that constitute together most of the U.S. safety net, taking consumption and lifetime income as proxies for need. Next, we analyze theoretically how marginal changes

¹Before the end of Medicaid auto-enrollment, U.S. Department of Health and Human Services (2022) forecasted that 8.2 million people would lose benefits, of which 6.8 million (83 percent) were expected to be from non-take-up among the eligible. Some policy reforms have already occurred in the context of school meals: As of spring 2024, ten states have extended the pandemic-era expansion (see <https://frac.org/healthy-school-meals-for-all>).

²This topic is discussed in international comparisons of welfare states (e.g., Esping-Andersen, 1990; Bartels and Neumann, 2021). However, most countries see incomplete take-up in their voluntary programs (Eurofound, 2015).

to voluntary and automatic transfers affect social welfare, deriving formulas that characterize the trade-off in program design between targeting benefits and ordeal costs. Last, we calibrate these welfare formulas using values from our data and external estimates.

The first part of our analysis measures self-targeting. Using data from the Panel Study of Income Dynamics (PSID) and the Consumer Expenditure Survey (CEX), we compare the consumption levels of transfer recipients to those of eligible nonrecipients with similar current annual incomes. We also examine self-targeting with respect to lifetime income. These selection measures follow in a tradition in public finance dating to [Vickrey \(1947\)](#), one that takes as a premise that consumption is a superior measure of living standards to current income, insofar as households are at least partially insured against transient income shocks.³ From this perspective, a key limitation on redistribution and social insurance is the absence of nonlinear taxes on consumption and lifetime income, likely reflecting practical or political constraints. Self-targeting may relax these constraints. Furthermore, with data on both consumption and lifetime income, we can distinguish how self-targeting affects between-lifetime redistribution versus its effects on within-lifetime insurance value.

Our findings show substantial self-targeting on both consumption and lifetime income across the eight transfers studied. For example, recipients of the Supplemental Nutrition Assistance Program (SNAP) have an average consumption rank 15 percentiles lower than eligible nonrecipients with similar incomes—a gap equivalent to \$8,000 per person per year, or half of recipients’ average per-capita consumption. These gaps are present in both the PSID and the CEX, although they appear somewhat smaller for other transfers. Self-targeting on consumption mostly reflects differences in lifetime income rather than within-lifetime differences, meaning that take-up largely reveals the lifetime poor rather than people facing temporary drops in consumption. Due to these consumption gaps, our calibrated model suggests that society should be willing to spend up to \$2 in resources to redistribute a marginal \$1 lump-sum from eligible nonrecipients to current recipients within the same transfer program.

Our results relate closely to a body of research on the targeting properties of ordeals, which has thus far reached mixed conclusions.⁴ We depart from this literature in three ways. First, we measure selection for the average transfer recipient rather than the marginal one, focusing on a welfare-relevant question that follows naturally from prior research: If ordeals are costly but not

³Empirical research has consistently confirmed this premise (e.g., [Poterba, 1989, 1991](#); [Cutler and Katz, 1992](#); [Blundell and Preston, 1998](#)), including for low-income households. More recent work highlights income mismeasurement among low-income households as an additional reason to favor consumption as a welfare measure ([Meyer and Sullivan, 2003](#); [Brewer et al., 2017](#)). These reasons often lead researchers (e.g., [Deshpande and Lockwood, 2022](#)) to use consumption as a proxy for marginal utility.

⁴Recent papers studying the targeting properties of ordeals and information include [Bhargava and Manoli \(2015\)](#), [Armour \(2018\)](#), [Ganong and Liebman \(2018\)](#), [Deshpande and Li \(2019\)](#), [Finkelstein and Notowidigdo \(2019\)](#), [Gray \(2019\)](#), [Lieber and Lockwood \(2019\)](#), [Domurat et al. \(2021\)](#), [Unrath \(2024\)](#), [Wu and Meyer \(2024\)](#), [Shepard and Wagner \(2024\)](#), [Ericson et al. \(2023\)](#), [Giannella et al. \(2024\)](#), [Arbogast et al. \(2024\)](#), and [Naik \(2025\)](#).

clearly beneficial, why not shift resources to automatic transfers? Second, we assess whether take-up reveals unobserved need rather than observed characteristics, a key distinction when eligibility rules or taxes can be adjusted alongside ordeals. Prior studies have typically focused on heterogeneity in take-up responses based on observable characteristics like age or income. Third, our approach draws welfare implications from descriptive regressions, avoiding the need for causal estimates from policy changes. Together, these differences lead to a distinct economic conclusion: Take-up appears to achieve desirable self-targeting on average, whereas prior work suggests the marginal impacts are more ambiguous.

The second part of the paper evaluates the welfare implications of self-targeting in transfer programs by asking a central policy question: When society allocates an additional dollar, should it be used to raise the benefit size of a voluntary transfer or to make an automatic payment to all eligible individuals?⁵ More precisely, this “voluntary-or-automatic” question compares a flat marginal increase to a voluntary transfer to providing the same total amount of resources at each income level across all eligible individuals with that income, including current nonrecipients. This approach uniquely isolates self-targeting by enforcing budget neutrality and comparing two versions of the same program. It is also distributionally neutral with respect to current income, stripping out incidental changes in progressivity from the welfare analysis.⁶ Of course, automatic transfers would raise myriad other issues in the real world, so we use this hypothetical reform as a thought exercise to isolate a single fundamental trade-off in transfer design.

Evaluating this “voluntary-or-automatic” comparison requires measuring both the marginal social benefits and costs of take-up as a targeting mechanism. We quantify these forces within a theoretical model of redistribution with nonlinear income taxes and voluntary transfers. First, we prove that the difference in marginal social benefits between voluntary and automatic transfers can be summarized by a self-targeting regression coefficient. Second, we show the difference in marginal social costs is captured by a take-up elasticity with respect to benefit levels. Third, recognizing that the “voluntary-or-automatic” comparison is not generically labor-supply neutral, we derive an additional term to account for the welfare impact of labor-supply responses.

Calibrating our formula yields three principal conclusions. First, self-targeting generates quantitatively significant social benefits: Our baseline calibration estimates gains of approximately 31 cents per transfer dollar, taking a dollar-weighted average across programs. Consistent with our descriptive findings, approximately nine-tenths of these benefits arise from increased progressivity with respect to lifetime income. The remainder reflects better insurance against within-lifetime consumption fluctuations.

⁵We focus on welfarist rationales for transfers. Non-welfarist normative frameworks, such as specific egalitarianism, can also justify transfers—as can externalities, paternalism, or market imperfections.

⁶The comparison can equally be viewed as a marginal shift from a fully-voluntary transfer b to a \$1 automatic transfer with a voluntary “top-up” of $(b - 1)$ multiplied by the voluntary transfer’s take-up rate.

Second, across transfers, the social benefits of self-targeting generally outweigh the associated ordeal costs. Our theoretical framework obtains upper-bound estimates of ordeal costs through the envelope theorem, as when take-up choices are made optimally, the welfare-relevant ordeal cost of marginal recipients equals their benefit level. These upper bounds are large, ranging from 4 to 25 cents per transfer dollar across programs. Models in which incomplete take-up is, at least in part, a result of non-optimizing behavior imply smaller ordeal costs, as the marginal costs paid by non-optimizers must be less than their marginal fiscal costs. Behavioral frictions that reduce take-up would then further weaken the case for automatic redistribution. Our results thus help advance self-targeting as an argument for existing U.S. transfers and against going automatic.

Third, the welfare effects of shifting toward automatic transfers vary considerably across programs. For SNAP and housing assistance, self-targeting is substantial, valued at approximately 39 and 22 cents per transfer dollar—substantially exceeding ordeal costs and rendering automatic transfers socially undesirable. Conversely, several transfers exhibit minimal self-targeting value while still imposing ordeal costs, suggesting potential benefits from automatic provision.

Survey data are subject to important concerns about measurement error in transfer receipt, income, and consumption (Meyer et al., 2009, 2015). We take this challenge seriously, as U.S. administrative data lack consumption and transfer receipt linked across programs. First, on misreporting of receipt, we adopt corrections from a recent literature that estimates how misreporting probabilities vary with observable characteristics (Davern et al., 2019; Mittag, 2019; Meyer et al., 2020). These corrections actually strengthen our results. Second, on consumption misreporting, self-targeting holds for consumption categories thought to be well-measured and for durable goods ownership (Meyer and Sullivan, 2023). Third, on lifetime income, we extend methods in Haider and Solon (2006) to address potential bias from incomplete income histories. Our analysis contains several implicit replications through its analysis of both consumption and lifetime income, take-up in eight transfer programs, and two distinct survey datasets.

We also implement several tests to address measurement error in transfer eligibility. Eligibility imputation is a difficult and pervasive challenge in analyses of U.S. transfers, whether using surveys or administrative data, as both lack eligibility information about nonrecipients. First, our results are robust to reclassifying simulated-ineligible recipients as eligible. Second, we show that our results hold among the poorest subsets of the population that are almost certainly eligible and are most welfare-relevant. Third, results persist after further controlling for any characteristic used in any eligibility rule across our eight transfers. These sensitivity analyses suggest it is unlikely that measurement issues in survey data explain our findings.

Our paper relates most closely to several other analyses of the welfare consequences of selective take-up in social programs. Alatas et al. (2016)'s experiment on transfers in Indonesia and Deshpande and Lockwood (2022)'s study of disability insurance both perform analyses similar to our

main test of self-targeting. Consistent with our results, they find participants in voluntary programs have lower consumption than similar non-participants. Two papers, [Lieber and Lockwood \(2019\)](#) and [Shepard and Wagner \(2024\)](#), develop sufficient-statistics formulas similar to ours, which they respectively use to study ordeal-based targeting in Medicaid home care and markets for subsidized health insurance. Compared to these papers, we study the full breadth of the U.S. safety net and generalize the sufficient-statistics result.⁷ [Finkelstein and Notowidigdo \(2019\)](#) present a welfare analysis of changes in ordeal costs, and we discuss later the key differences between their reform and ours, in which benefit levels change and not ordeals. Finally, in a quite different income-maintenance framework, [Kleven and Kopczuk \(2011\)](#) also argue that take-up costs may be justified by their targeting benefits.

This paper is also connected to two literatures that study redistribution from the perspectives of consumption and lifetime income. The first literature often views the safety net as a form of insurance against consumption fluctuations (e.g., [Blundell and Preston, 1998](#); [Gruber, 2000](#); [Blundell and Pistaferri, 2003](#); [Blundell et al., 2008](#); [Hoynes and Luttmer, 2011](#)). In studying consumption insurance from SNAP, for instance, [Blundell and Pistaferri \(2003\)](#) estimate a consumption-savings model and find that take-up is strongly influenced by “long-run income innovations.” While we do not estimate such a model, our results also highlight the connection to lifetime income, developing its implications from a sufficient-statistics perspective ([Chetty, 2006](#)).⁸

A second relevant literature examines lifetime incidence, or incidence with respect to proxies of lifetime resources ([Fullerton and Lim Rogers, 1993](#); [Liebman, 2002](#); [Bhattacharya and Lakdawalla, 2006](#); [Hoynes and Luttmer, 2011](#); [Blundell et al., 2015](#); [Bengtsson et al., 2016](#); [Roantree and Shaw, 2018](#); [Levell et al., 2021](#); [Auerbach et al., forthcoming](#)). Our results imply that take-up, and not only eligibility rules, influences transfers’ insurance value and lifetime progressivity.

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. We then explain three imputation procedures that augment the PSID data: for cash-equivalent values of in-kind transfers, transfer eligibility, and lifetime income. Appendix B contains additional data details.

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

⁷We study a rich dynamic environment with ability heterogeneity and nonlinear taxation, and we newly incorporate ordeal costs and labor supply responses into the welfare analysis.

⁸[Chetty \(2009\)](#) and [Kleven \(2021\)](#) contrast structural and sufficient-statistic approaches, with specific reference to consumption–savings behavior and discussion of applications of both approaches.

lifetime income, its consumption data, and its information on the receipt of all major U.S. transfer programs. Its major limitations are the reporting issues that we discuss in depth in Section 3. We also replicate the PSID analysis in Consumer Expenditure Survey (CEX) microdata from 1997 to 2019. We prepare the CEX data according to similar procedures as below.

2.1 Income, Consumption, and Transfer Receipt

Current Income. We define household income as the 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 adjust for household size using the equivalence scale $e_h = (N_{h,adult} + 0.7N_{h,child})^{-0.7}$, where $N_{h,adult}$ and $N_{h,child}$ respectively denote the numbers of adults and of children in household h . We compute income ranks within year, pooling across birth-year cohorts.

Current Consumption. The PSID has extensive coverage of consumption expenditures since 1999. Expenditure categories include food, housing, health, transportation, education, child care, and several smaller topics. We adjust the data in two ways to better reflect consumption rather than expenditure, following Meyer and Sullivan (2023). These adjustments convert durable-goods ownership into service flows. 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 for differences in household size. Consumption ranks are also computed within year, implicitly adjusting for changes over time in price levels.

Transfer Receipt. The PSID records self-reported household-level receipt of ten means-tested transfers.⁹ 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); 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.” The CEX covers the first six of these programs and thus omits WIC, LIHEAP, and school meals. Table 1 reports summary statistics.

⁹The measures are for any receipt within the calendar year, though we also observe monthly receipt for three transfers (SNAP, SSI, and TANF). Our analysis excludes contributory social-insurance programs such as unemployment insurance, worker’s compensation, and Social Security. These programs differ fundamentally from noncontributory transfers, and accordingly, we see differences in self-targeting (see Appendix Figure A1). No other transfers are consistently available in the PSID, with the main omissions being tax credits and federal education subsidies.

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	10.9	16.5	5.3	1.0	6.8	4.7	4.2	11.9	28.7	n.a.
Take-Up Rate, Simulated Eligibles	43.4	54.0	10.0	12.3	66.5	40.9	17.0	46.5	47.1	n.a.
Mean Annual Benefit, Recipients	3,428	5,462	5,581	14,201	2,768	572	530	625	n.a.	n.a.
Characteristics of Households or Heads of Recipient Households										
Mean Age, Head	41.6	41.7	40.6	35.5	45.2	33.9	44.6	39.8	42.0	43.1
% Married	19.6	32.7	10.7	15.7	27.3	40.3	26.3	40.7	33.9	48.2
Mean Household Size	3.1	3.4	2.5	3.6	2.6	4.3	3.1	4.2	3.2	2.6
% Children at Home	52.3	60.3	40.7	91.2	29.3	92.2	49.6	94.3	52.0	32.3
% Nonwhite or Hispanic	65.8	63.8	74.4	75.4	62.8	71.8	60.2	69.9	63.9	42.2
% H.S. Graduate	71.3	74.2	74.8	61.4	74.2	72.3	71.6	72.4	76.8	89.0
Mean Household Income	21,350	34,672	21,749	13,879	25,873	39,380	22,923	41,575	38,094	84,885
% Employed	46.2	54.8	50.4	39.9	39.0	69.7	45.7	70.9	59.7	78.9
Mean Rank, Equivalized Households										
Current Income	17.0	22.6	19.7	12.4	19.1	24.4	17.6	25.6	25.5	50.0
Consumption	17.0	22.7	16.8	12.0	27.8	19.3	19.8	22.4	26.3	50.0
Lifetime Income	29.9	34.9	28.3	25.2	32.2	38.3	30.7	37.4	37.2	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 data is from the PSID (waves 1997–2019). Average lifetime ranks are computed as means weighted by life-years of transfer receipt. Monetary values are expressed in 2020 constant dollars.

2.2 Imputation of Other Variables

Cash Equivalents of In-Kind Transfers. We augment PSID/CEX data on the dollar values of transfers with values from the Supplemental Poverty Measure module of the U.S. Current Population Survey (CPS). For SNAP, TANF, SSI, 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 values in-kind transfers other than Medicaid dollar-for-dollar with expenditures.¹⁰

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:

$$E[y_{it} | X_{it}] = \exp(\alpha_i \lambda_a + X'_{it} \beta_a), \quad (1)$$

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, explained in Appendix B, before using the regression results to impute 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 households they experience as adults, without discounting for time. That is, the lifetime household income of individual i is

$$\bar{y}_i^h = \sum_t e(\hat{y}_{it}^h) = \sum_t e(\hat{y}_{it} + \hat{y}_{j(i,t),t}) \quad (2)$$

where t is again summed over the years in which i is between ages 18 and 65, \hat{y} is a predicted income, 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. We compute lifetime-income ranks within birth-year cohorts.

¹⁰For Medicaid, the CPS uses Census estimates of household-level “fungible values” and individual-level “market values.” We use fungible values, so as to remain at the household level. The PSID variables for Medicaid also include state medical-assistance programs.

Simulated Eligibility. Studying transfer take-up among the eligible requires measures of transfer eligibility, so as to distinguish the ineligible from eligible nonrecipients. 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 from survey-data variables. In Appendix B, we show that our simulated-eligibility measure is strongly predictive of transfer receipt, though misclassification is apparent. Considerable fractions of recipients are simulated to be ineligible, and take-up rates are lower than official estimates. Both are routine issues in microsimulations of eligibility.¹¹ Mismeasured eligibility predisposes us to understate the importance of eligibility rules relative to self-targeting among the eligible, and so we consider this threat carefully.

2.3 Preliminary Evidence of Self-Targeting

Before turning to the main analysis, we show evidence of self-targeting in nearly the raw data. We focus on SNAP as an illustrative example. Table 2 reports rates of receipt, simulated eligibility, and take-up among the eligible, doing so jointly by quintile of income and consumption.

Panel A shows that, holding income fixed, low-consumption households are more likely to receive SNAP than high-consumption households. For instance, 35 percent of households in the first quintile of income receive SNAP on average. But 51 percent of households in both the first quintile of income and the first quintile of consumption receive SNAP. Holding income fixed, take-up is thus concentrated among people who have low consumption. This finding differs from the well-known fact that transfer recipients have lower incomes and are generally needier (e.g., Tiehen et al., 2017), in that we show gaps between recipients and non-recipients persist within income. Transfer receipt is thus more targeted on need than its distribution by income might suggest.

¹¹Appendix B provides two supplementary analyses that probe why our estimated take-up rates are lower than figures published by government agencies. First, focusing on SNAP, we decompose the gap into three components: differences in receipt, differences in simulated eligibility, and a “misalignment” component. This third component reflects a little-discussed flaw in official figures that biases them upward: As they simply divide total enrollment by total simulated eligibility, there is no adjustment for recipients in the administrative data who would be simulated ineligible in survey data. We find the underreporting of receipt (see Section 3.4) and the misalignment issue are both quantitatively important to the gap in take-up rates. Second, we review estimates of take-up rates for other programs. Our estimates appear in line with other survey-based estimates but are below administrative-data estimates, consistent with transfer underreporting and misalignment.

Table 2: SNAP Receipt, Eligibility, and Take-Up Rates by Income and Consumption Quintile

<i>Panel A: Receipt Rate</i>							
		Income Quintile					
		1	2	3	4	5	Avg.
Consumption Quintile	1	51.0	21.3	8.9	7.2	7.2	35.3
	2	23.8	9.4	3.0	1.3	0.5	8.5
	3	13.0	5.8	2.1	0.7	0.5	3.4
	4	5.9	4.0	1.1	0.3	0.2	1.3
	5	5.5	2.8	1.8	0.3	0.1	0.9
	Avg.		33.9	11.7	2.9	0.7	0.2
<i>Panel B: Simulated Eligibility Rate</i>							
		Income Quintile					
		1	2	3	4	5	Avg.
Consumption Quintile	1	71.4	13.4	2.1	1.8	2.8	42.7
	2	45.3	8.6	1.8	0.7	0.4	11.7
	3	34.3	8.1	1.6	1.3	0.3	6.2
	4	29.4	8.1	1.3	0.7	0.6	4.0
	5	29.8	9.1	2.3	1.3	0.6	4.0
	Avg.		55.2	10.1	1.7	1.0	0.6
<i>Panel C: Take-Up Rate Among Simulated Eligibles</i>							
		Income Quintile					
		1	2	3	4	5	Avg.
Consumption Quintile	1	58.4	42.8	.	.	.	56.6
	2	34.5	23.6	.	.	.	30.2
	3	23.5	16.7	.	.	.	19.3
	4	13.5	7.6	.	.	.	10.9
	5	14.5	13.8	.	.	.	11.5
	Avg.		48.3	29.2	.	.	.

Notes: This table reports the shares of PSID 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 Table A1 for a tabulation by income and lifetime income, and Appendix Table A2 for a replication in the CEX.

There are two reasons why transfer receipt might be sensitive to both consumption and income. First, transfers may have eligibility criteria that correlate with consumption, even conditional on income, such as asset tests or categorical eligibility for some groups (e.g., people with disabilities). Making many high-consumption households ineligible would directly reduce receipt in that population. Second, such patterns may arise due to “self-targeting,” or rates of transfer take-up among the eligible that depend on their consumption. Self-targeting is of particular interest, as then take-up may reveal private information and thus relax incentive constraints on redistribution.

Panels B and C show the pattern we highlight in Panel A results from both self-targeting and eligibility rules. Take-up among the simulated eligible drops sharply in consumption given income, as much as does the rate of simulated eligibility.¹² These patterns suggest that transfers may indeed reveal something about need that is not already contained in their income. In the rest of the paper, we establish self-targeting more carefully as an empirical fact and assess the welfare consequences of a marginal shift from voluntary toward automatic transfers.

3 Estimates of Self-Targeting in Transfers

This section measures self-targeting in transfers. First, we introduce and implement an empirical framework to study self-targeting. Second, we provide four facts about self-targeting. Third, we explore heterogeneity in self-targeting. Fourth, we assess the sensitivity of our results to measurement issues. In turn, these measurements of self-targeting yield our estimated social benefits of take-up, which we will compare to estimated social costs of ordeals.

3.1 Approach

Motivation and Specification. The ideal measure of self-targeting in a given transfer is the average difference in the social marginal utilities of income (“need”) between the transfer’s recipients and eligible nonrecipients who are otherwise identical to the government. More precisely, the desired quantity is

$$\frac{E[u'_{c,1}] - E[u'_{c,0}]}{E[u'_c]}, \quad (3)$$

where $u'_{c,1}$ and $u'_{c,0}$ respectively indicate the marginal utilities of recipients and eligible nonrecipients. The denominator, average marginal utility, expresses this difference in money-metric terms. That is, the measure in Equation 3 captures a society’s estimated willingness to pay (per dollar) to redistribute resources from eligible non-recipients to recipients, under the assumption that fiscal

¹²While 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.

externalities from this action are offset via lump-sum taxes on all individuals. Section 4 derives this measure of self-targeting formally within our model.

To operationalize this concept of need, we begin with descriptive rank–rank regressions as a first step, followed by theory-consistent regressions. We estimate the following:

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

where \bar{R}_{it} is the consumption rank or lifetime-income rank for household i in year t , R_{it} is i 's current-income rank at that time, $f(R_{it})$ is a flexible function of this rank, and D_{it} indicates i 's receipt status for a given transfer program.¹³ Our baseline specification limits the sample to the simulated-eligible ($E_{it} = 1$), thus dropping the simulated-ineligible. The coefficient β is therefore the average difference in consumption rank or lifetime-income rank between transfer recipients and eligible nonrecipients with similar incomes. When we turn from descriptive to normative analysis, we replace the rank outcome \bar{R}_{it} with calibrated values of marginal utility.

Discussion. Translating the concept of marginal utility differences in Equation 3 into an empirical framework like Equation 4 involves several key choices. Many of these choices have parallels in the literature on measuring poverty and material hardship (e.g., Citro and Michael, 1995). Here we discuss four worthy of focus: the choice of consumption as a proxy for need, the use of rank-based measures, the control for current income, and the focus on simulated eligibility. Section 3.4 further explores the robustness of our results to these choices.

First, marginal utilities of income are unobserved, requiring a model to link the data to marginal utility. Here, we take consumption as a proxy for marginal utility, as others have commonly assumed (e.g., Finkelstein et al., 2019; Deshpande and Lockwood, 2022). Two broad motivations for this choice are the partial insurance of household consumption against income shocks and the mismeasurement of low incomes in both survey and administrative data. This choice aligns with consumption–savings models where marginal utility is monotonic in consumption.¹⁴ Under this monotonicity, consumption ranks align with marginal-utility ranks in a cross-section of households. Lifetime-income ranks would also align in this sense if consumption is fully insured against income shocks. We therefore describe transfers as “advantageously self-targeted” if take-up among eligible households predicts lower ranks in the consumption or lifetime-income distributions.

Nonetheless, imputing marginal utility from consumption has inherent limitations. Perhaps the most significant is that it is infeasible to fully capture all variation in preferences when performing this imputation. For instance, factors like aging, parenthood (Blundell and Preston, 1998), opportunity costs of time (Aguiar and Hurst, 2005), and habit formation (Chetty and Szeidl, 2016)

¹³We parameterize $f(R_{it})$ using cubic splines with knots at the 10th, 25th, and 50th percentiles of income.

¹⁴This includes a special case of our model. For further discussion, see Deaton (1992).

all influence the relationship between consumption and marginal utility. Heterogeneity in labor-supply elasticities or nonseparability of labor and consumption create similar difficulties. A second set of challenges arises from measurement issues in observed consumption, including systematic underreporting of certain consumption types (Meyer and Sullivan, 2023).

Second, we focus on rank-transformed outcomes to provide a clear, interpretable summary of self-targeting. The common use of rank-transformed outcomes is motivated by their robustness to outliers and scale differences.¹⁵ In our context, one drawback is that regressions on rank-transformed outcomes implicitly weigh the outcome distributions differently than regressions in levels or marginal utilities. We later show similar results using both of those measures. Reweighting the data would also address this issue. By implication, little is lost in moving from descriptive facts, presented in ranks, to welfare implications that are not in rank form.

Third, the inclusion of flexible controls for current income reflects the theoretical perspective that transfers are designed to complement, not duplicate, income taxation. Of course, real-world income tax schedules depend on variables other than income, most notably aspects of household composition (e.g., household size and marital status). Further controls can be warranted to the extent the government would adjust eligibility rules in transfer reforms.

Fourth, the ideal estimate of self-targeting compares recipients with eligible nonrecipients, but we only have imperfect measures of simulated eligibility. By implication, our sample will inevitably exclude some true-eligibles and include some true-ineligibles. Measurement error in eligibility is, moreover, unlikely to be independent of consumption. The exclusion of some eligibles and the inclusion of some ineligibles are thus potential sources of bias.

Connection to Prior Research. Our investigation of self-targeting differs significantly from prior literature reviewed in Section 1. This distinction emerges because we consider different hypothetical reforms, resulting in different objects of interest.

Much of the prior literature evaluates the welfare consequences of changes in ordeal costs. As an empirical matter, it has measured the characteristics of people whose take-up choices respond to policy changes. The theoretical relevance of these marginal recipients is that heterogeneity in their take-up responses reveals the selection effects of ordeals on need at the margin.

By contrast, we study a budgetary shift from voluntary to automatic transfers. For a sufficiently small change, a budget reform redistributes exclusively between a voluntary transfer's inframarginal recipients and its eligible inframarginal nonrecipients. Marginal recipients are exactly indifferent to this reform, since their benefits and costs of take-up are equal, and the characteristics of indifferent people are immaterial for welfare. Instead, the selection that matters for welfare is between

¹⁵Our approach is also related to a storied literature on consumption-function estimation (Friedman, 1957). With consumption as an outcome, our self-targeting coefficients can be interpreted as differences in permanent income between recipients and eligible nonrecipients.

inframarginal recipients and nonrecipients, as in Equation 4.

Both ordeal changes and benefit shifts can offer valuable insights. Governments routinely consider changes in benefit levels as well as changes to eligibility and ordeals. Considering a larger space of transfer reforms has two other potential benefits. First, the literature finds mixed results on targeting under the first approach (e.g., [Deshpande and Li \(2019\)](#) versus [Finkelstein and Notowidigdo \(2019\)](#)). This uncertainty may reflect heterogeneity in marginal households across settings and ordeals. They also give reason to examine other policy hypotheticals, like our “voluntary-to-automatic” reform, which might (and indeed does) offer a more-robust empirical conclusion. Second, our analysis requires only observational data, since it is fundamentally descriptive, whereas identifying marginal recipients requires exogenous variation in ordeals.

3.2 Four Facts About Self-Targeting

Fact 1: Self-targeting is advantageous with respect to consumption. In the typical transfer program, recipients rank around 10 percentiles lower in the consumption distribution than eligible nonrecipients with similar current incomes (see the yellow diamonds in Panel A). These consumption-given-income rank differences are so large that they often exceed the raw differences in income ranks between recipients and simulated-eligible nonrecipients (Appendix Figure A2). In level terms, these results are consistent with differences of approximately 20 to 60 percent in these outcomes, or around \$2,500 to \$10,000 per person per year in consumption (see Appendix Table A3). While consistent in sign, self-targeting does appear to vary in magnitude across programs, from zero in SSI to 15 percentiles in SNAP. This heterogeneity is our focus in Section 3.3.

Fact 2: Self-targeting among the eligible often appears more important than eligibility rules. In Panels A and B, we also show differences in consumption between recipients and *all* similar-income nonrecipients. That is, results in blue circles include simulated-ineligible people in the estimation sample. These results therefore combine two distinct sources of selection, eligibility rules and self-targeting by the eligible.

Overall, we find somewhat more selection in the pooled sample than when we restrict the sample to the simulated-eligible. This pattern suggests that both eligibility rules and self-targeting tend to select advantageously on consumption and lifetime income. Across transfers, self-targeting is often the primary force behind selection, and eligibility rules are generally secondary. This interpretation requires eligibility to be well-measured, a topic to which we return below.

Fact 3: Adverse shocks within life, as well as low lifetime income, induce transfer take-up. Panels B and C study the extent to which self-targeting on consumption is on lifetime income (and thus between lifetimes) or is on consumption given lifetime income (thus within lifetimes). In Panel B, we find roughly half as much self-targeting on lifetime income as we found on consumption in

Panel A. That is, some of what take-up reveals is who among the eligible is lifetime-poor and thus has little past income or future income to smooth during a period of transfer eligibility.

Panel C reaches a similar conclusion in a different way. Here we augment the self-targeting regression (Equation 4) with person fixed effects and restrict to the simulated-eligible. For comparison, the blue series in Panel C repeats results from Panel A (originally in yellow) which do not include fixed effects and thus pool between- and within-person self-targeting. We continue to find self-targeting on consumption: Within-lifetime consumption drops correlate with changes in take-up, holding current income fixed. Importantly, these results continue to condition on current income, so they go beyond well-documented empirical patterns of take-up around observable shocks (e.g., [Wu and Zhang, 2025](#)). These empirical patterns lead to our conclusion in Section 4 that voluntary transfers have some insurance value.

We further show the relevance of permanent income to self-targeting by plotting consumption functions as in [Friedman \(1957\)](#) for recipients, eligible nonrecipients, and ineligible (see Appendix Figure A3). Group differences in this permanent-income concept, or the incomes at which the average savings rate is zero, are large and accord with the rank differences shown in Figure 1.

Fact 4: Marginal utilities of income appear higher among transfer recipients than among eligible nonrecipients with similar incomes. Panel D connects our descriptive self-targeting regressions to the welfare analysis, and specifically to the welfare benefits of self-targeting. We do so by replacing the household-rank dependent variable with a calibrated value for each household’s marginal utility of income. The resulting coefficients can be interpreted as the additional value of moving a marginal dollar from a transfer’s eligible nonrecipients to its recipients.

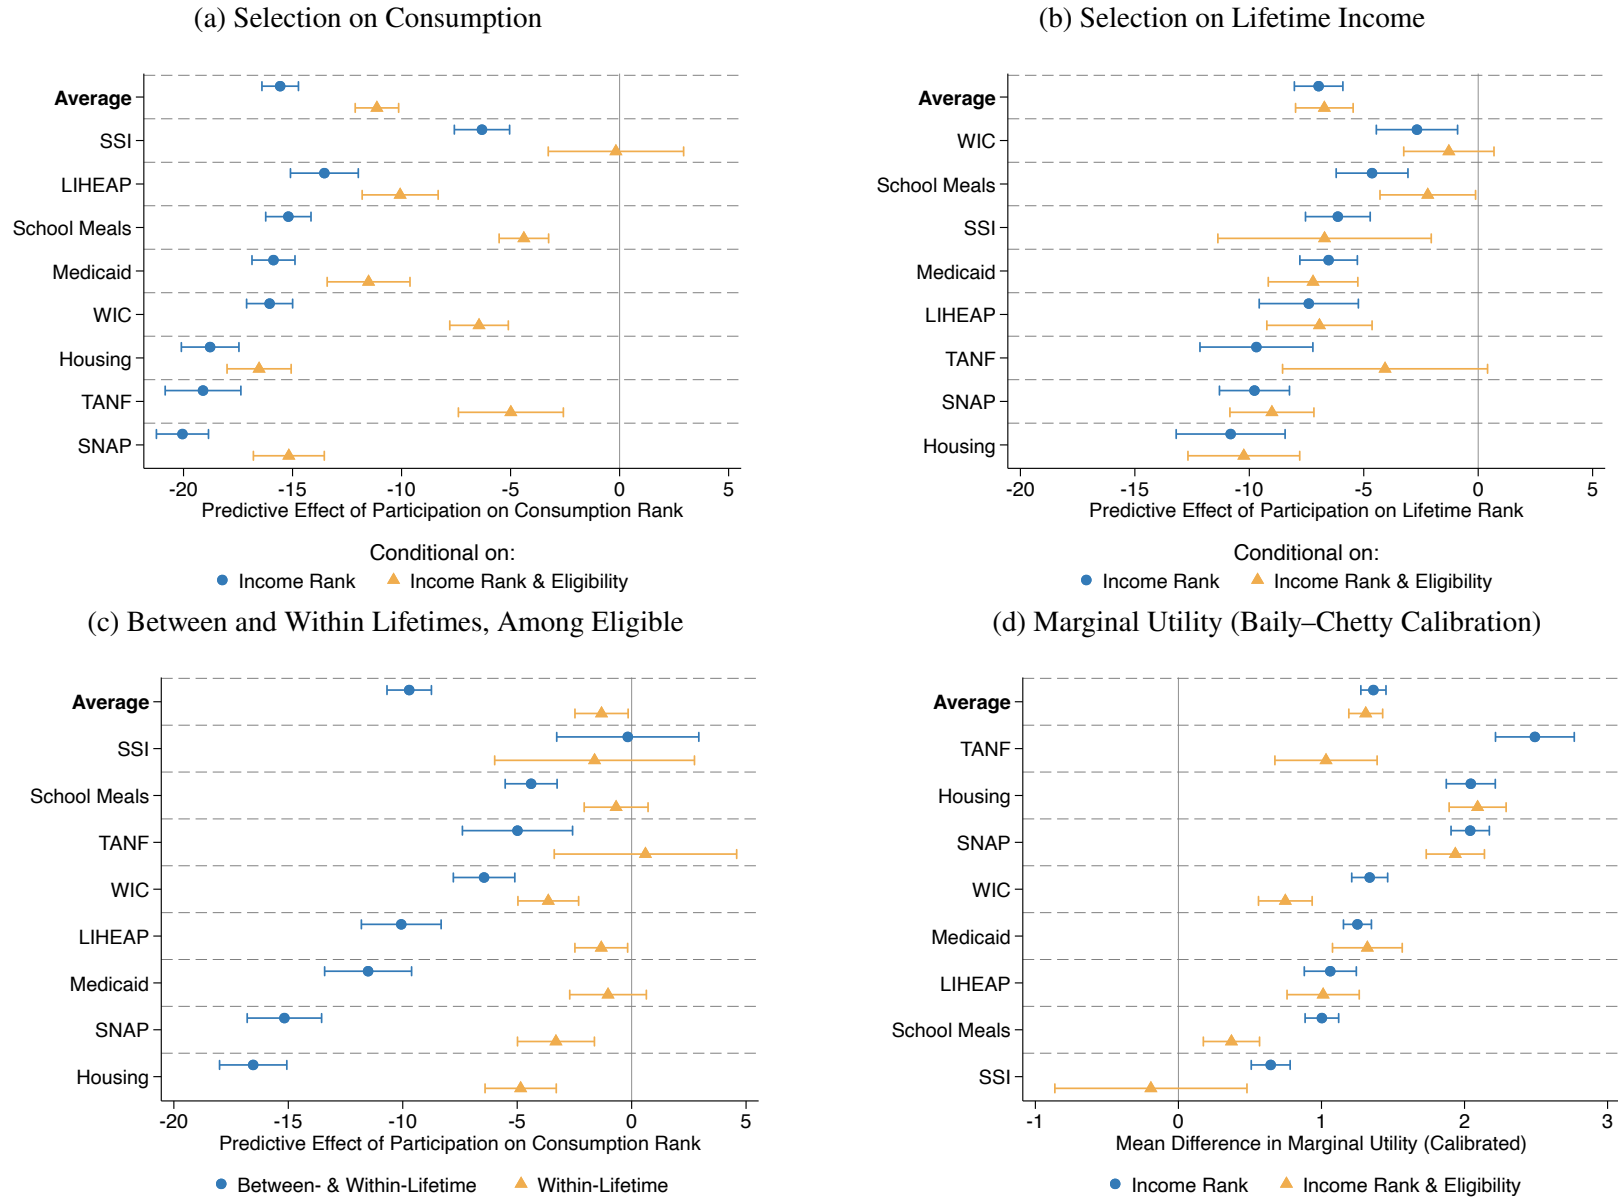
To calibrate marginal utilities for each household, we use the following preferences over consumption and labor hours from [Greenwood et al. \(1988\)](#):

$$u(c, \ell; \psi_i) = \frac{1}{1 - \gamma} \left(c - \frac{\psi_i}{1 + 1/\eta} \ell^{1+1/\eta} \right)^{1-\gamma}, \quad (5)$$

where c is the household’s real annual equivalized consumption and ℓ is the annual labor hours of the household head or spouse. We will use the same preferences in our welfare analysis (Section 4) and discuss them further there.¹⁶ We express our results as a money-metric by rescaling each marginal utility using the population-average marginal utility.

¹⁶We calibrate $\gamma = 3$ (capturing both individual risk aversion and social welfare weights) and $\eta = 0.3$ as constants across households. The disutility parameter is calibrated internally as $\psi_{it} = w_{it}/\ell_{it}^{1/\eta}$, which we obtain by inverting the household’s labor supply function $\ell_{it}^*(w_{it}, \psi_{it})$. Following [Finkelstein et al. \(2019\)](#), we also impose a “consumption floor” at the fifth percentile in calculating marginal utility u'_c .

Figure 1: Self-Targeting in Transfer Programs



16

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 β from Equation 4). For the yellow diamonds, we estimate the regression only on people whom we simulate to be eligible. Panel C augments the specification with person-level fixed effects. The “average” row of Panels A–D stacks the program-level regressions, weighting each program by total amount transferred. Panel D adapts Equation 4 by replacing the ranks outcome with household-level calibrated levels of marginal utility from our theoretical model, using consumption, wages, and hours. In all panels, 95-percent confidence intervals reflect clustered standard errors by household.

This regression yields coefficients which are exactly the object in Equation 3, that is, the left-hand side of the Baily–Chetty equation for the optimal level of social insurance (Baily, 1978; Chetty, 2006). We thus capture the social benefits of self-targeting, which Section 4 compares to calibrated estimates of the social costs of take-up. Panel D shows the estimates with and without a sample limitation to the simulated-eligible.

Panel D shows that, in most transfers, the marginal dollar of transfer to recipients is worth substantially more than the marginal dollar to eligible nonrecipients. For instance, society would be willing to pay \$1.94 in resources to make a nondistortionary transfer of \$1 from the average SNAP-eligible nonrecipient to the average SNAP recipient. We estimate this Baily–Chetty premium to be \$1.31 per transfer dollar on average over the eight transfers. The premium arises for two reasons. First, calibrated estimates of marginal utility are high for the bottom of the consumption distribution. Second, as shown in the other three panels, transfers tend to attract such households, even conditional on income and eligibility. Indeed, in transfers where we found no advantageous self-targeting on consumption and lifetime income, we do not find such a premium.

3.3 Explanations and Extensions

Does take-up reveal exogenous earnings ability, or only a possibly endogenous component of “need”? How does self-targeting vary across demographic groups? What might explain the across-program variation in self-targeting? And how does self-targeting relate to economic models of consumption behavior? We explore these questions here.

Consumption as a Measure of Need. We have so far described our results in ways that imply take-up reveals some exogenous concept of need. Behavioral responses to taxes and transfers, however, suggest that household need is not fully exogenous. The responses would affect the social costs but not the social benefits of self-targeting in transfers. Three pieces of evidence suggest take-up does indeed reveal information about households’ exogenous resources.

First, we find self-targeting in the most plausibly exogenous correlates of need (see Appendix Table A4). In most programs, take-up among the eligible is strongly associated with high-school dropout, single parenthood, disability, minority race or ethnicity, and lower rates of savings. As these groups are high-need, these patterns also provide support for the assumption that low consumption reveals marginal utility.

Second, we show that consumption patterns in the years around take-up are consistent with self-targeting on persistent earnings ability, rather than behavioral responses to transfers (see Appendix Figure A4). Among *current* eligible nonrecipients of a given transfer, *future* recipients have on average a lower current-consumption rank than similar *future* eligible nonrecipients.¹⁷ We find

¹⁷The regression of interest is $\bar{R}_{it} = \alpha_{ct} + \beta D_{i,t+k} + f(R_{it}) + u_{it}$ within the subsample such that $D_{it} = 0$, for $k > 0$.

strong self-targeting even for take-up in the very distant future (10–20 years), ruling out potentially strategic changes in consumption just before transfer take-up.

Third, we show that the behavioral responses in consumption required to fully explain our results are very large (see Appendix Figure A5). This addresses the concern that recipients may adjust reported or actual consumption in response to transfers. For most transfers, the response must exceed the dollar value of transfers received. In addition, our focus on the targeting of the “last dollar” of transfers is conservative. Receiving transfers itself enables households to raise consumption and thus lowers imputed marginal utility, and so the “first-dollar” targeting of transfers would thus be even better than in our main results (Deshpande and Lockwood, 2022).

We also considered several reasons why preferences may lack a monotonic relationship between consumption and marginal utility: lifecycle dynamics, household composition, nonseparability in labor, and health state-dependence. While these forces are undoubtedly present among some households, they appear unable to reverse our main finding that transfer recipients have lower marginal utilities than similar eligible non-recipients.

Reviewing these checks, first we find similar self-targeting with and without equivalence-scale adjustments (Appendix Figure A6) and when we control flexibly for household composition (Appendix Figure A7). By implication, our results are insensitive to the equivalence scale or relationship of preferences to household composition. Second, our self-targeting estimates also change little when we adjust consumption for regional variation in price levels (Appendix Figure A8). Third, we show in two ways that lifecycle shifts in preferences are unlikely to overturn self-targeting. We allow for cohort-specific slopes of the income control (Appendix Figure A9), and we estimate self-targeting separately by age group (Appendix Figure A10). Finally, we show the robustness of self-targeting to utility functions with state-dependence in health and or with nonseparability of consumption and labor (Appendix Figure A11). We find self-targeting across levels of health and work hours.

Demographic Heterogeneity. People may have greater or lesser tendencies to self-target. Even if take-up predicts lower average consumption and lifetime, this relationship could be weaker or inverted in some demographic groups. Such heterogeneity is of interest for multiple reasons. For instance, some readers might wish to apply social welfare functions other than ones we consider. Alternatively, heterogeneity might motivate policy accommodations, such as social workers and native-language forms, to adjust ordeals in a targeted manner. Finally, heterogeneity could speak to potential mechanisms by which self-targeting arises.

We first extend our analysis by dividing the population by education, race, and ethnicity (Appendix Figures A12 and A13). We also break out three disadvantaged populations: single parents, non-native English speakers, and car non-owners. Examining SNAP and Medicaid, we find some suggestive heterogeneity in self-targeting. In general, we see more self-targeting among

better-off groups (e.g., college graduates) than worse-off groups (e.g., single parents). No single demographic category, however, drives our results or stands out strongly from this broad pattern.

We also examine heterogeneity according to the number of distinct transfers received (see Appendix Figure A14). In both the PSID and the CEX, eligible households that receive multiple transfers are more self-targeted on consumption and lifetime income than those receiving only one transfer. This heterogeneity might rationalize rules that make recipients of one transfer categorically eligible for other transfers, such as “adjunctive eligibility” in WIC.

Finally, we look more granularly at where self-targeting emerges across the distributions of consumption, lifetime income, and current income in the PSID and CEX. Households with the lowest consumption and lifetime income drive our self-targeting estimates (see Appendix Figure A15). Self-targeting on consumption also appears slightly stronger at the bottom of the income distribution than elsewhere (see Appendix Figures A16 and A17). Taken together, these results confirm self-targeting is present even among the neediest households that are most influential in social-welfare calculations.

Across Transfers. We provide a correlational analysis of variation in self-targeting across transfers. In particular, we correlate our estimates with the following program characteristics: the eligible share of the population, and the mean consumption and income ranks of the eligible. There is some suggestive evidence for the view that ordeals and eligibility rules act like substitutes: Programs with tighter eligibility rules have weaker self-targeting (see Appendix Figure A18). However, this pattern is not causal and may reflect other program attributes.¹⁸

One further challenge in studying across-transfer heterogeneity in self-targeting is that it could arise spuriously from differences in eligibility (Appendix Figure A19).¹⁹ To test this explanation, we build a “stacked” dataset with one observation for each person and each transfer for which that person is eligible. We then estimate models of take-up with and without person-level fixed effects.²⁰ These specifications yield similar results, implying across-transfer heterogeneity is not spurious (i.e., is genuinely within-person).

Microfoundations of Self-Targeting. Theories of consumption behavior suggest several potential microfoundations for self-targeting. In particular, take-up might be driven by differences in expected future consumption, future-consumption risk, borrowing constraints, and tastes. Here we explore

¹⁸We also do not find self-targeting in “contributory” social insurance programs like unemployment insurance, worker’s compensation, and Social Security (see Appendix Figure A1). Self-targeting on need would not be expected in these programs, as benefit levels are functions of past earnings.

¹⁹By “spurious,” we mean the following scenario. When eligible, every person could be equally likely to take up Transfer A as Transfer B, but they vary in their probability of taking up any transfer. If different populations were eligible for A and for B, we may find different intensities of self-targeting in A and B.

²⁰The person-effects specification is identified only by people who are simultaneously eligible for multiple transfers, analogous to movers identifying firm or place effects.

which of these appears important, in the hope of offering some directions for future research.

Differences in two-years-ahead future consumption account for much of self-targeting (see Appendix Figure A20). That is, among households with similar current income *and* future consumption, take-up becomes much less informative of current consumption. Beyond future consumption, additional controls seem to matter little, including the household’s stock of liquid assets or the occupation and industry of household members. These findings provide little indication that borrowing constraints or precautionary-savings motives are important reasons for take-up. Similar to Blundell et al. (2008), we conclude that a simple life-cycle model with permanent and transitory components of income is likely to generate realistic self-targeting on consumption.

3.4 Sensitivity to Mismeasurement

Survey data are imperfect. Here we consider the potential for bias in our results due to measurement error in simulated eligibility, self-reported transfer receipt, current income, consumption, and lifetime income. Overall, measurement issues seem unlikely to explain apparent self-targeting, given its magnitude and robustness.


We are especially careful with the survey data because our analysis is mostly infeasible in U.S. administrative data. First, such datasets lack appropriate measures of consumption. Second, they rarely link information across transfer programs. Third, while administrative data would improve measurement for some inputs to the analysis (e.g., income and transfer receipt), it would be harder to impute eligibility. Administrative data largely does not record eligibility information for nonrecipients, lacks the detailed covariates of survey data useful for imputing it, and may not capture some eligible nonrecipients who may appear in surveys (e.g., income-tax nonfilers).

Data Comparison. We estimate self-targeting on consumption for the five transfers that appear in both the PSID and CEX. Overall, the CEX shows a similar degree of self-targeting as the PSID (see Figure 2). However, we find less heterogeneity across programs in the CEX than in the PSID. These results suggest that data limitations that are specific to the PSID, such as panel attrition or less-thorough survey questions on consumption, cannot explain our results.

Eligibility. We mostly analyze self-targeting among people who are simulated-eligible for transfers, a group differs from the truly-eligible through both errors of inclusion and exclusion. The two-sidedness of the sample selection issue generates a bias of ambiguous sign.²¹ We assess this

²¹Appendix C derives the bias of our estimator. We express the bias in three terms. First, self-targeting could be weaker or stronger among the simulated-eligible population relative to the simulated-ineligible. Second, unobserved variation in true eligibility among the simulated-eligible might correlate with consumption. Third, there may be residual informativeness of simulated eligibility for consumption, to the extent we do not control flexibly for eligibility-rule variables. We expect the first term leads us to understate self-targeting, and the second and third terms to overstate self-targeting, leaving the overall bias ambiguous.

Figure 2: Estimates of Self-Targeting by Data Source (PSID Versus CEX)



Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank, in both the Panel Study of Income Dynamics (PSID, blue circles) and the Consumer Expenditure Survey (CEX, yellow diamonds). We estimate all regressions on the simulated-eligible population. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

bias through several robustness checks. Adjustments to our measures of simulated eligibility can shrink our estimates of self-targeting, but they leave standing our conclusion that advantageous self-targeting is a key force in most transfers.

First, simple adjustments to the eligibility simulations, such as imposing income limits or liquid-asset tests, have modest impacts on our estimates of self-targeting (Appendix Figures A21 and A22). These results suggest that further improvements to the eligibility simulations are unlikely to change our estimates materially. Furthermore, when we reclassify all simulated-ineligible recipients of a given program as eligible, we still find that self-targeting concentrates incidence among the consumption-poor and lifetime-poor (Appendix Figure A23). This reclassification test is quite demanding, as by construction it “overfits” eligibility rules to explain receipt.

Second, when we restrict the sample to demographic groups with very high rates of simulated eligibility, we continue to find self-targeting (Appendix Figures A24 and A25). Among such groups, patterns of transfer receipt are less likely to be confounded by eligibility mismeasurement and thus more credibly isolate self-targeting.

Third, we are mostly limited to annual data, despite the possibility of correlated within-year fluctuations in receipt, eligibility, income, and consumption. For three transfers, however, we can

show that measuring receipt in terms of the share of months in the year, or comparing only zero-month to twelve-month households, matters little for our PSID results (Appendix Figure A26). In addition, monthly data on food insecurity and SNAP take-up allow us to test directly for aggregation bias from annual data (Appendix Table A5). We can also address income changes directly: We find similar self-targeting when we look only at households without any change in employer in the last two years, including the non-employed (Appendix Figure A27). Such households are likelier to have stable incomes, reducing one source of potential aggregation bias from annual data.

Fourth, we test measurement error by examining how the self-targeting coefficient moves when we add additional controls to proxy for unobserved eligibility rules. Following Oster (2019), coefficient stability suggests mismeasurement in eligibility rules is unlikely to reverse our conclusions. We first augment Equation 4 with controls for any variable that enters into the eligibility simulation for any transfer we study.²² This attenuates but does not eliminate selection into receipt for most transfers (Appendix Figure A28). We then also include controls for variables that do not enter any eligibility simulation but predict both consumption or lifetime income and transfer receipt: race, education, and marital status. Adding in these further controls does little to move our estimates. This test builds confidence in our results, as it seems unlikely that any unobserved variable would be much more important than these “unused observables.”

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 variables.

Their corrections consistently increase our estimates of advantageous self-targeting (Appendix Table A6). There are two reasons why. First, under constant misreporting rates, our estimates are attenuated. Consider misreporting probabilities $p_0 = \Pr(\tilde{D}_i = 0 | D_i = 1)$ and $p_1 = \Pr(\tilde{D}_i = 1 | D_i = 0)$, where D_i indicates true receipt and \tilde{D}_i indicates reported receipt. Comparing the feasible regression of $y_i = \tilde{\beta}\tilde{D}_i + u_i$ to the infeasible regression $y_i = \beta D_i + u_i$, one can show that $\beta = \tilde{\beta}/(1 - p_0 - p_1)$.²³ Second, the estimates in Mittag (2019) and Davern et al. (2019) both imply that underreporting of transfers is more common among households with low consumption and lifetime income, holding income constant. Thus, their adjustments amplify the increase in the

²²These are the household’s state of residence by year, 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).

²³In the PSID, Meyer et al. (2009) find rates of net under-reporting (false negative minus false positive) of 7 to 27 percent across programs we study. This suggests a presumption that our estimates in Figure 1 are understated, even for transfer programs where heterogeneous-misreporting corrections have not yet been estimated.

selection that we would find under constant misreporting rates.

We also compute the rates of transfer 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.

Income. Income is poorly measured at the bottom of the distribution. While there is also mismeasurement in consumption, it appears to be less severe than in income (Meyer and Sullivan, 2003; Brewer et al., 2017). This point threatens our analysis: To the extent we under-control for income due to attenuation bias, our estimates of self-targeting may be overstated.²⁴

We address this threat in two ways. We first supplement our control for income rank with a “predicted” income rank derived from other labor variables, such as weekly hours and occupation.²⁵ This strategy lessens the threat of misreporting if errors are not perfectly correlated across labor variables. We find this additional control for predicted income has a modest impact on our results (see Appendix Figure A29), consistent with limited bias from misreporting. While this method cannot completely address measurement error, it mitigates the concern by effectively consolidating many labor variables into additional controls.

Next, we directly probe the sensitivity of self-targeting to any bias transmitted from income mismeasurement. In a linear version of Equation 4,²⁶ we constrain the coefficient on income rank to higher values than the unconstrained estimate. We claim there exists a threshold value for the constrained coefficient that is implausible a priori, as it would imply an excessive noise share of income variance or unrealistically little smoothing of consumption relative to income.

Self-targeting in most transfers appears robust to substantial mismeasurement of income (see Appendix Tables A7, A8, and A9). In the PSID, when we constrain the rank–rank slope between consumption and income to 0.7, up from 0.51 when unconstrained, we find that SNAP recipients are on average 13.6 percentiles lower in the consumption distribution than eligible nonrecipients with the same income. This amounts to a downward adjustment of only two percentiles from our baseline estimate, despite allowing for substantial attenuation bias in the rank–rank slope.

Consumption. Meyer and Sullivan (2023) describe consumption as comparatively “well-measured”

²⁴The underreporting of income on surveys would affect our analysis only insofar as transfer recipients and nonrecipients have different underreporting rates. In fact, one plausible story of income misreporting runs counter to our results. In this story, transfer recipients have incentives to underreport income to maintain eligibility. They may do so in any quasi-official setting, including surveys. These incentives might apply less strongly to consumption and to nonrecipients. All else equal, transfer recipients would thus appear *positively* selected on consumption given income.

²⁵Using the March Supplements to the Current Population Survey that match our PSID data years, we estimate Poisson regression models of individual income based on occupation, industry, weeks worked per year, weekly hours, self-employment, and basic demographic information. We then apply these predicted incomes to our PSID data.

²⁶ $\bar{R}_i = \beta D_i + \gamma R_i + u_i$, where \bar{R}_i is consumption rank or lifetime rank, R_i is income rank, and D_i is transfer receipt.

for some goods and services. Appendix Tables A10 and A11 document self-targeting in the PSID and CEX when we look exclusively at these consumption categories. For instance, in the PSID, Medicaid recipients consume 39 percent less in housing, 33 percent less in vehicles, and 33 percent less in food at home than similar-income eligibles who are not on Medicaid.

Similar patterns manifest in PSID and CEX measures of household ownership of consumer durable goods. Following Meyer and Sullivan (2012), we consider whether the household owns a home, car, and computer, as well as the number of rooms and presence of central air conditioning in the home (see Appendix Tables A12 and A13). In the PSID, SNAP recipients are 13 percentage points less likely to own a home, 9 percentage points less likely to own a car, and 9 percentage points less likely to own a computer than similar-income eligibles not on SNAP. These proxies for consumption seem highly unlikely to suffer from meaningful misreporting.

Lifetime Income. Inferring lifetime income from “snapshots” is known to be difficult (Haider and Solon, 2006). We employ a new, transparent check against this issue. Lifetime-income ranks for households with many years-in-sample are mostly data, whereas those with few years-in-sample are mostly imputation. If the imputation process has imparted a systematic bias in our results, then the estimated extent of selection into transfers would be correlated with a household’s number of 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 A30 shows that selection on lifetime income is essentially constant across households when split by their years-in-sample. Moreover, the minor “drift” effects sometimes appear for consumption, which is not imputed in this way. The phenomenon may reflect considerations other than a bias in lifetime-income estimation, such as sample attrition.²⁷

4 Welfare Analysis

This section conducts a welfare analysis of voluntary take-up as a targeting device for transfers. For this analysis, we derive a new theoretical result that contrasts the welfare costs and benefits of voluntary and automatic transfers. We calibrate this cost–benefit formula using our data and external estimates, and we discuss its implications for social policy.²⁸

²⁷Attrition from the PSID is, of course, also a source of concern. Yet it is not obvious that households that stay in the sample longer are more representative than those who attrit more quickly.

²⁸Appendix C reports estimates of the Marginal Value of Public Funds (MVPF) for voluntary and automatic transfers that are consistent with our model. Appendix D contains proofs of all theoretical results.

4.1 Basic Environment

To make the economics of our theoretical result especially clear, we first illustrate it in a basic environment without labor supply or risk aversion. Suppose there is a unit mass of people indexed by $i \in [0, 1]$. Of these, a share $M(s)$ choose to take up a voluntary transfer, receiving benefit s and paying a hassle cost $\kappa(i)$ to do so. The complementary share $1 - M(s)$ do not take up because their ordeal cost strictly exceeds the benefit, $\kappa(i) > s$.

Let us consider the welfare impacts of the following two reforms. In the first, we increase the voluntary transfer by an amount ds . Raising the benefit causes a share $\frac{M(s)}{s}\varepsilon_b ds$ to take up the transfer, where the take-up elasticity with respect to the benefit amount is $\varepsilon_b = d \log M(s)/d \log s$. The fiscal externality from this behavioral response is $\varepsilon_b M(s) ds$. In the second reform, we give all people $(1 + \varepsilon_b)M(s) ds$ dollars automatically, thus spending exactly the same amount as in the first reform. We want to determine which reform is better for welfare.

Let $\alpha(i)$ be the marginal welfare weight of household i , setting the population average social value to unity without loss of generality. The difference in the welfare effects of the reforms is

$$\begin{aligned} \frac{dW}{ds} = & \underbrace{M(s)E[\alpha(i) \mid \kappa(i) \leq s]}_{\text{recipients}} \cdot ((1 + \varepsilon_b)M(s) - 1) \\ & + \underbrace{(1 - M(s))E[\alpha(i) \mid \kappa(i) > s]}_{\text{nonrecipients}} \cdot (1 + \varepsilon_b)M(s), \end{aligned} \quad (6)$$

which we compute as automatic minus voluntary. This expression follows from weighting the net changes in payments to recipients and nonrecipients by these groups' respective population shares and welfare weights. Ordeal costs do not appear, as they only change for a marginal type who was exactly indifferent to taking up. Equation 6 simplifies to

$$\frac{dW}{ds} = -\beta\sigma_M^2(s) + M(s)\varepsilon_b, \quad (7)$$

where the variance of take-up is $\sigma_M^2(s) = M(s)(1 - M(s))$. The coefficient β equals $E[\alpha(i) \mid \kappa(i) \leq s] - E[\alpha(i) \mid \kappa(i) > s]$, which is the difference between the marginal social values of income of the recipients and nonrecipients.²⁹

Equation 7 captures our first theoretical result. The first term is the social benefit of self-targeting from the marginal dollar of voluntary transfer. If take-up identifies higher-welfare-weight people on average, then a voluntary transfer has desirable redistributive properties. The targeting is more socially valuable when take-up is more informative about welfare weights $\alpha(i)$ and thus

²⁹Our welfare formula thus belongs to a ‘‘Baily–Chetty’’ class of sufficient-statistics formulas that contrast a targeting gain (in the form of a difference in social marginal utilities) with a behavioral response (in the form of an elasticity).

when β is larger. When all or none take up the transfer, targeting is impossible, explaining why the reform's benefits depend on the variance of the take-up rate.

The second term represents the social costs of ordeals. In response to the marginal dollar, some people take up, since the benefit now just exceeds their ordeal cost. We apply the envelope theorem to infer these costs. With privately optimal take-up choices, the fiscal savings from people who no longer take up equals the change in the social cost of ordeals.

Example. Suppose there is a voluntary transfer with a 50-percent take-up rate. Imagine that on average, recipients consume 57 percent as much as non-recipients, and the take-up elasticity is 0.5. Society has constant-relative-risk-aversion (CRRA) preferences over consumption with parameter $\gamma = 2$. On the margin, should funds go to a voluntary transfer or an automatic one?

The answer is voluntary. Each dollar taken from the recipients is worth about three times ($1/0.57^2 \approx 3$) that of a dollar given to nonrecipients. Normalizing the average welfare weight to one, given 50-percent take-up, yields a regression coefficient $\beta = (3/(3 \times 0.5 + 1 \times 0.5)) - (1/(3 \times 0.5 + 1 \times 0.5)) = 1$. That is, society should accept an efficiency cost of up to one dollar in order to redistribute one dollar of resources from nonrecipients to recipients. The variance of take-up is $\sigma_M^2 = (0.5)(1 - 0.5) = 0.25$. Then $dW/ds = -(1)(0.25) + (0.5)(0.4) = -0.05$, so shifting the marginal dollar ds toward the automatic transfer and away from the automatic one reduces social welfare by approximately five cents.

4.2 Setup of Full Model

We next propose a dynamic model of optimal redistribution, which we use below to derive our theoretical result. The dynamics allow for a concept of consumption that is meaningfully distinct from current income. The model also incorporates risk aversion and endogenous labor supply, and thus insurance and efficiency implications of potential transfer reforms.

Households. Time t is discrete, and households live forever. Each period, a household has assets $a_t \in \mathbb{R}$ and a multidimensional type (w_t, κ_t, ξ_t) . The household earns a wage $w_t \in \mathbb{R}^+$, has a cost of transfer take-up $\kappa_t \in \mathbb{R}^+$, and has a persistence component $\xi_t \in \mathbb{R}$. The conditional distribution of κ_t is assumed to be continuous, that is, without mass points. Households apply a constant discount rate $\rho \in (0, \infty)$ in discounting utility flows over time.

Household heterogeneity beyond wages has two purposes. First, the type ξ_t encodes the persistence or variability of wages and take-up costs over time. This model element allows for heterogeneous dynamics across households in types (w_t, κ_t) . An important example is household income processes that vary in both a persistent component and a transitory component.³⁰

³⁰Blundell et al. (2008), for instance, consider an income process of the following form: $w_t = P_t(\xi_t) + \epsilon_t$, where ϵ_t is a moving-average process and $\Delta\xi_t = \xi_{t+1} - \xi_t$ is a martingale difference sequence.

Second, the take-up cost κ_t sets up the joint choice of labor supply and transfer take-up. Each period, households first choose how much labor $l_t \in \mathbb{R}^+$ to supply to generate labor income $z_t = w_t l_t$. They make this choice knowing their wage w_t but not their take-up cost κ_t beyond its conditional distribution. Let $\theta_t = (w_t, \xi_t)$ represent the household's information set in choosing labor supply. They then draw κ_t and make their take-up choice, which we denote as $\mathbb{1}_S = 1[S(z_t) \geq \kappa_t]$.

Households' current labor incomes are taxed according to the nonlinear schedule $T(z_t) : \mathbb{R}^+ \rightarrow \mathbb{R}$. Households divide their after-tax labor income, along with any assets or debts carried into the period, between consumption c_t and net assets a_{t+1} . There is also a voluntary transfer with nonlinear schedule $S(z_t) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. Households value a dollar of automatic transfer as equivalent to a cash dollar, so there is no distinction with income taxes.³¹ Under the informational assumptions, the household θ views taxes and expected transfers as a consolidated system in the labor supply choice. Letting $M(z; \theta) = \Pr(S(z) \geq \kappa \mid \theta)$ denote its take-up probability, the household faces an expected marginal "keep" rate of $1 - T'(z) - M(z; \theta)S'(z)$ with respect to labor income z .³² Negative taxes on labor income are possible, and capital income is untaxed.

Each period, households choose consumption, labor income, and savings to maximize its ex-ante value function, according to the Bellman equation:

$$V(a_t, \theta_t) = \max_{c_t, z_t, a_{t+1}} \left\{ u(c_t, z_t; \theta_t) + \frac{1}{1 + \rho} \mathbb{E}_t[V(a_{t+1}; \theta_{t+1}) \mid a_{t+1}; \theta_t] \right\} \quad (8)$$

subject to period budget constraints

$$c_t + a_{t+1} = z_t - T(z_t) + R_t a_t + \int_0^{S(z_t)} [S(z_t) - \kappa] \mu(\kappa \mid \theta_t) d\kappa \quad (9)$$

and a borrowing constraint $a_{t+1} \geq \bar{a}_{t+1}(\theta_t)$. The variable R_t is the gross interest rate. Transversality and no-Ponzi conditions are enforced by further imposing $\lim_{t \rightarrow \infty} (1 + \rho)^{-t} a_{t+1} = 0$.

Our key assumption on preferences over consumption and work hours is that they take the form in [Greenwood et al. \(1988\)](#): $u(c, z; \theta) = U(c - v(z; \theta); \theta)$. These preferences allow for risk aversion and rule out income effects, as is common in the nonlinear optimal income tax literature. For each household, the choice of hours is one-to-one with labor income z , and we therefore model the household as directly choosing z .

³¹Appendix D discusses the implications of valuing in-kind transfers differently than cash. In the model, we also do not explicitly model eligibility rules, treating them as infinite take-up costs. Households therefore do not know more than the eligibility rate at θ_t , although we relax this in our empirical implementation.

³²These assumptions, particularly on the timing of the realization of κ_t , are for tractability. Otherwise, the household's labor supply would depend upon their take-up choice, which is discontinuous and would itself depend on labor supply. These assumptions ensure that the transfer system enters the labor-supply choice smoothly, and in particular through the conditional probability of receipt and the benefit amount.

Welfare. We define social welfare as a weighted sum of the ex-ante value functions using the welfare weights $\alpha(\theta)$:

$$W = \int_{\Theta} \alpha(\theta) V(a; \theta) d\mu(a, \theta),$$

where μ is the joint distribution of assets and types, and the government's instruments are the automatic and voluntary transfer schedules $T(z)$ and $S(z)$. We normalize the population-average welfare weight $E[\alpha(\theta)]$ to one. When we report money-metric welfare, we rescale W by the population average social marginal value of income $E[\alpha(\theta) V'_a(a; \theta)]$.

Before proceeding with our analysis, we restate our previous definition of self-targeting (Equation 3) within this model.

Definition 1. *In the cross-section of eligible people with an income z , transfer take-up is “advantageously self-targeted” if the expected social marginal value of income decreases with take-up cost: $E[\alpha(\theta) V'_a(a; \theta) | \kappa, z]$ is a decreasing function of κ for each z .*

This definition embeds two motives for targeting some households with transfers: insurance and redistribution. First, risk-averse households vary in their marginal utilities of income $V'_a(a; \theta)$ over time, due to uninsurable shocks to their types (w, ξ) . Second, households vary in welfare weights $\alpha(\theta)$, which we introduce following [Saez and Stantcheva \(2016\)](#). We use these weights to penalize inequality in lifetime incomes that goes beyond household risk aversion alone.

For either motive, social benefits may arise when take-up costs κ lead households to make take-up choices that are correlated with the type ξ , over which society may have preferences. A specific example would be if households can smooth consumption, and ξ encodes a household's long-run average wage and w encodes their current wage. At any w , low- ξ households would consume less than high- ξ households. If, all else equal, take-up costs κ are rising in ξ , then selection into take-up on ξ would be advantageous, consistent with our results in Section 3. Intuitively, self-targeting is advantageous when, within sets of households that otherwise look identical to the government, the first to take up transfers are the “neediest” in the society's judgment.

Reforms. In the voluntary reform, the available transfer $S(z)$ is decreased by ds at every income z . In the automatic reform, taxes at z are cut by $\tau(z) = \bar{M}(z) ds$, where $\bar{M}(z) = E_{\theta} [M(z; \theta) | z]$ is the take-up rate at income z . This is equivalent to raising an automatic transfer of $\bar{M}(z) ds$ valued at par with cash. By consequence, people at each income level receive the same average change in taxes minus transfers between the voluntary and automatic reforms. However, unlike the level-shift voluntary reform, the automatic reform changes marginal rates by $\tau'(z) = \frac{d}{dz} \bar{M}(z) ds$ at z . In both reforms, the fiscal costs of marginal transfer recipients and labor supply responses are covered through lump-sum taxes on all households.³³

³³In a more general class of marginal transfer reforms, one would have to account for redistribution both between

We see this specific contrast as the natural way to quantify the welfare gains from self-targeting. First, by fixing the transfer's budget, we focus attention on optimal design within a program, rather than the program's merits versus another use of funds. Second, as in [Kaplow \(2011\)](#), requiring neutrality with respect to the income distribution removes incidental impacts of the policy change on the overall progressivity of taxes and transfers. Our welfare calculations thus do not reflect changes in progressivity that could, in principle, be achieved by income taxes alone. Third, a flat change in the transfer has an intuitive real-world analog: changing a fully-voluntary transfer into one with a small automatic transfer with a large top-up provided upon application.

4.3 Theoretical Result

Proposition 1. *The difference in the welfare effects of increases in automatic transfers versus voluntary transfers (automatic minus voluntary) is*

$$\begin{aligned} \frac{dW}{ds} = & \underbrace{-\beta\sigma_M^2}_{\text{lost value of self-targeting}} + \underbrace{\bar{M}\bar{\varepsilon}_b}_{\text{fiscal cost of marginals}} \\ & + \underbrace{\int_z \frac{\bar{M}'(z)z\bar{\varepsilon}_\tau(z)}{1-T'(z)} \frac{d}{dz} \left(S(z)\bar{M}(z) - T(z) \right) dH(z)}_{\text{labor supply effect}}, \end{aligned} \quad (10)$$

where β is the coefficient on take-up from a regression of welfare weights on take-up controlling for income, $H(z)$ is the income distribution, $\bar{M}(z) = \int_\xi M(z; \theta) d\mu(\xi|z)$ is the take-up rate at income z , $\bar{M} = \int_z \bar{M}(z) dH(z)$ is the overall take-up rate, $\sigma_M^2 = \int_z \bar{M}(z)(1 - \bar{M}(z)) dH(z)$ is the within-income variance in take-up, $\bar{\varepsilon}_b$ is an average take-up elasticity with respect to benefit size, and $\bar{\varepsilon}_\tau(z) = \int_\xi \frac{1-T'(z)}{z} \frac{\partial z}{\partial(1-\tau)}(\theta) d\mu(\xi|z)$ is the elasticity of income with respect to a small change τ' in the marginal tax rate of those with initial income z , as in [Jacquet and Lehmann \(2014\)](#).

Proposition 1 shows our result in the basic environment carries over into the Mirrleesian setting with a labor-supply choice. A regression coefficient still summarizes the targeting benefits, and a take-up elasticity still summarizes the ordeal costs. We show formally in [Appendix D](#) that, when self-targeting is advantageous and $S(z)$ is positive, the first term in Equation 10 is negative. That is, society loses some benefits of self-targeting to move toward automatic transfers.

The welfare-relevant coefficient β is a weighted comparison of marginal social value of income of recipients and similar nonrecipients. In particular, $\beta = \Delta E[\alpha(\theta)V'_a(a; \theta)]$, where α is the social welfare weight, V'_a is the marginal utility of income, r is the lifetime-income rank, and z is

and across incomes, fiscal savings from marginal recipients, and labor-supply effects. With non-marginal changes to voluntary transfers, one cannot apply the envelope theorem to reveal ordeal costs.

the current-income rank. The difference operator is defined to be $\Delta E[y] = \int_z \omega(z)(E[y | z, D = 1] - E[y | z, D = 0]) dH(z)$ for some outcome y , income z , and transfer receipt D . The weights are according to the variance of receipt: $\omega(z) = \bar{M}(z)[1 - \bar{M}(z)] / \int_z \bar{M}(z)[1 - \bar{M}(z)] dH(z)$.

Yet there are some differences with the basic environment. First, we incorporate insurance value, which appears through the welfare-relevant β coefficient. Second, the formula adds a term to account for the fiscal impact of the labor supply response to the reform. As we prove in Appendix D, the term is negative when the tax system is optimal and take-up decreases in income ($\bar{M}'(z) < 0$). Under these assumptions, the automatic reform requires higher marginal tax rates to offset the cut to the voluntary transfer, reducing labor supply.

4.4 Interpretation

We make five observations about how we understand this theoretical result, related to program scale, insurance value, ordeal costs, and its application to in-kind transfers and non-marginal reforms.

First, the difference in welfare between the reforms (dW/ds) varies in an intriguing way with take-up in the voluntary transfer. To illustrate this, observe that universal take-up ($\bar{M} = 1$) has no targeting benefits and that take-up is costly. On the other hand, take-up can also be “too low”: Inducing the first person to take up has no social value, as by the envelope theorem, their take-up cost exactly equals the benefits they receive. That is, only after the first person takes up are there any inframarginal households and thus are first-order welfare gains possible. This suggests an inverse U -shaped relationship between the welfare impact dW/ds and take-up $M(s)$, although it clearly depends upon the labor-supply effect, as well as how β and $\bar{\varepsilon}_b$ might vary across levels of s .

Second, Proposition 1 allows for distinctions between the redistribution-related and insurance-related welfare benefits. Holding fixed the other sufficient statistics, risk aversion enters only through the self-targeting coefficient β . Our decomposition is as follows:

$$\beta = \underbrace{\Delta E[\alpha E[V'_a]]}_{\text{redistribution}} + \underbrace{\Delta E\left[\alpha E[V'_a] \left(\frac{V'_a}{E[V'_a]} - 1\right)\right]}_{\text{insurance}}, \quad (11)$$

where $E[V'_a]$ is a person’s average marginal utility over their lifetime. This decomposition follows our approach in Section 3 of distinguishing between- and within-lifetime self-targeting, while further incorporating welfare weights. That is, people can be currently needy (high $\alpha V'_a$) either because they have high need in all states of life (high $\alpha E[V'_a]$) or they are high current need relative other states (high $V'_a/E[V'_a]$).

Third, our reform differs fundamentally from changes to ordeals. We instead take the ordeal as given and reallocate resources between voluntary and automatic transfers. Our welfare formula

therefore weighs the value of transfers to inframarginal recipients against ordeal costs to *marginal* recipients that disenroll when the voluntary transfer is cut. By contrast, 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](#); [Naik, 2025](#)).

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. However, envelope-based estimates of ordeal costs require households to make transfer take-up decisions optimally. This assumption might be seen as problematic, since research has found non-optimizing behavior in take-up ([Bhargava and Manoli, 2015](#); [Finkelstein and Notowidigdo, 2019](#); [Anders and Rafkin, 2022](#)). Yet the optimizing assumption works against our conclusion, giving 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 at a lower real resource cost than if households were optimizing.³⁴

Fourth, our results extend immediately to in-kind transfers, or other transfers that households may value differently from their cost of provision. Suppose households are willing to pay $\lambda < 1$ for each transfer dollar. Our expression for dW/ds from Proposition 1 would then be multiplied by λ , as we show in Appendix D. The welfare effect dW/ds remains correct in units of households' willingness-to-pay for the transfer, rather than in dollars. The parameter λ simply rescales the welfare effect and cannot reverse its sign.

Fifth, we study marginal reforms, in which marginal recipients only matter for welfare through fiscal externalities, as they are indifferent to take up. However, in a non-marginal reform, people whose take-up choices change due to the reform are not necessarily indifferent. By consequence, who they are—that is, self-targeting among marginal recipients—would affect welfare.

4.5 Calibration

We draw on our estimates and external inputs to calibrate Proposition 1 (see Appendix D for additional details). We also review where, in our assessment, reasonable scholarly disagreement persists about the appropriate calibrated values, along with its implications for our analysis.

Policy Parameters. We obtain transfer receipt rates $\bar{M}(z)$, per-capita average benefits $S(z)$, and the income distribution $H(z)$ all from the PSID. The receipt-rate variance σ_M^2 comes from the estimates of $\bar{M}(z)$. We impose a piecewise-linear tax schedule $T(z)$ fit using average marginal rates that reflect federal and state taxes on income and payroll ([Congressional Budget Office, 2015](#)).

Preferences. The next input for welfare analysis, the self-targeting coefficient β , requires us to

³⁴Other behavioral forces could raise take-up, such as limited self-control ([Chan, 2017](#)). If larger in magnitude than misperceptions, these would lead us to understate ordeal costs.

select welfare weights and the household’s period utility function. Our baseline estimates impose the same assumptions as in Section 3. These were a [Greenwood et al. \(1988\)](#) specification of household period utility, with a risk-aversion parameter over consumption of $\gamma = 2$ ([Chetty and Finkelstein, 2013](#)) and labor supply elasticity $\eta = \bar{\epsilon}_\tau = 0.3$ ([Saez et al., 2012](#)). We draw on the PSID microdata for household equivalized consumption as c and labor hours as l .

We set the take-up elasticity $\bar{\epsilon}_b$ to 0.4, consistent with literature reviews by [Bound and Burkhauser \(1999\)](#) and [Krueger and Meyer \(2002\)](#) on disability and unemployment insurance respectively. Other papers estimating U.S. take-up elasticities include [McGarry \(1996\)](#) ($\bar{\epsilon}_b = 0.5$, for SSI) and [Pukelis \(2024\)](#) ($\bar{\epsilon}_b = 0.1$, for SNAP).

Welfare Function. Our baseline results use a constant-elasticity social welfare function with parameter $\gamma_{sp} = 1$, implying a moderate social value of redistribution. Stronger concavity in the social welfare function (higher γ_{sp}) would strengthen the redistributive motive and thus favor voluntary transfers if self-targeting is advantageous. Having fixed this parameter, we use the joint distribution of income and consumption to compute the welfare weight for each household. We then estimate differences in welfare weights between transfer recipients and non-recipients conditional on income, similar to the self-targeting analysis in Section 3.

4.6 Welfare Results

Results. Panel A of Table 3 reports our primary estimates of the welfare effects of reallocating resources from a given voluntary transfer to an automatic one. This shift incurs a welfare cost due to weaker targeting, and Columns 1 and 2 decompose this cost into its redistribution and insurance components, following Equation 11. Column 3 shows the fiscal savings on marginal people who exit the transfer (the second term in Equation 10). Due to the envelope argument explained above, Column 3 can also be interpreted as the social savings on ordeal costs among marginal people. Column 4 shows the labor-supply effect (the third term in Equation 10). Column 5 shows the net welfare effect of the reform.

Taking SNAP as an example, we find a shift toward automatic transfers loses some benefits of self-targeting. On the margin, self-targeting is worth about 44 cents per dollar of SNAP. About 39 cents of these benefits are attributed to between-lifetime redistribution and 5 cents to within-lifetime insurance. The insurance share of benefits is small, consistent with Panel C of Figure 1. When voluntary transfers are cut, the government saves 17 cents per SNAP dollar from marginal people who no longer take up. Finally, the automatic transfer increases marginal tax rates, which reduces labor supply and imposes a small fiscal externality (0.9 cents). Together, the net effect of making SNAP more automatic on the margin is a net social loss of 28 cents per dollar.

We can check these conclusions with back-of-the-envelope calculations. In Panel D of Figure 1,

we found social benefits of \$1.94 per dollar redistributed to SNAP recipients from SNAP-eligible nonrecipients. Table 1 also reported a SNAP take-up rate of 43 percent. Applying our welfare formula (Equation 10), we come close to the social benefits: $(1.94)(0.43)(1 - 0.43) = 0.48$. Similarly, we match the social costs of ordeals by multiplying the take-up rate and the take-up elasticity: $(0.43)(0.4) = 0.17$. Finally, labor-supply effects are complex to calculate, but their small size reflects that the reforms have little net impact on average marginal tax rates.

Overall, we find a stark trade-off between the benefits of self-targeting and the costs of ordeals. Looking across transfers, 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 are not well-approximated by the social savings on ordeal costs alone. That result is reflected in the dollar-weighted average, which shows that the social benefits of self-targeting actually exceed the social costs of ordeals. In summary, self-targeting appears viable as an argument for status-quo voluntary transfers over automatic ones.

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. Automatic benefits, by contrast, appear highly 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 generating self-targeting.

4.7 Sensitivity Analysis and Limitations

Panel B examines the sensitivity of our results to the calibrated parameters. Across these versions of our analysis, self-targeting remains an important advantage of voluntary transfers. Indeed, self-targeting typically offsets most if not all of the social costs of ordeals, even at upper-bound values. We then raise several limitations of the welfare analysis.

Risk Aversion and Social Preferences. The more risk-averse are people, and the more society cares about inequality, the larger are the welfare losses from forgoing self-targeting in transfers. Put another way, automating transfers is likely to be desirable only when people care less about consumption fluctuations or when society cares less about lifetime inequality.

Elasticities. The take-up elasticity helps to determine fiscal externalities and thus the implied ordeal costs. If take-up is more responsive to benefits than we assume, this implies larger ordeal costs on the margin and thus could motivate automatic transfers. Results are less sensitive to the taxable-income elasticity, as the net of the reforms leaves marginal tax rates mostly unchanged.

Table 3: Welfare Effects of Budgetary Shifts Toward Automatic Transfers (Cents per Transfer Dollar)

	Self-Targeting Gains		Other Forces		Total
	Redistribution (1)	Insurance (2)	Upper Bound on Ordeals (3)	Labor-Supply Effects (4)	(5)
<i>Panel A: Primary Estimates</i>					
Dollar-Weighted Average	-27.5	-3.2	15.8	-0.7	-15.6
SNAP	-38.9	-5.4	17.4	-0.9	-27.8
Medicaid	-29.7	-2.2	21.0	-0.7	-11.6
Housing Assistance	-22.2	-3.2	6.1	-0.5	-19.8
TANF	-7.5	-0.8	4.0	-0.6	-4.9
SSI	6.6	-2.4	25.0	-0.8	28.4
School Lunch	-7.9	-0.8	18.2	-0.6	8.9
WIC	-14.2	-3.4	19.1	-1.5	-0.1
LIHEAP	-13.3	-1.5	7.5	-0.4	-7.8
<i>Panel B: Sensitivity (Dollar-Weighted Average)</i>					
Less progressive social preferences ($\gamma_{sp} = \frac{1}{2}$)	-22.5	-2.6	15.8	-0.7	-10.0
More progressive social preferences ($\gamma_{sp} = 2$)	-35.6	-4.4	15.8	-0.7	-24.8
Less risk averse ($\gamma = 1$)	-17.8	-1.4	15.8	-0.7	-4.0
More risk averse ($\gamma = 3$)	-35.0	-5.0	15.8	-0.7	-24.8
Take-up elasticity $\eta = 0.2$	-27.5	-3.2	7.9	-0.7	-23.5
Take-up elasticity $\eta = 0.6$	-27.5	-3.2	23.8	-0.7	-7.7
Elasticity of taxable income $\varepsilon = 0.15$	-27.6	-3.2	15.8	-0.4	-15.3
Elasticity of taxable income $\varepsilon = 0.6$	-27.5	-3.2	15.9	-1.4	-16.2

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 $\gamma = 1$. We calibrate the fiscal cost of marginals by assuming the take-up elasticity is $\bar{\varepsilon}_b = 0.4$. We calibrate the elasticity of taxable income at $\varepsilon_\tau = 0.3$. All columns report the money-metric welfare gains in cents per transfer dollar. Columns correspond to the terms of Equation 10, where we divide each term by the average social marginal utility, which yields a money-metric interpretation.

Limitations. First, our reforms are hypothetical. We consciously ignore many practicalities of implementing automatic transfers, which may be quite important (see, e.g., [Wu and Meyer, 2024](#)). Our analysis therefore cannot directly address the welfare impacts of real-world policy proposals, such as making SNAP fully automatic without offsetting fiscal adjustments. These policies involve additional parameters including inframarginals’ ordeal costs, reduced stigma from fully-automatic programs, and the marginal cost of public funds. Nor do such policies need to be marginal or budget-balanced by cutting spending within the program. Second, the analysis 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 they likely favor automatic transfers. Third, we ignore behavioral responses to transfers beyond take-up and labor supply, such as cross-program take-up spillovers or dynamic incentives for human-capital investment. Fourth, we assume homogeneous labor supply elasticities, which precludes consideration of self-targeting on elasticities rather than levels.

5 Conclusion

This paper studies the positive and normative implications of “self-targeting” in U.S. transfer programs, in which eligible people self-select into take-up on the basis of need. We find substantial self-targeting on consumption and lifetime income across eight U.S. transfers. That is, transfer recipients appear needier on average than eligible nonrecipients, even when holding income fixed. This result is surprising in light of prior evidence, which has reached mixed conclusions as to how changes in take-up costs (“ordeals”) have affected program targeting. Combining our estimates with a new social-welfare formula, we estimate that in some transfers, the benefits of using take-up as a targeting device may plausibly exceed its costs.

Several limitations mean this study is best understood as a first effort to evaluate whether self-targeting can rationalize the widespread use of voluntary take-up in U.S. transfer programs. Among them are the difficulties of measuring income, consumption, transfer receipt, and transfer eligibility in our data. In addition, we have relied upon an adjusted measure of household consumption as an imperfect proxy for marginal utility. Our welfare test, while transparent, does not allow us to contemplate optimal transfer design or non-marginal reforms. Finally, we have abstracted from many dimensions of transfers. Among the ignored aspects are administrative costs, other ways to acquire information (e.g., disability reviews), and some aspects of household heterogeneity (e.g., selection on elasticities). These limitations of our study offer directions for future research.

References

- Aguiar, Mark and Erik Hurst**, “Consumption Versus Expenditure,” *Journal of Political Economy*, 2005, 113 (5), 919–948.
- 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 Ralfkin**, “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 and CHIP Enrollments,” *American Journal of Health Economics*, 2024, 10 (2), 237–271.
- Armour, Philip**, “The Role of Information in Disability Insurance Application: An Analysis of the Social Security Statement Phase-In,” *American Economic Journal: Economic Policy*, 2018, 10 (3), 1–41.
- Auerbach, Alan J, Laurence J Kotlikoff, and Darryl Koehler**, “U.S. Inequality and Fiscal Progressivity: An Intragenerational Accounting,” *Journal of Political Economy*, forthcoming.
- Baily, Martin Neil**, “Some Aspects of Optimal Unemployment Insurance,” *Journal of Public Economics*, 1978, 10 (3), 379–402.
- 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.
- Bhattacharya, Jay and Darius Lakdawalla**, “Does Medicare Benefit the Poor?,” *Journal of Public Economics*, 2006, 90 (1-2), 277–292.
- Blundell, Richard and Ian Preston**, “Consumption Inequality and Income Uncertainty,” *The Quarterly Journal of Economics*, 1998, 113 (2), 603–640.
- **and Luigi Pistaferri**, “Income Volatility and Household Consumption: The Impact of Food Assistance Programs,” *Journal of Human Resources*, 2003, pp. 1032–1050.

- , — , and **Ian Preston**, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921.
- , **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.
- Bound, John and Richard V Burkhauser**, “Economic Analysis of Transfer Programs Targeted on People with Disabilities,” *Handbook of Labor Economics*, 1999, 3, 3417–3528.
- Brewer, Mike, Ben Etheridge, and Cormac O’Dea**, “Why Are Households that Report the Lowest Incomes So Well-Off?,” *Economic Journal*, 2017, 127 (605).
- Chan, Marc K**, “Welfare Dependence and Self-Control: An Empirical Analysis,” *Review of Economic Studies*, 2017, 84 (4), 1379–1423.
- Chetty, Raj**, “A General Formula for the Optimal Level of Social Insurance,” *Journal of Public Economics*, 2006, 90 (10-11), 1879–1901.
- , “Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods,” *Annu. Rev. Econ.*, 2009, 1 (1), 451–488.
- and **Adam Szeidl**, “Consumption Commitments and Habit Formation,” *Econometrica*, 2016, 84 (2), 855–890.
- and **Amy Finkelstein**, “Social Insurance: Connecting Theory to Data,” in “Handbook of Public Economics,” Vol. 5, Elsevier, 2013, pp. 111–193.
- 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.
- Cutler, David M and Lawrence F Katz**, “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980’s,” *The American Economic Review*, 1992, 82 (2), 546–551.
- 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.
- Deaton, Angus**, *Understanding Consumption*, Oxford University Press, 1992.
- Deshpande, Manasi and Lee M Lockwood**, “Beyond Health: Nonhealth Risk and the Value of Disability Insurance,” *Econometrica*, 2022, 90 (4), 1781–1810.

- **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.
- Domurat, Richard, Isaac Menashe, and Wesley Yin**, “The Role of Behavioral Frictions in Health Insurance Marketplace Enrollment and Risk: Evidence from a Field Experiment,” *American Economic Review*, 2021, 111 (5), 1549–1574.
- 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.
- Esping-Andersen, Gosta**, *The Three Worlds of Welfare Capitalism*, Princeton University Press, 1990.
- 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.
- , **Nathaniel Hendren, and Erzo FP Luttmer**, “The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment,” *Journal of Political Economy*, 2019, 127 (6), 2836–2874.
- Friedman, Milton**, *A Theory of the Consumption Function*, Princeton University Press, 1957.
- Fullerton, Don and Diane Lim Rogers**, *Who Bears the Lifetime Tax Burden?*, Brookings Institution Washington, DC, 1993.
- 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.
- Giannella, Eric, Tatiana Homonoff, Gwen Rino, and Jason Somerville**, “Administrative Burden and Procedural Denials: Experimental Evidence from SNAP,” *American Economic Journal: Economic Policy*, 2024, 16 (4), 316–340.
- Gray, Colin**, “Leaving Benefits on the Table: Evidence from SNAP,” *Journal of Public Economics*, 2019, 179, 104054.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman**, “Investment, Capacity Utilization, and the Real Business Cycle,” *American Economic Review*, 1988, pp. 402–417.
- Gruber, Jonathan**, “Cash Welfare as a Consumption Smoothing Mechanism for Divorced Mothers,” *Journal of Public Economics*, 2000, 75 (2), 157–182.
- 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.

—, **Hilary Hoynes, Jamila Michener, and Donald Moynihan**, “Introduction: Administrative Burden as a Mechanism of Inequality in Policy Implementation,” *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2023, 9 (4), 1–30.

Hoynes, Hilary W and Erzo FP Luttmer, “The Insurance Value of State Tax-and-Transfer Programs,” *Journal of Public Economics*, 2011, 95 (11-12), 1466–1484.

Isaacs, Julia, “The Costs of Benefit Delivery in the Food Stamp Program: Lessons From a Cross-Program Analysis,” Contractor and Cooperator Report 39, USDA Economic Research Service 2008.

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

Kaplow, Louis, *The Theory of Taxation and Public Economics*, Princeton University Press, 2011.

Kleven, Henrik J, “Sufficient Statistics Revisited,” *Annual Review of Economics*, 2021, 13 (1), 515–538.

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.

Ko, Wonsik and Robert A Moffitt, “Take-Up of Social Benefits,” *Handbook of Labor, Human Resources and Population Economics*, 2024, pp. 1–42.

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.

Liebman, Jeffrey B, “Redistribution in the Current US Social Security System,” in “The Distributional Aspects of Social Security and Social Security Reform,” University of Chicago Press, 2002, pp. 11–48.

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.
- Naik, Canishk**, “Mental Health and the Targeting of Social Assistance,” Working Paper 2025.
- Nichols, Albert L and Richard J Zeckhauser**, “Targeting Transfers Through Restrictions on Recipients,” *American Economic Review*, 1982, 72 (2), 372–377.
- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 2019, 37 (2).
- 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.
- Pukelis, Kelsey**, “SNAP Policies and Enrollment Following the COVID-19 Pandemic,” Working Paper 2024.
- 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.

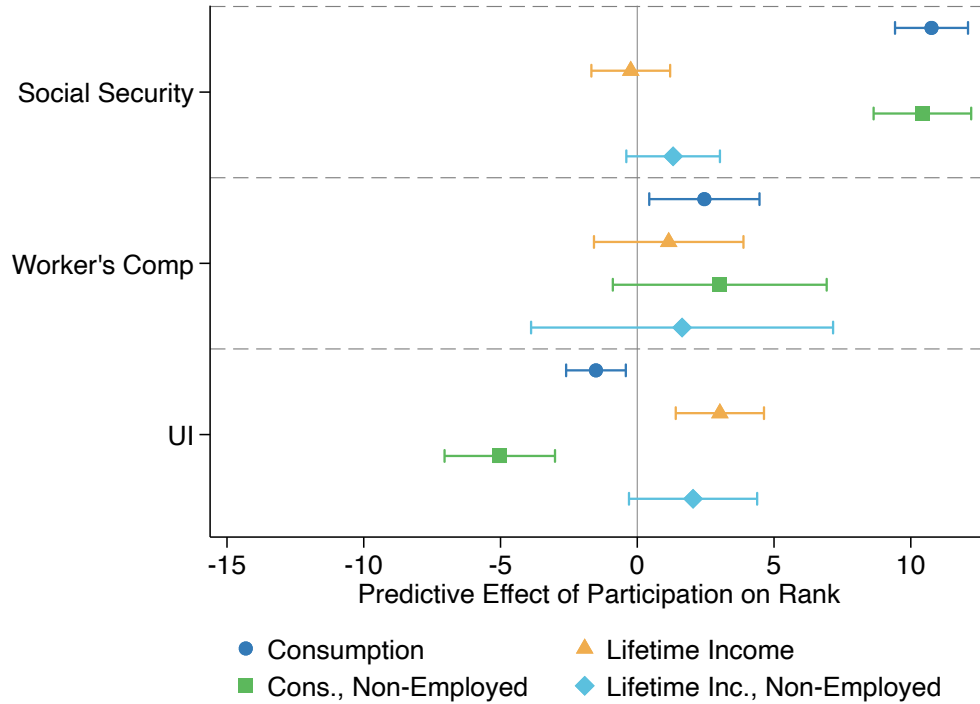
- Shepard, Mark and Myles Wagner**, “Reducing Ordeals through Automatic Enrollment: Evidence from a Subsidized Health Insurance Exchange,” Working Paper 30781, National Bureau of Economic Research 2024.
- Tiehen, Laura, Constance Newman, and John A Kirlin**, “The Food-Spending Patterns of Households Participating in the Supplemental Nutrition Assistance Program: Findings from USDA’s FoodAPS,” Economic Information Bulletin 176, U.S. Department of Agriculture 2017.
- Unrath, Matthew**, “Targeting, Screening, and Retention: Evidence from California’s Food Stamps Program,” Working Paper 2024.
- 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 2024.
- **and Jonathan Zhang**, “Sliding into Safety Net Participation: A Unified Analysis Across Multiple Programs,” *National Tax Journal*, 2025, 78 (1), 000–000.
- Zedlewski, Sheila and Linda Giannarelli**, “TRIM: A Tool for Social Policy Analysis,” Research Report, Urban Institute 2015.
- Zimmermann, Wendy and Karen C Tumlin**, “Patchwork Policies: State Assistance for Immigrants Under Welfare Reform,” Occasional Paper 24, Urban Institute 1999.

Appendices for Online Publication

A Additional Tables and Figures	43
B Data Appendix	93
C Further Results	104
D Theory Appendix	108

A Additional Tables and Figures


Figure A1: Absence of Self-Targeting in Contributory Programs



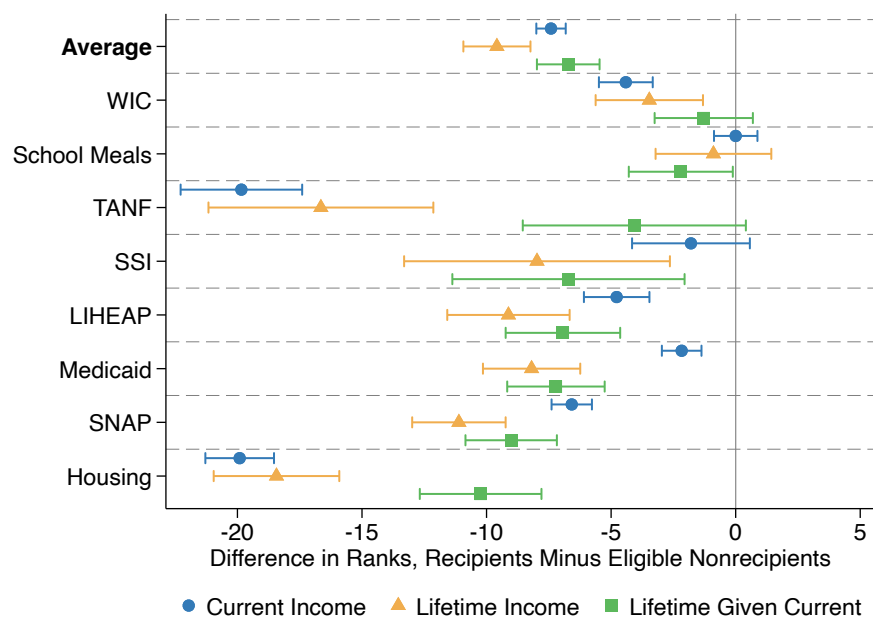
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). We do not have eligibility simulators for these programs. We therefore show these effects for all people (blue circles and yellow triangles), as well as the non-employed (green squares and teal diamonds), which we take as an approximate proxy for eligibility. In this non-elderly adult population, most Social Security recipients are likely to be on Disability Insurance. PSID data do not allow us to distinguish old-age benefits from disability benefits until 2005. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A2: Raw Differences in Income, Consumption, and Lifetime Income by Transfer: Benchmarks for Magnitude of Self-Targeting

Panel A: Consumption

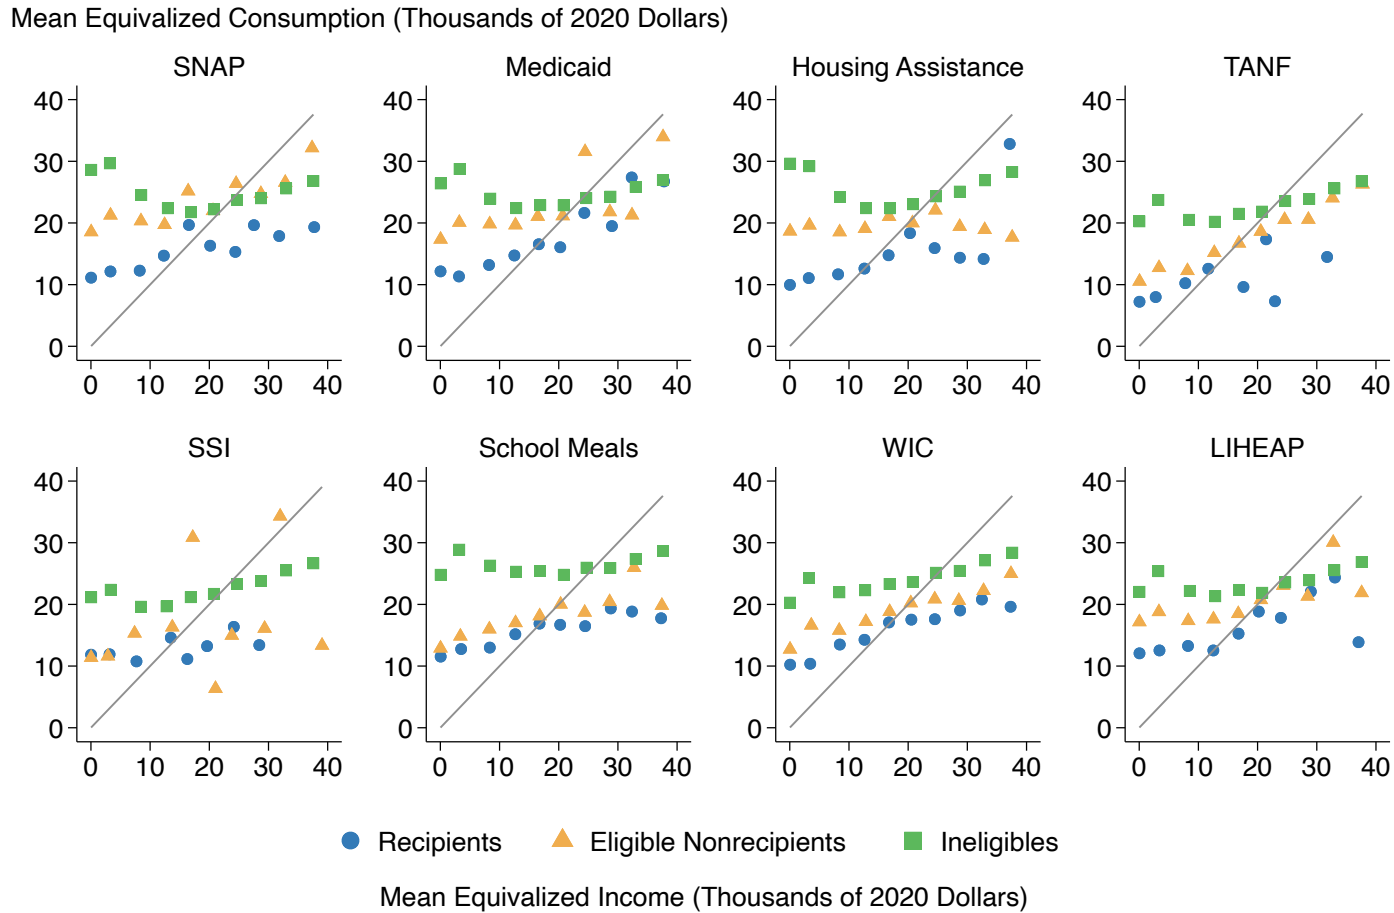


Panel B: Lifetime Income



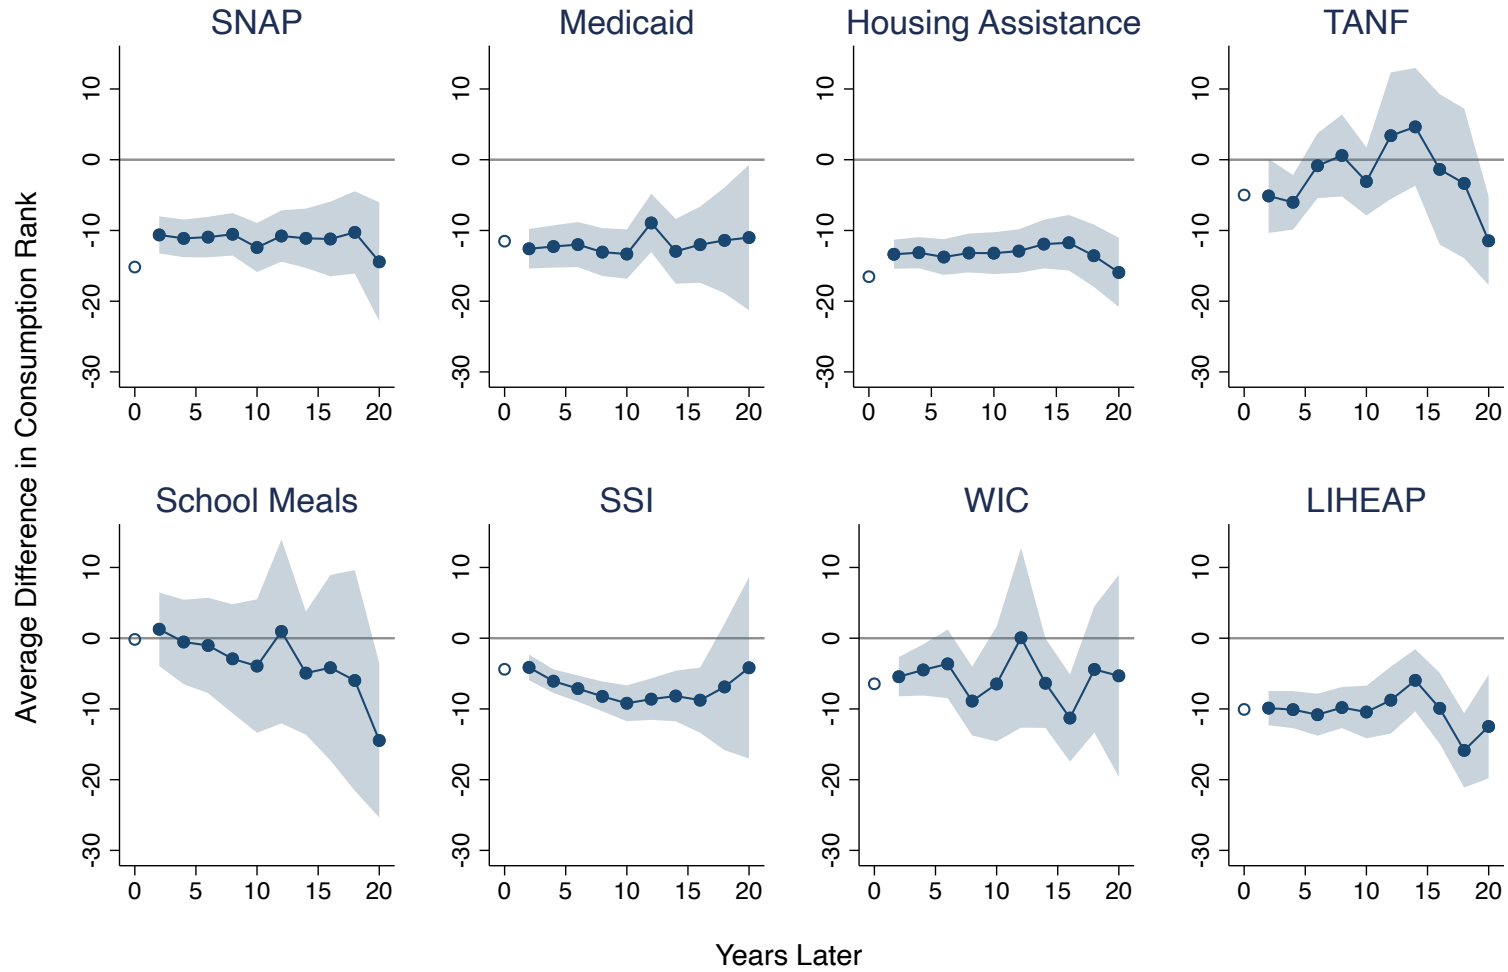
Notes: Each panel of the figure displays three differences in ranks for each transfer. Blue circles show raw differences in average income ranks between a transfer's recipients and its simulated-eligible nonrecipients. Yellow triangles show raw differences in average ranks of consumption (Panel A) or lifetime income (Panel B) between the same groups. Green squares show differences in average ranks of consumption (Panel A) or lifetime income (Panel B), further controlling for current-income rank as in Equation 4. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A3: Relationship Between Consumption and Income by Transfer Receipt and Eligibility




Notes: This figure displays binned scatterplots of the relationship between equivalized real household consumption and income, splitting the data by transfer receipt and eligibility. Each panel presents the consumption–income relationship for three groups (recipients in blue circles, eligible nonrecipients in red triangles, and ineligibles in green squares) for a specific transfer.

Figure A4: Self-Targeting 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 $\bar{R}_{it} = \alpha_{ct} + \beta D_{i,t+k} + f(R_{it}) + u_{it}$, where we plot β for each horizon k . The estimation sample is always restricted to current eligible nonrecipients, $D_{it} = 0$. Shaded regions reflect bootstrapped 95-percent pointwise confidence intervals, with clustering by household.

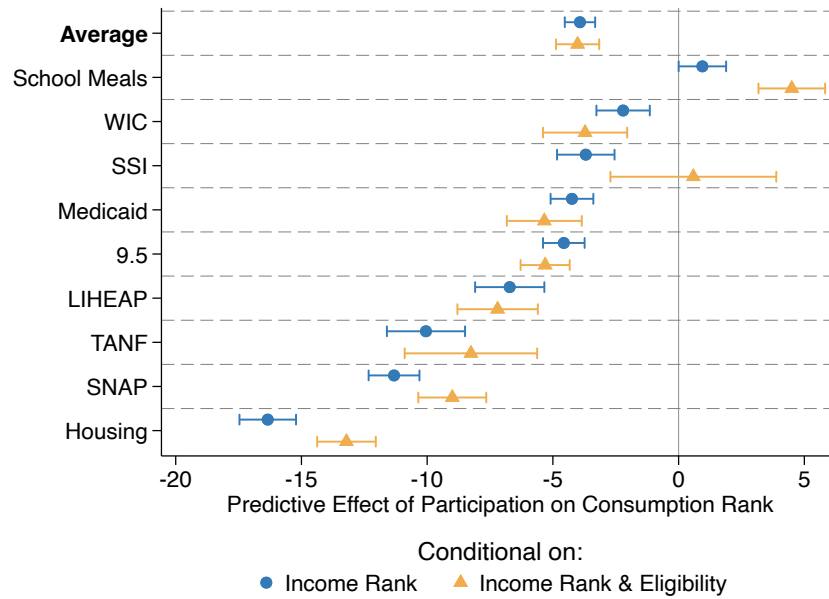
Figure A5: Bounds Analysis of Self-Targeting on Consumption




Notes: This figure displays alternative estimates of self-targeting, adjusting our consumption measure by subtracting or adding the dollar value of the transfers to each household. Blue dots report our baseline estimates, while yellow squares (“minus transfer”) and green diamonds (“plus transfer”) present the resultant bounds. Subtracting transfers from consumption yields an estimate of “first-dollar” self-targeting under the assumption of a marginal propensity to consume out of transfer income of one. Adding transfers onto consumption is informative if households reduce actual or reported consumption, relative to their income, in response to transfers. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A6: Self-Targeting in Transfer Programs (Non-Equivalized Households)

Panel A: Selection on Consumption

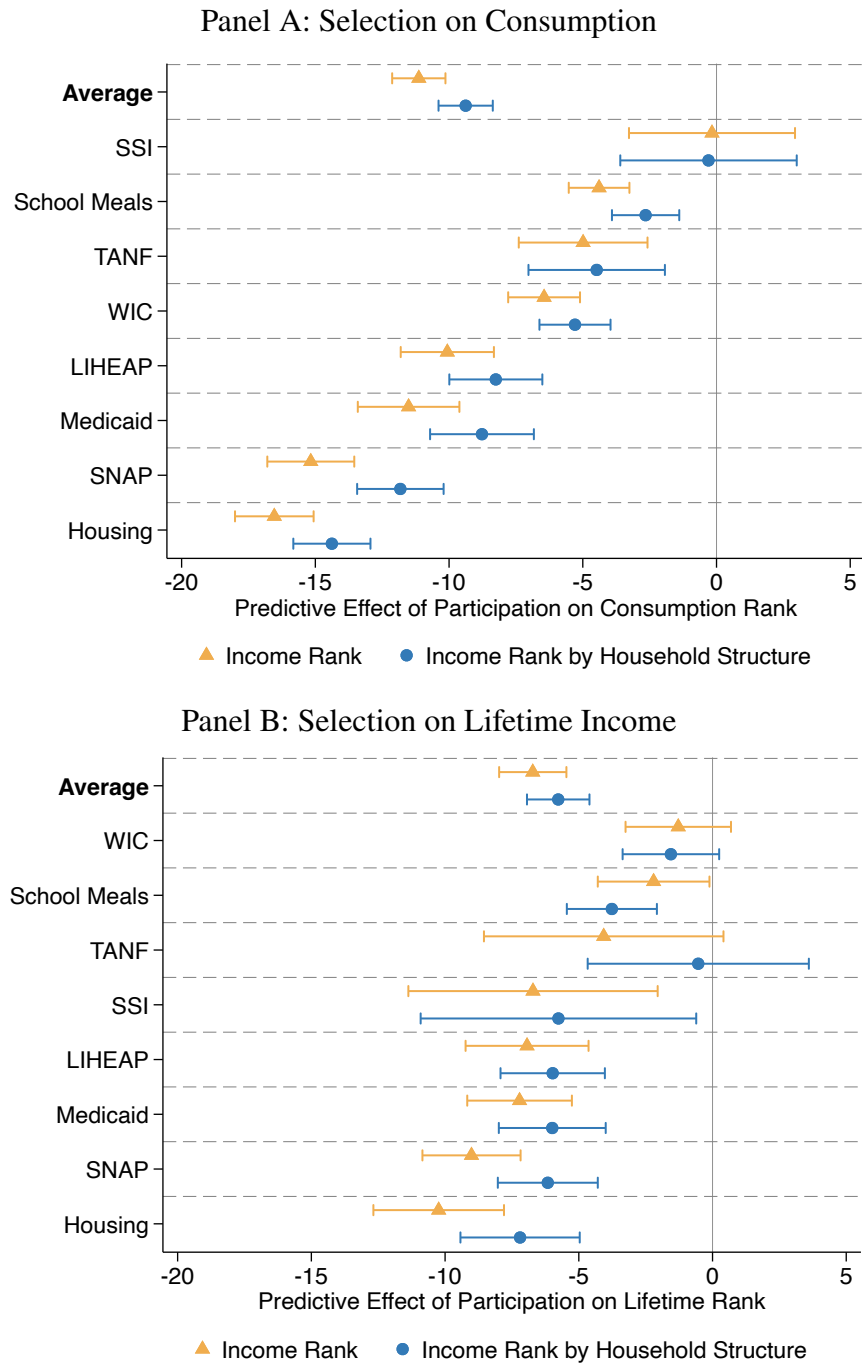


Panel B: Selection on Lifetime Income




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 β from Equation 4). No equivalence-scale adjustment is applied in ranking households on consumption and lifetime income. For the yellow diamonds, we estimate the regression only on people whom we simulate to be eligible. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A7: Selection into Transfer Receipt: Income Interacted with Household Structure



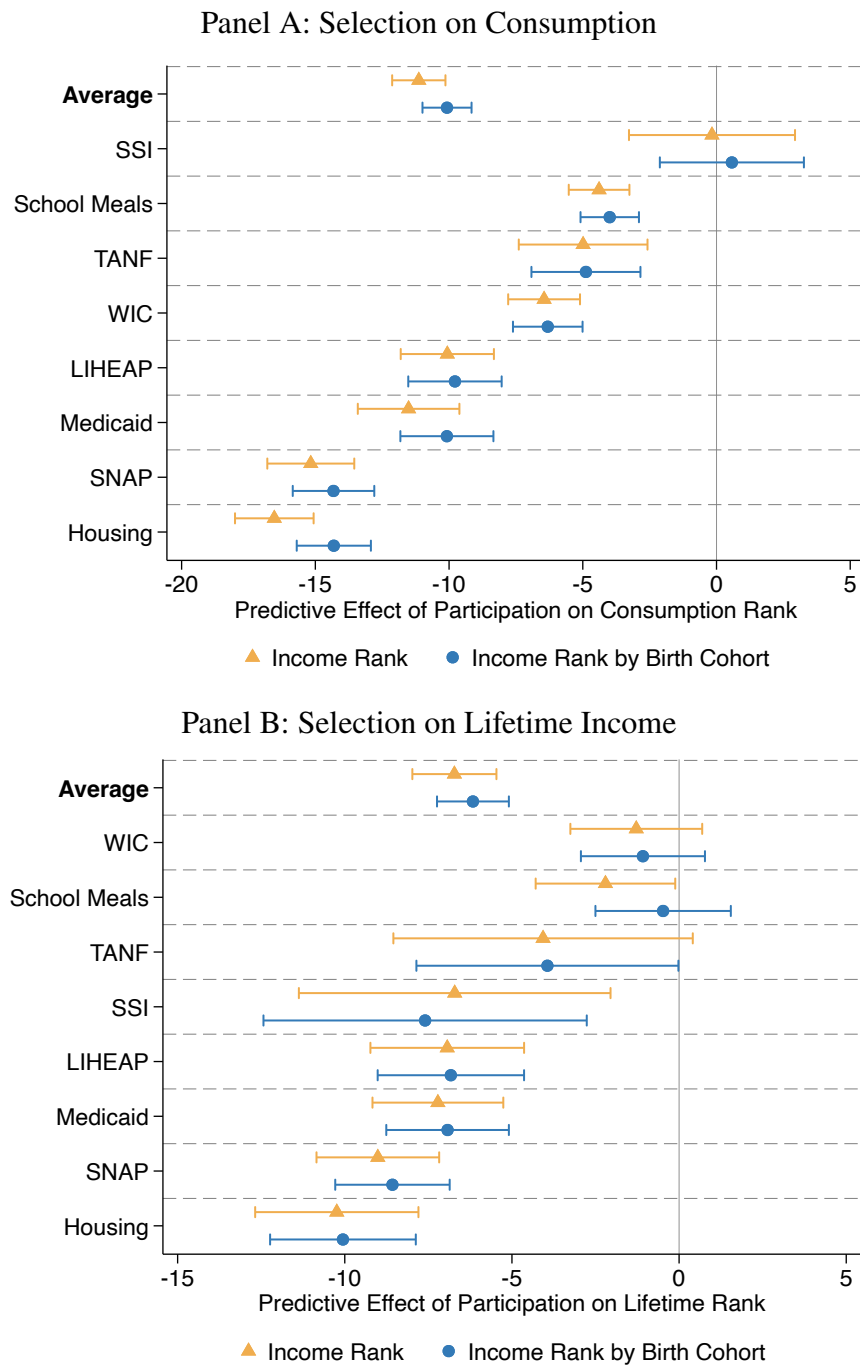
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). All estimates limit the sample to the simulated-eligible. For estimates represented by blue circles, we allow for different splines in household income rank for each category of “household structure.” To approximate the income tax schedule, we define household-structure categories by unique combinations of the number of adults, the number of children (under age 18), and marital status. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A8: Selection on Consumption into Transfer Receipt: Adjusted for Regional Price Parity



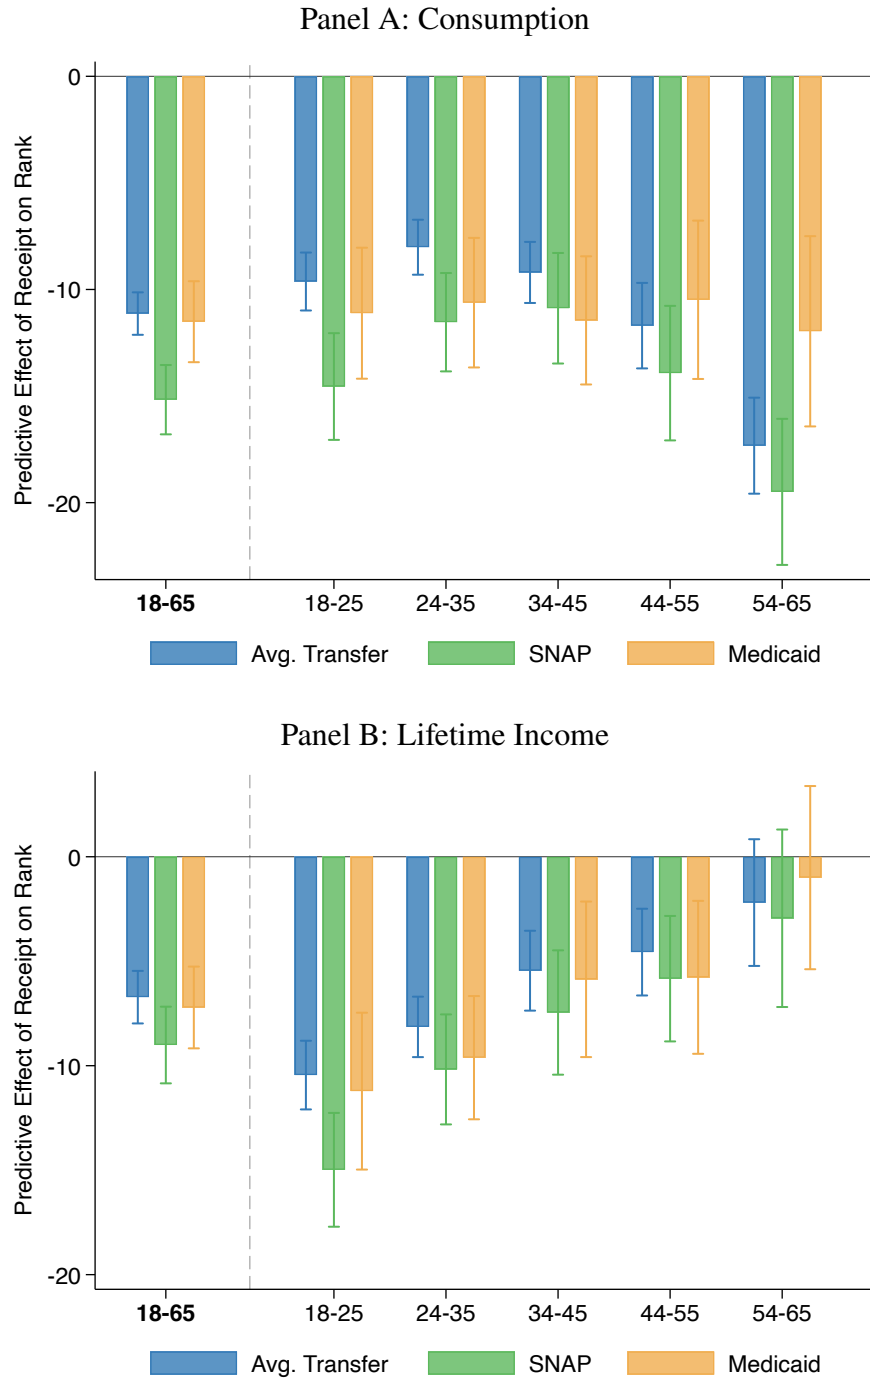
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4). Both yellow diamonds and blue circles restrict the sample to simulated eligibles. For estimates represented by blue circles, we adjust household consumption to account for regional differences in purchasing power, using indices from the American Chamber of Commerce Researchers Association (ACCRA COLI, 1997–2008) and the U.S. Bureau of Economic Analysis (BEA RPP, 2009–2019). The income-rank control is similarly adjusted. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A9: Selection into Transfer Receipt: Income Interacted with Cohort Indicators



Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). All estimates limit the sample to the simulated-eligible. For estimates represented by blue circles, we further control for birth-year cohort fixed effects as well as cohort-specific linear slopes in current-income rank. We continue to control flexibly for current-income rank, but the spline is not cohort-specific due to data limitations. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A10: Self-Targeting by Age




Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4). Results are split into five age groups to the right of the dashed line, with our main-sample results to the left. We report results for the receipt the average transfer in blue, SNAP only in green, and Medicaid only in orange. The average reflects estimates from a stacked specification across the eight programs, weighting by total amounts in 2019. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A11: Robustness of Self-Targeting to Alternative Utility Functions

Panel A: State Dependence in Health Status




Panel B: Nonseparability in Hours



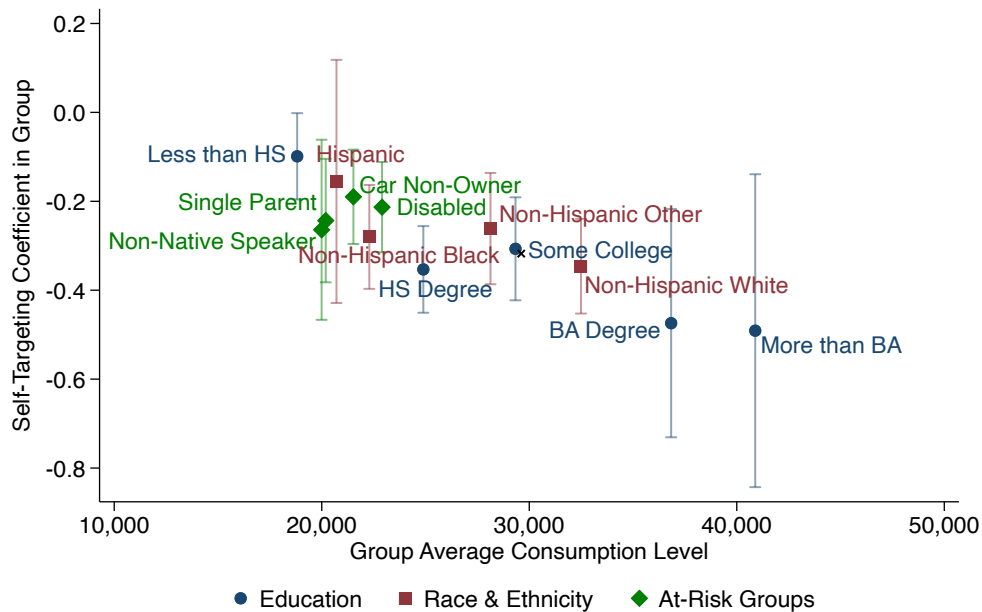
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4). Panel A responds to potential state dependence of marginal utility in self-reported health status. We estimate self-targeting separately by health status: poor or fair (blue circles), good (yellow diamonds), very good or excellent (green squares). Panel B responds to potential nonseparability in labor hours. We estimate self-targeting separately by work status: nonemployed (blue circles), part-time work (yellow diamonds), full-time work (green squares). All regressions limit the sample to the simulated-eligible. The “average” bars reflect a dollar-weighted average over the eight transfers. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A12: Demographic Heterogeneity in Self-Targeting (PSID)

(a) SNAP



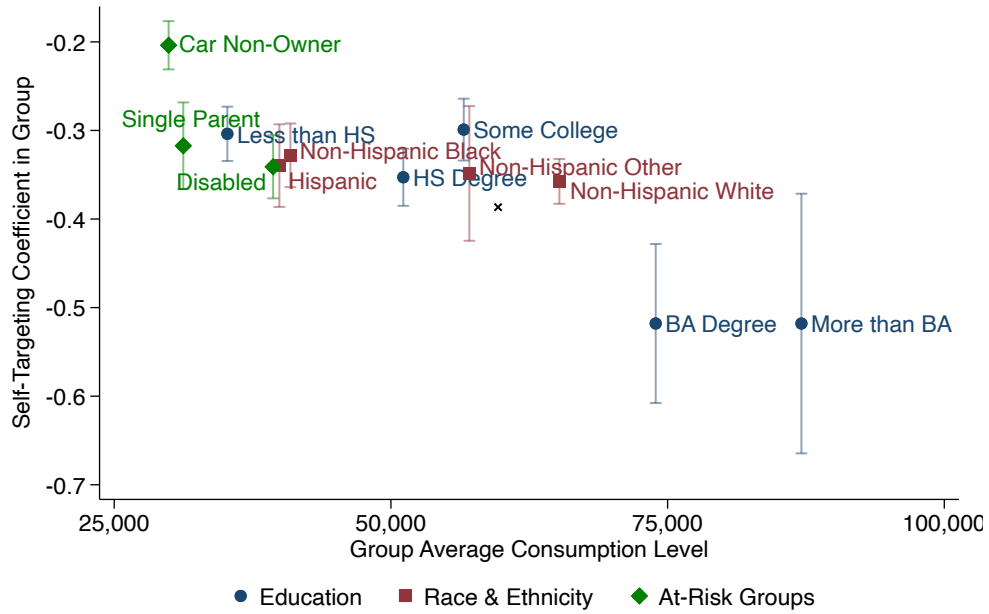
(b) Medicaid




Notes: This figure displays estimates of the predictive effects of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). We split the sample by education (less than HS, HS degree, some college, BA only, more than BA), race/ethnicity (non-Hispanic white, non-Hispanic black, Hispanic, non-Hispanic other), and by “at-risk” group (single parent, disabled, non-native English speaker, car non-owner). The data are simulated-eligible members of the listed demographic group in the PSID. Panels A and B show results SNAP and Medicaid respectively. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household. The black “x” point indicates the population-average self-targeting coefficient and consumption rank.

Figure A13: Demographic Heterogeneity in Self-Targeting (CEX)

(a) SNAP

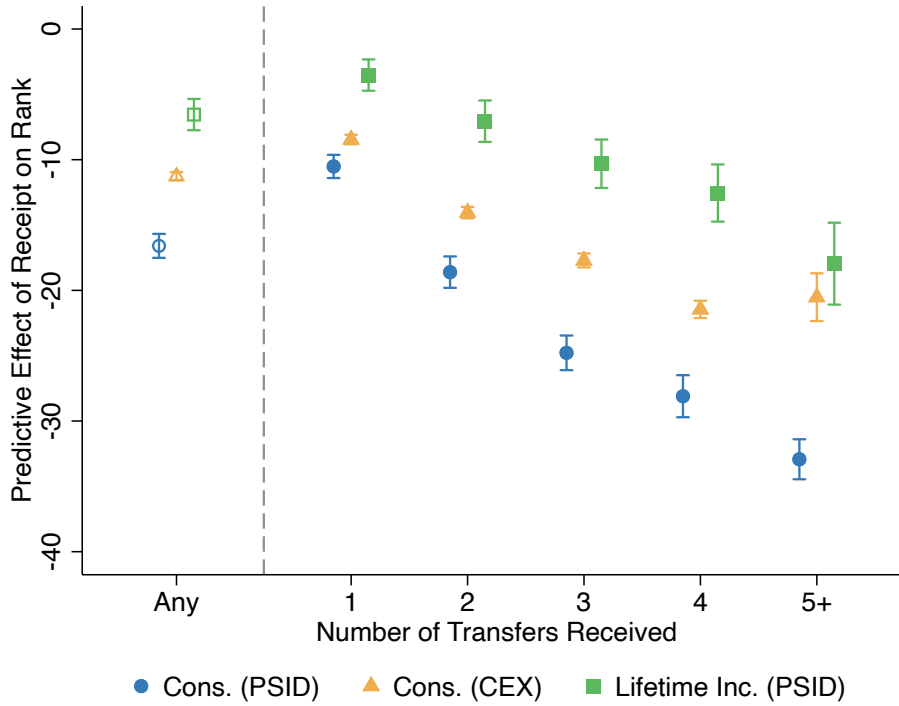


(b) Medicaid



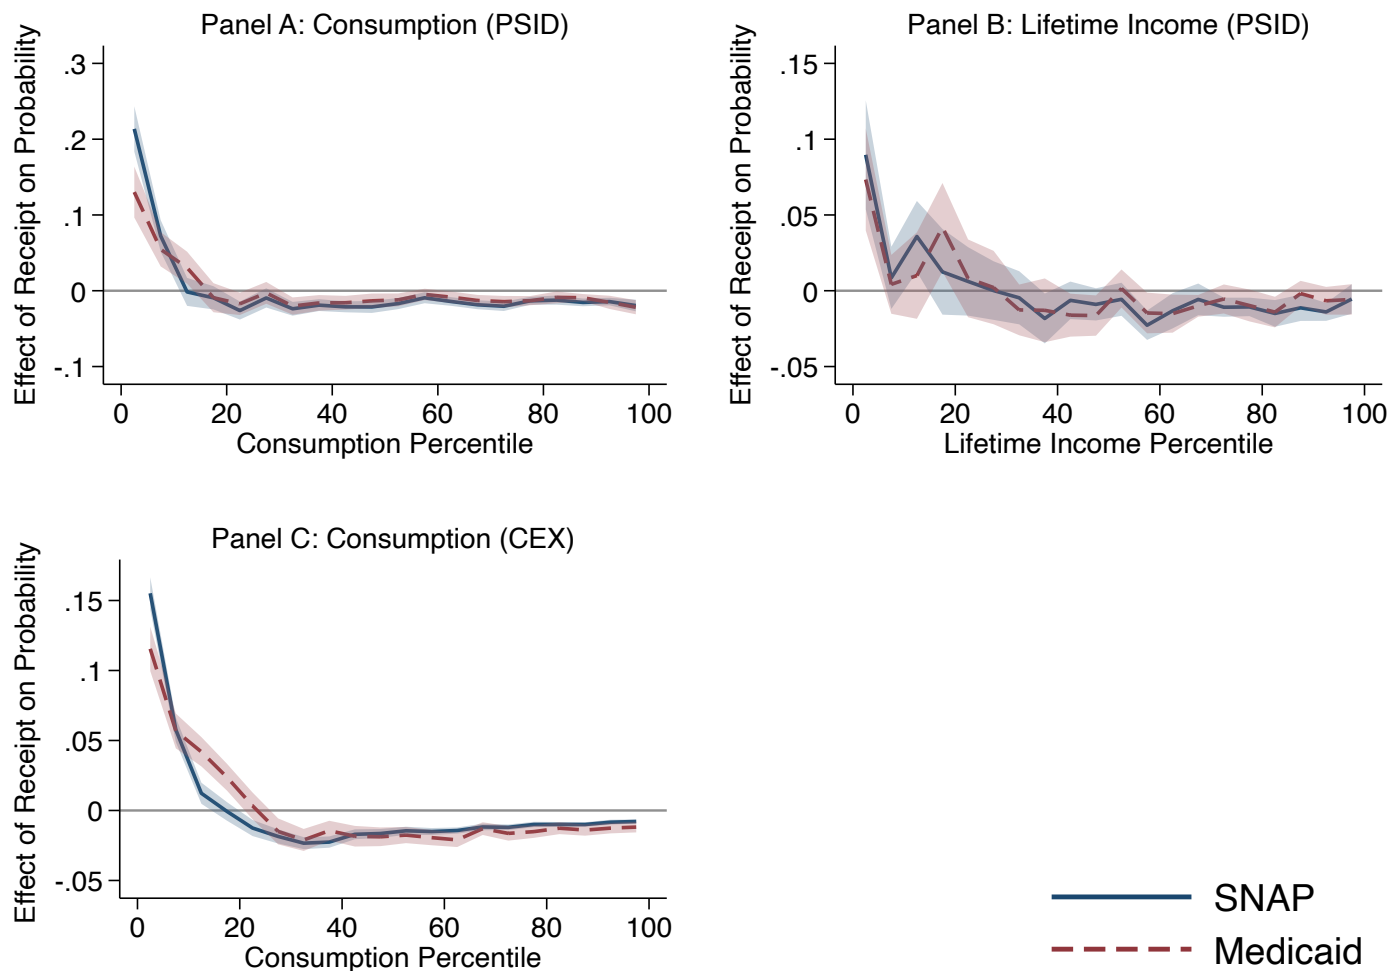
Notes: This figure displays estimates of the predictive effects of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). We split the sample by education (less than HS, HS degree, some college, BA only, more than BA), race/ethnicity (non-Hispanic white, non-Hispanic black, Hispanic, non-Hispanic other), and by “at-risk” group (single parent, disabled, non-native English speaker, car non-owner). The data are simulated-eligible members of the listed demographic group in the CEX. Panels A and B show results SNAP and Medicaid respectively. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household. The black “x” point indicates the population-average self-targeting coefficient and consumption rank.

Figure A14: Self-Targeting by Number of Distinct Transfers Received



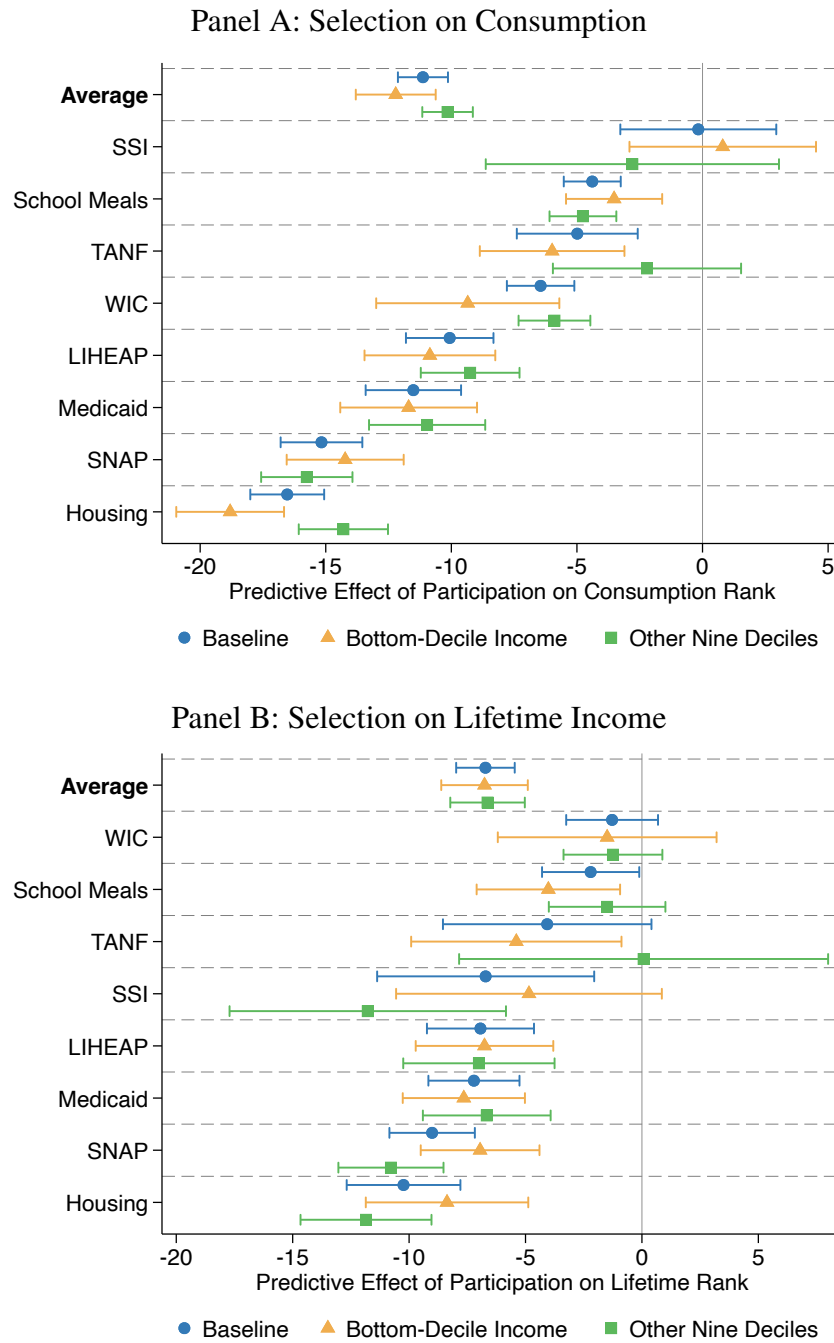
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4), distinguishing by the number of distinct transfers received. We report estimates using PSID consumption data (blue circles), CEX consumption data (yellow triangles), and PSID lifetime-income data (green squares). We also include results for receiving any transfer to the left of the dashed vertical line. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A15: Distributional Decompositions of Self-Targeting by Consumption Ventile




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 β from Equation 4). We estimate this specification using ventiles of consumption or lifetime-income ranks as outcomes—that is, each point is a coefficient from a separate regression with an outcome $1(\bar{R}_i \in [a, b])$. Confidence bands are pointwise, at the 95-percent level, and reflect clustered standard errors by household.

Figure A16: Heterogeneous Self-Targeting by Current Income Rank (PSID)



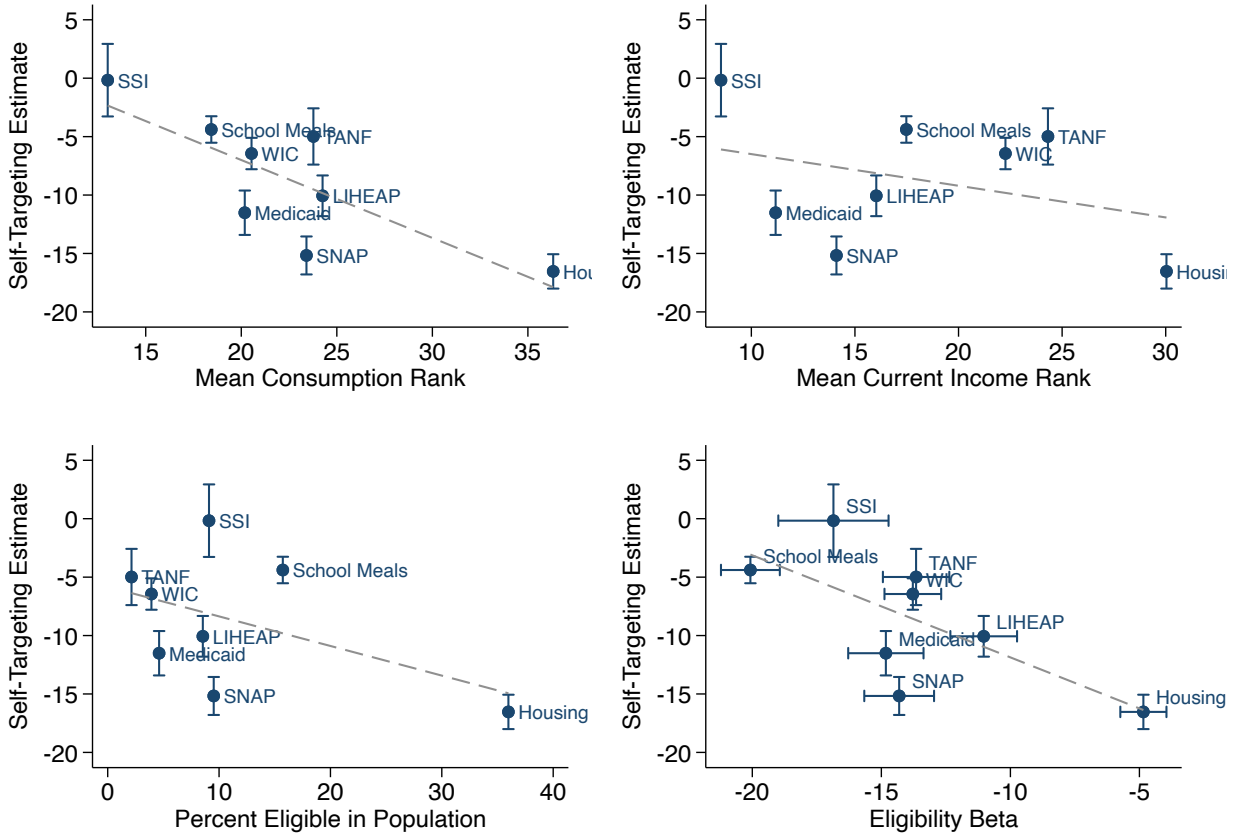
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). Estimates in blue circles repeat our main estimates in Figure 1. For estimates in yellow triangles or green squares, we respectively split the sample at the tenth percentile of the distribution of equivalized current household income. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A17: Heterogeneous Self-Targeting on Consumption by Current Income Rank (CEX)



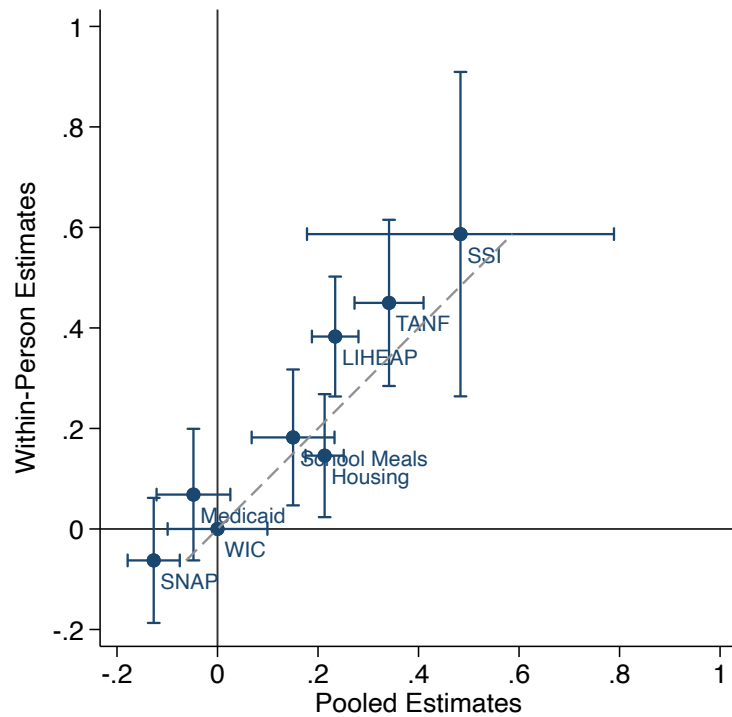
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4). Estimates in blue circles repeat our main estimates in Figure 1. For estimates in yellow triangles or green squares, we split the sample at the tenth percentile of the distribution of equivalized current household income. The data source is the Consumer Expenditure Survey. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A18: Relationship of Self-Targeting Estimates to Transfer Characteristics



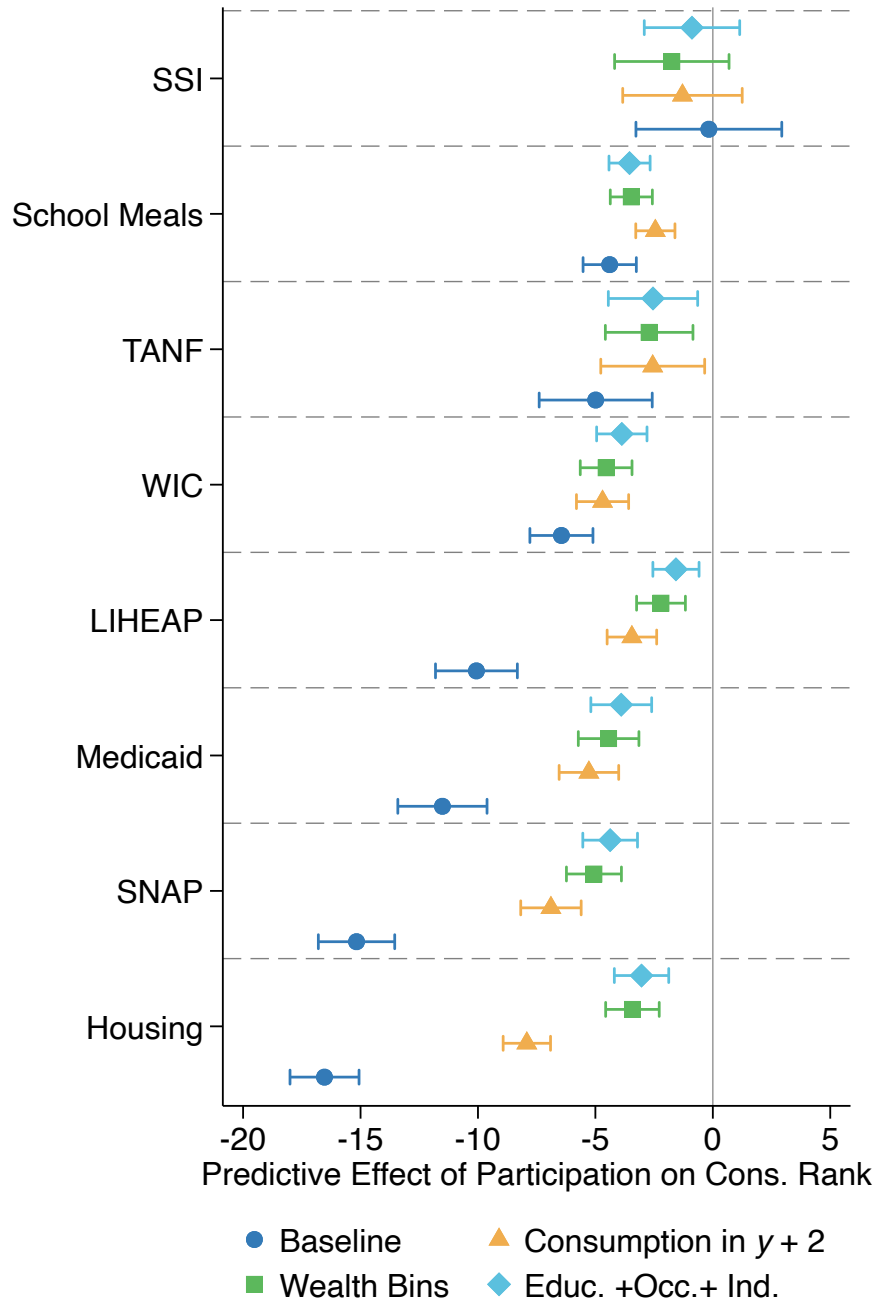
Notes: This figure presents four scatterplots that compare our self-targeting estimates (coefficient β from Equation 4) by transfer program to a program-level characteristic. We consider the mean consumption rank of transfer recipients (top left), the mean current-income rank of recipients (top right), the simulated-eligibility rate in the population (bottom left), and the selectivity of simulated eligibility rules (bottom right). The final characteristic is measured by the coefficient β in the following regression: $\bar{R}_i = \beta E_{it} + f(R_{it}) + u_{it}$, where E_{it} is simulated eligibility status. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A19: Variation in Eligible Populations Does Not Explain Across-Program Heterogeneity in Self-Targeting



Notes: This figure displays point estimates and confidence intervals for the following model of transfer take-up among the eligible: $D_{i,k} = \alpha_i + \beta_k \bar{R}_i + \gamma_k R_i + u_{i,k}$. In this expression, $D_{i,k}$ is a binary indicator of whether household i receives transfer k , α_i is a household fixed effect, \bar{R}_i is their consumption rank, and R_i is their current-income rank. We plot estimates and intervals of β_k for a specification with the fixed effects α_i (within-person, vertical axis) and without them (pooled, horizontal axis). The pooled estimates are computed relative to WIC, so as to be comparable to the within-person results, as these can only capture relative differences in self-targeting between transfers for the same household. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household. Pooled and within-person estimates, if exactly equal, would fall on the gray dashed 45-degree line.

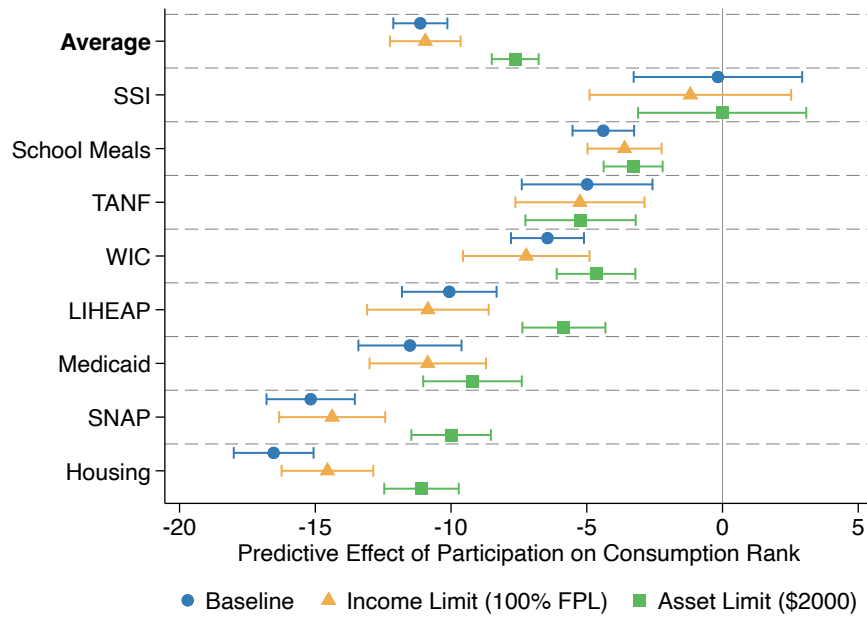
Figure A20: Relating Self-Targeting to Theories of Consumption Behavior



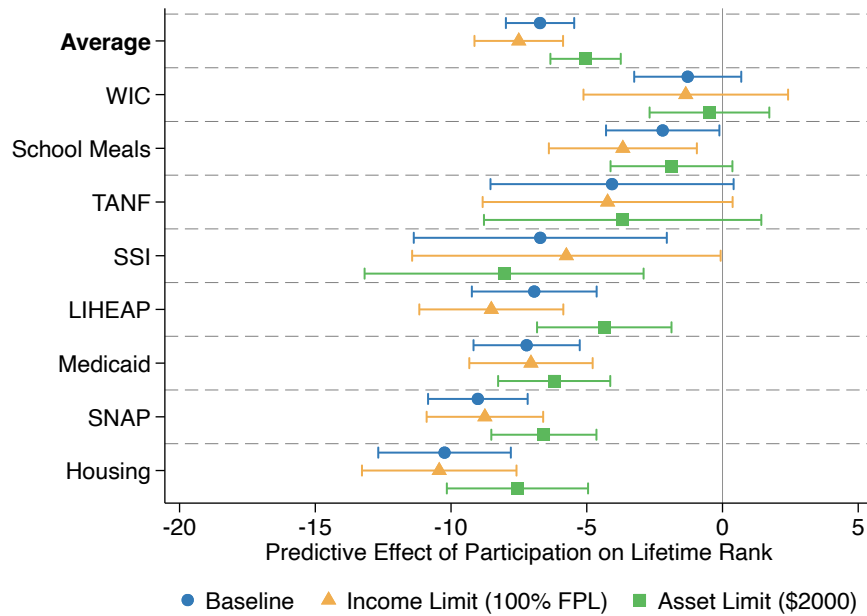
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4), but augmented with additional controls noted in the legend. Blue dots present our baseline estimates, yellow triangles include a control for the two-years-ahead consumption rank, green squares further include controls for wealth, and teal diamonds also include fixed effects for the person’s education, occupation, and industry. The wealth controls are indicators for the household’s decile of liquid savings, home equity, value of household automobiles, and other wealth. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A21: Adjustments to Eligibility Simulations (PSID)

Panel A: Selection on Consumption




Panel B: Selection on Lifetime Income



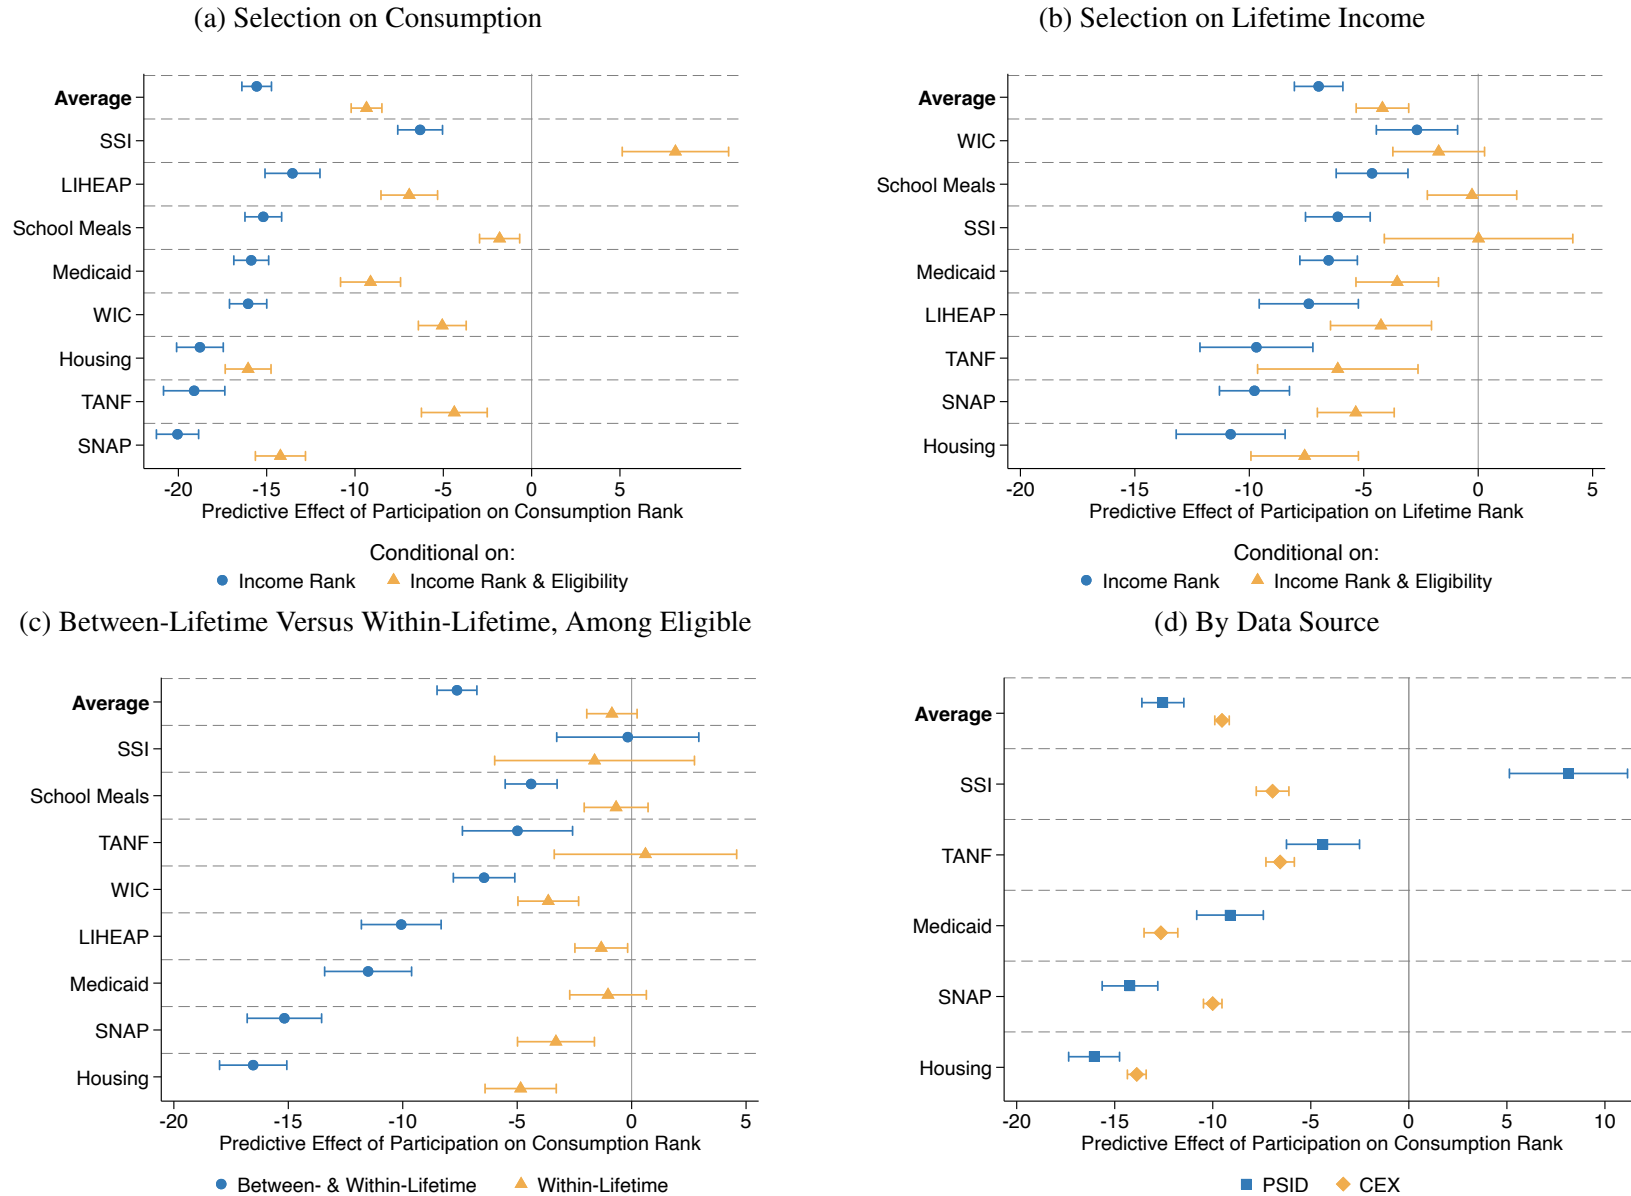
Notes: This figure displays estimates of the predictive effect of transfer benefit receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). For estimates represented by yellow triangles or green squares, we respectively impose an income limit (at 100% of the federal poverty level) or an asset limit (at \$2,000 in liquid assets, in 2020 dollars adjusted for the Consumer Price Index) on top of our eligibility simulations. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A22: Adjustments to Eligibility Simulations (CEX)



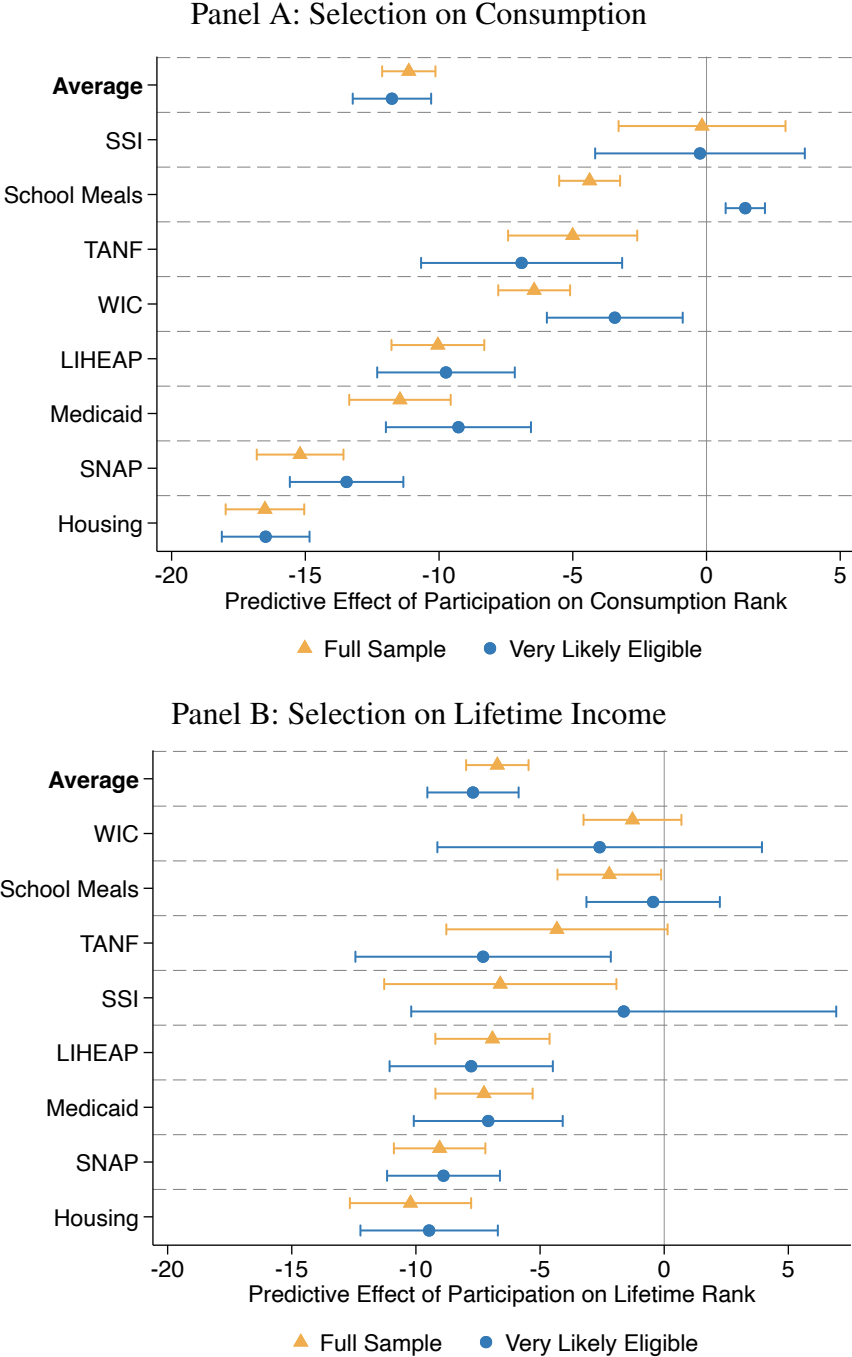
Notes: This figure displays estimates of the predictive effect of transfer benefit receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4). For estimates represented by yellow triangles or green squares, we respectively impose an income limit (at 100% of the federal poverty level) or an asset limit (at \$2,000 in liquid assets, in 2020 dollars adjusted for the Consumer Price Index) on top of our eligibility simulations. The data source is the Consumer Expenditure Survey. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A23: Self-Targeting in Transfer Programs: Reclassifying Simulated Ineligible Recipients




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 β from Equation 4). For the yellow diamonds, we estimate the regression only on people whom we simulate to be eligible. Panel C augments the specification with person-level fixed effects. The “any” row of Panels A–C is an indicator for receipt of at least one of the eight transfers. Panel D adapts Equation 4 by replacing the transfer indicator with indicators for the number of transfers received in that year. In all panels, 95-percent confidence intervals reflect clustered standard errors by household.

Figure A24: Self-Targeting Among the Very Likely Eligible (PSID)



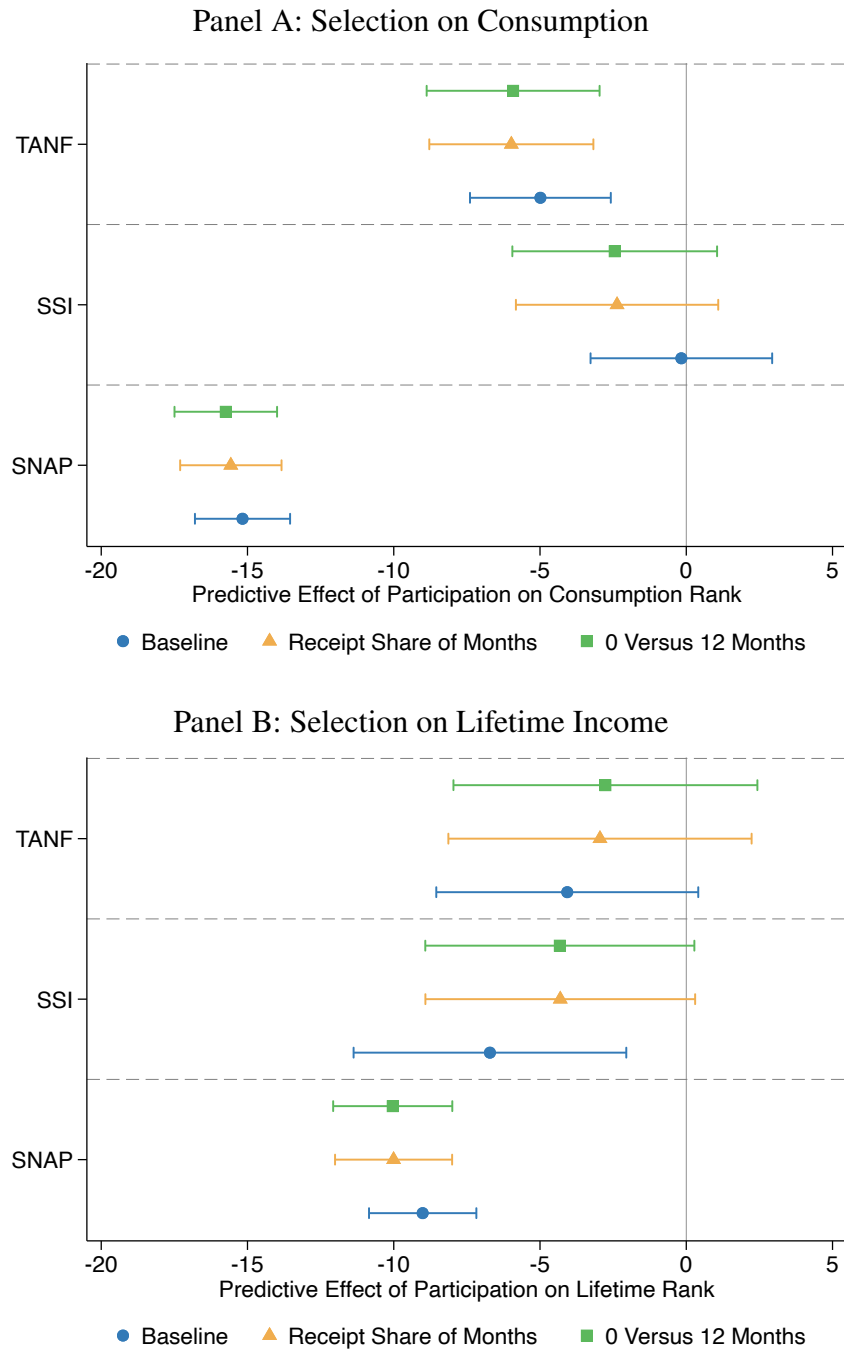
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). Both yellow diamonds and blue circles restrict the sample to simulated eligibles. For estimates represented 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 75 percent. The eligibility logit uses the following demographic variables: age (as a quadratic), 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, presence of a child in the household, income as a share of the federal poverty level, and rank-transformed current income, lifetime income, and consumption. The data source is the Panel Study of Income Dynamics. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A25: Self-Targeting on Consumption Among the Very Likely Eligible (CEX)



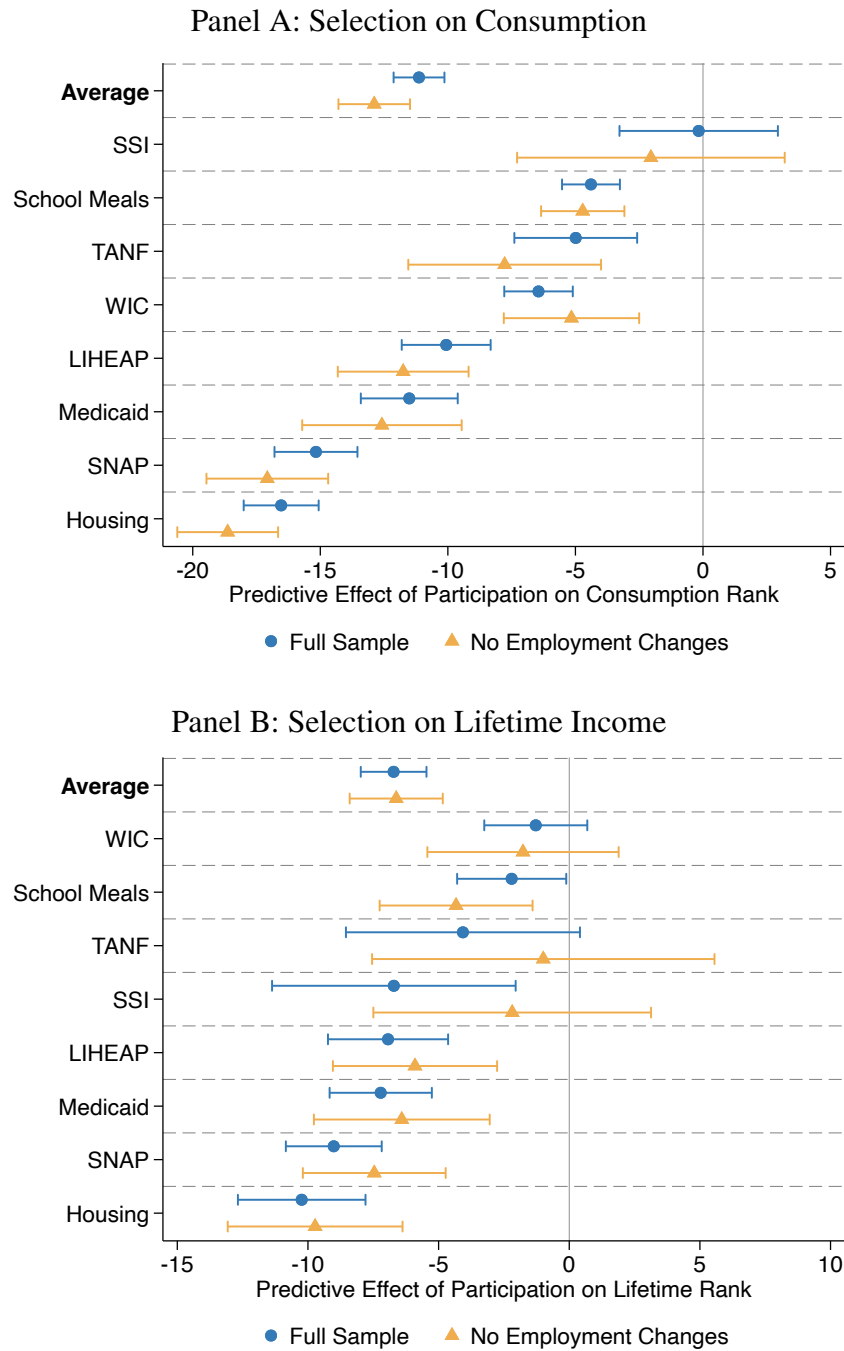
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank (coefficient β from Equation 4). Both yellow diamonds and blue circles restrict the sample to simulated eligibles. For estimates represented by blue circles, we further limit the sample to people who, in a logistic regression of simulated eligibility status on demographic observables, have a predicted probability of eligibility above 75 percent. The eligibility logit uses the following demographic variables: age (as a quadratic), 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, presence of a child in the household, income as a share of the federal poverty level, and rank-transformed current income, lifetime income, and consumption. The data source is the Consumer Expenditure Survey. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A26: Accounting for Months of Transfer Receipt



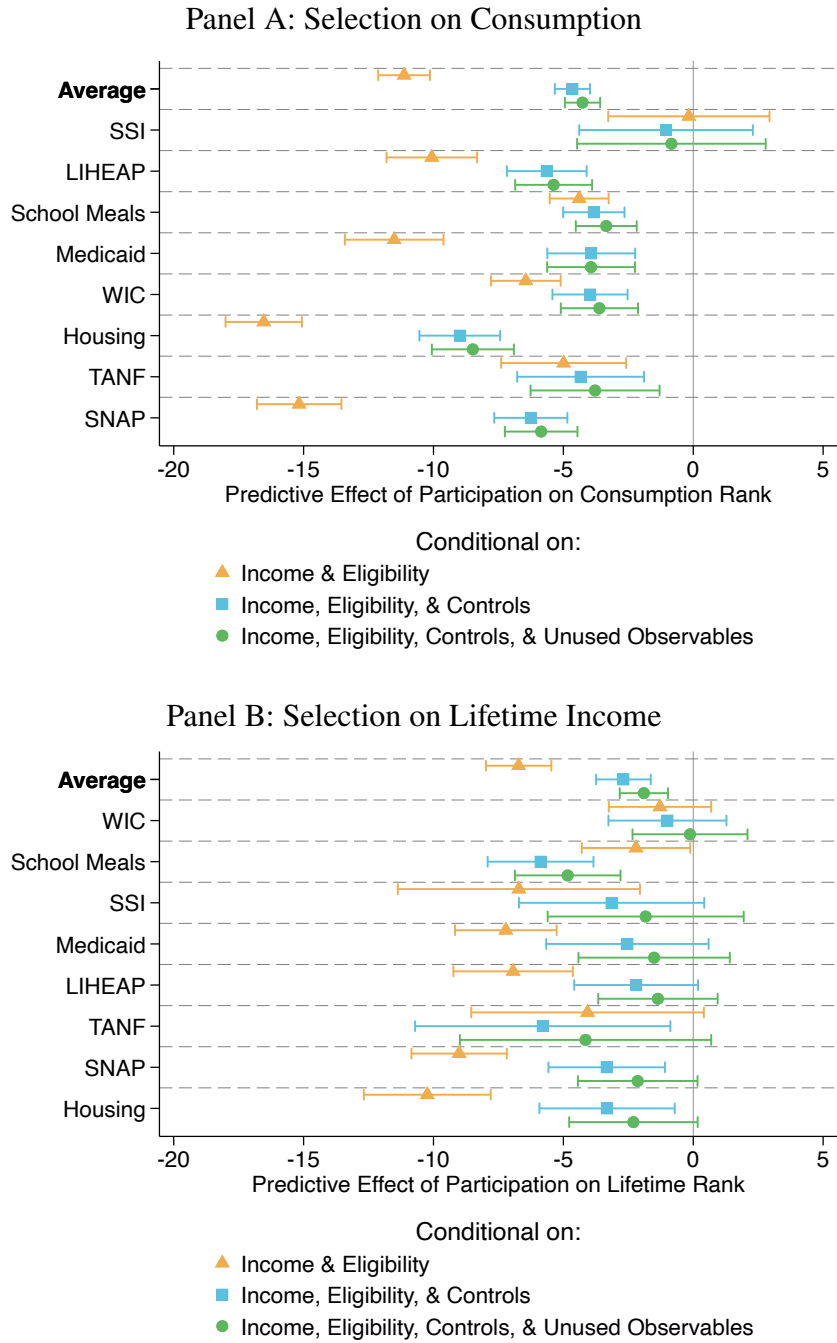
Notes: This figure displays estimates of the predictive effect of transfer benefit receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4). Estimates in blue circles repeat our main estimates in Figure 1. Estimates in yellow diamonds measure receipt as a share of months of the calendar year in which the household received the transfer, whereas the estimates in green squares drop households with partial-year receipt. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household. Partial-year recipient households respectively account for 4, 17, and 21 percent of receipt-months in SSI, SNAP, and TANF respectively.

Figure A27: Test for Aggregation Bias in Income Using Stable Employer Subsample




Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank or lifetime-income rank, conditional on current-income rank (coefficient β from Equation 4), distinguishing by the number of distinct transfers received. Estimates in blue circles reflect the full sample, whereas estimates in yellow diamonds are for the subsample of households with no change in employer in the last two years. We consider both the head and any spouse in determining employment stability, and include cases in which there is stable nonemployment. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A28: What Explains Selection into Transfer Receipt? With Controls



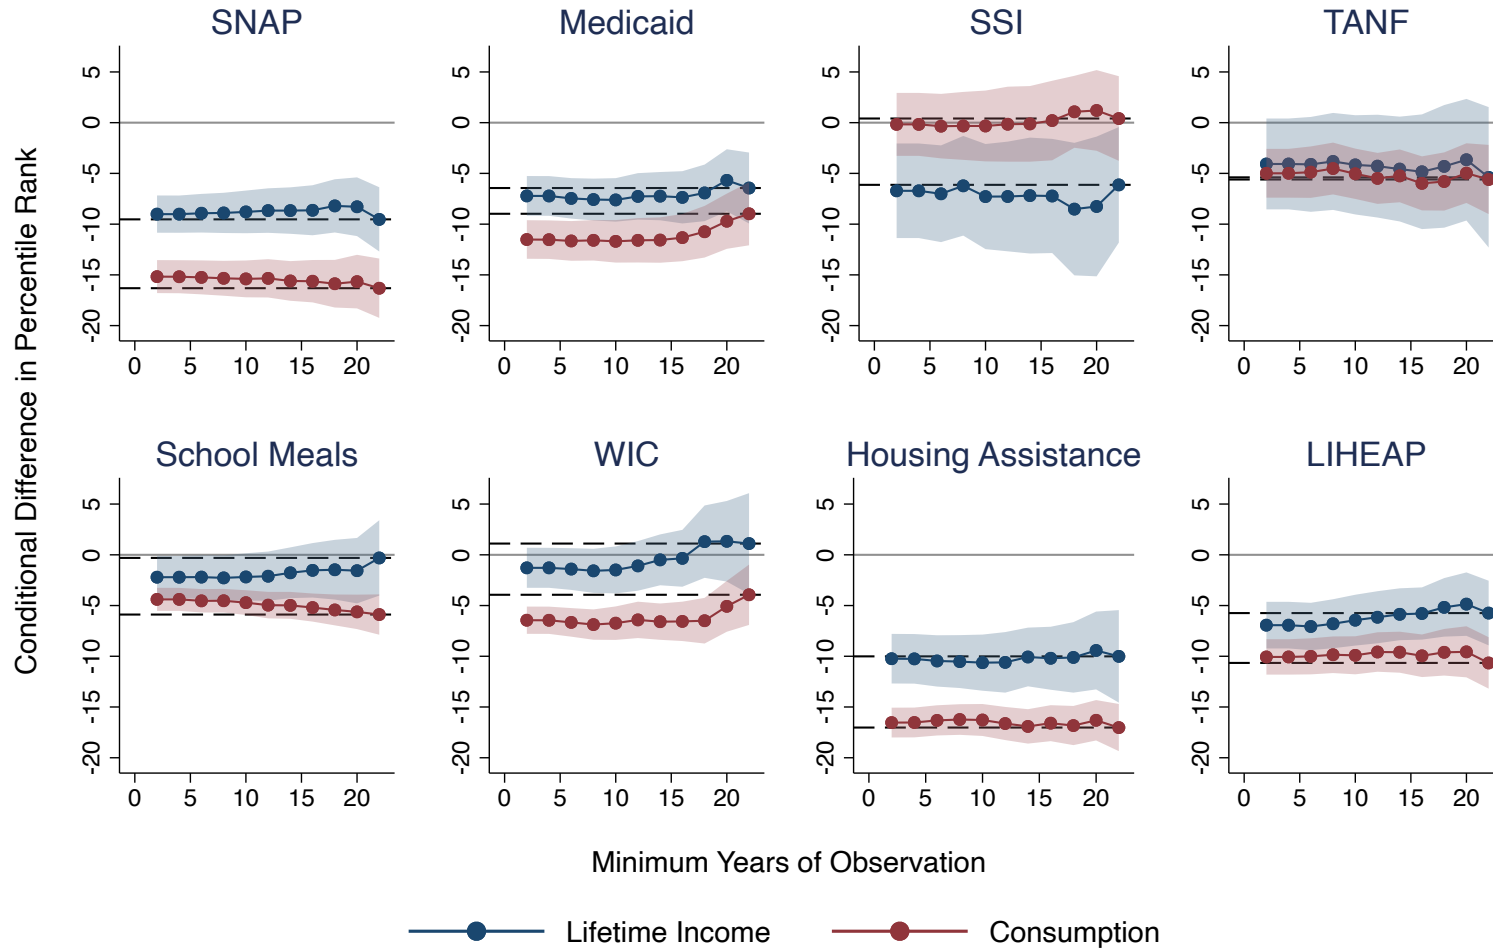
Notes: This figure displays estimates of the predictive effect of transfer receipt on consumption rank (Panel A) or lifetime-income rank (Panel B), conditional on current-income rank (coefficient β from Equation 4). For estimates represented by the yellow diamonds, we estimate the regression only on people whom we simulate to be eligible. For estimates represented by the teal squares, we condition on eligibility and characteristics that enter any eligibility rule. For estimates represented by green circles, we condition on eligibility, eligibility characteristics, and several demographic characteristics (race, education, and marital status). The “any” row is an indicator for receipt of at least one of the eight transfers. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A29: Selection into Transfers, with Predicted-Income Control




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 yellow diamonds, we include the predicted-income control. Both samples are limited to simulated-eligible people. The “any” row of Panel A is an indicator for receipt of at least one of the eight transfers. Confidence intervals are at the 95-percent level and reflect clustered standard errors by household.

Figure A30: Self-Targeting by Minimum Years of Observation



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 β from Equation 4). Moving to the right of each panel, we restrict the sample to eligible people with progressively more years of observation in the PSID. Dashed lines indicate our baseline estimates. Shaded regions reflect confidence intervals at the 95-percent level with clustering by household.

Figure A31: Selection into Transfer Receipt Over Time

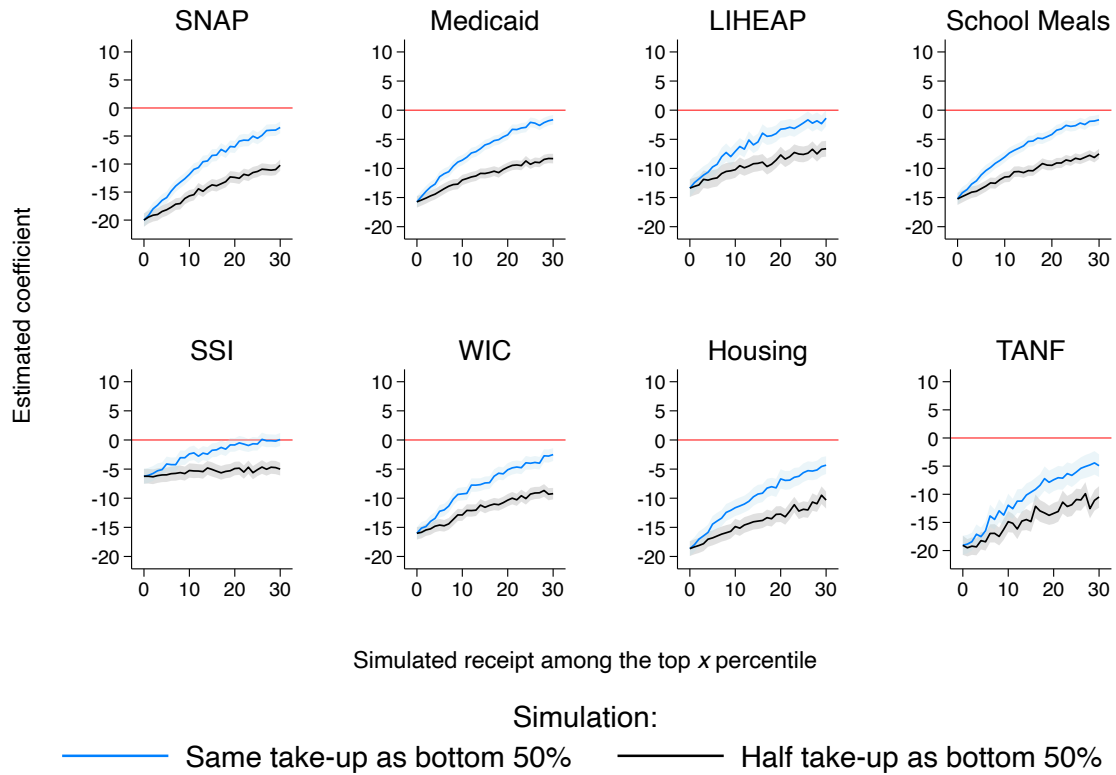


Notes: This figure displays coefficients from the following regression specification:

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

where i denotes households, t denotes years, and s denotes transfer programs. The outcome \bar{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 . Each sample is limited to that transfer's simulated-eligible population. 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 clustering by household. All regressions use PSID sample weights but are not otherwise adjust to account for variation in transfer program size.

Figure A32: 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 clustering by household.

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

Panel A: Receipt Rate

		Income Quintile					Avg.
		1	2	3	4	5	
Lifetime Income Quintile	1	44.5	17.3	4.0	0.9	0.4	35.3
	2	32.9	13.1	3.3	1.2	0.3	8.5
	3	26.0	9.7	3.0	0.7	0.3	3.4
	4	20.2	9.0	1.9	0.5	0.1	1.3
	5	18.8	6.0	2.2	0.7	0.2	0.9
	Avg.	33.9	11.7	2.9	0.7	0.2	

Panel B: Simulated Eligibility Rate

		Income Quintile					Avg.
		1	2	3	4	5	
Lifetime Income Quintile	1	66.0	14.3	2.5	1.2	0.4	29.5
	2	54.5	10.9	2.6	1.8	1.5	19.0
	3	50.9	9.7	1.7	0.9	0.9	10.2
	4	40.8	7.2	1.0	0.7	0.6	5.6
	5	33.4	5.3	0.9	0.8	0.3	4.2
	Avg.	55.2	10.1	1.7	1.0	0.6	

Panel C: Take-Up Rate Among Simulated Eligibles

		Income Quintile					Avg.
		1	2	3	4	5	
Lifetime Income Quintile	1	57.2	36.9	.	.	.	54.1
	2	45.9	30.5	.	.	.	41.5
	3	37.9	27.6	.	.	.	33.2
	4	33.8	17.6	.	.	.	27.5
	5	28.6	12.1	.	.	.	23.0
	Avg.	48.3	29.2	.	.	.	

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: SNAP Receipt, Eligibility, and Take-Up Rates by Income and Consumption Quintile:
Consumer Expenditure Survey

Panel A: Receipt Rate

		Income Quintile					Avg.
		1	2	3	4	5	
Consumption Quintile	1	40.2	16.0	5.8	2.5	2.2	28.5
	2	22.5	8.7	2.8	0.9	0.5	8.9
	3	13.4	4.6	1.8	0.7	0.4	3.1
	4	7.3	3.5	1.3	0.4	0.2	1.3
	5	3.7	1.4	0.7	0.3	0.1	0.4
	Avg.	30.3	9.0	2.2	0.6	0.2	

Panel B: Simulated Eligibility Rate

		Income Quintile					Avg.
		1	2	3	4	5	
Consumption Quintile	1	53.4	9.6	3.6	2.2	2.0	34.2
	2	42.8	6.0	2.2	1.1	0.8	12.4
	3	36.9	4.8	2.2	1.1	0.5	5.8
	4	35.9	4.0	1.8	1.0	0.7	3.4
	5	35.2	3.0	1.6	1.1	0.5	2.1
	Avg.	47.7	6.4	2.2	1.1	0.5	

Panel C: Take-Up Rate Among Simulated Eligibles

		Income Quintile					Avg.
		1	2	3	4	5	
Consumption Quintile	1	47.2	43.8	.	.	.	46.8
	2	28.1	30.9	.	.	.	28.1
	3	17.8	24.4	.	.	.	19.1
	4	10.3	17.8	.	.	.	11.9
	5	3.7	8.0	.	.	.	5.0
	Avg.	38.3	33.9	.	.	.	

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 equalized household consumption and income. Due to low rates of simulated eligibility, we do not report take-up rates for the top three income quintiles. The data source is the Consumer Expenditure Survey.

Table A3: Dollars and Percentage Differences Between Recipients and Similar-Income Non-Recipients

	Proportion Difference			Difference in 2020 Constant Dollars		
	Cons. (PSID) (1)	Cons. (CEX) (2)	Lifetime Inc. (PSID) (3)	Cons. (PSID) (4)	Cons. (CEX) (5)	Lifetime Inc. (PSID) (6)
SNAP	-0.549*** (0.039)	-0.497*** (0.013)	-0.913*** (0.155)	-8,253*** (639)	-11,607*** (326)	-17,797*** (3,886)
Medicaid	-0.491*** (0.044)	-0.485*** (0.021)	-0.761*** (0.113)	-6,617*** (668)	-11,849*** (558)	-14,369*** (2,549)
Housing Assistance	-0.451*** (0.035)	-0.479*** (0.016)	-0.474*** (0.175)	-7,476*** (618)	-15,876*** (530)	-14,527*** (5,356)
TANF	-0.342*** (0.095)	-0.434*** (0.027)	-0.200 (0.280)	-3,551*** (979)	-10,102*** (638)	-4,102 (5,722)
SSI	-0.004 (0.079)	-0.247*** (0.022)	-0.175 (0.256)	-36 (692)	-7,083*** (637)	-1,164 (1,657)
School Meals	-0.341*** (0.027)		-0.331** (0.167)	-3,774*** (325)		-5,552** (2,756)
WIC	-0.239*** (0.026)		-0.182* (0.108)	-2,570*** (298)		-4,319* (2,470)
LIHEAP	-0.374*** (0.037)		-0.515*** (0.142)	-5,298*** (594)		-9,569*** (2,669)

Notes: This table reports estimates of differences in consumption and lifetime income between transfer recipients 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 equalized current household income using cubic basis splines. Standard errors are clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A4: Self-Targeting on Additional Measures of Need

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)	School Meals (6)	WIC (7)	LIHEAP (8)
<i>Panel A: Log Wage (Mean: 1.742)</i>								
Receives Transfer	-0.089*** (0.022)	-0.042 (0.027)	-0.113*** (0.021)	-0.061 (0.047)	0.412*** (0.134)	-0.027 (0.019)	0.028 (0.019)	0.005 (0.029)
<i>Panel B: High-School Dropout (Mean: 0.186)</i>								
Receives Transfer	0.092*** (0.017)	0.053*** (0.016)	0.066*** (0.023)	0.048 (0.033)	-0.009 (0.053)	0.066*** (0.018)	0.023 (0.016)	0.058*** (0.020)
<i>Panel C: Single Parent (Mean: 0.147)</i>								
Receives Transfer	0.191*** (0.015)	0.165*** (0.014)	0.152*** (0.019)	0.195*** (0.031)	0.000 (0.000)	0.267*** (0.016)	0.009 (0.019)	0.078*** (0.019)
<i>Panel D: Disabled (Mean: 0.325)</i>								
Receives Transfer	0.105*** (0.015)	0.148*** (0.019)	0.049** (0.019)	0.050* (0.027)	0.200*** (0.036)	-0.123*** (0.013)	0.014 (0.010)	0.115*** (0.018)
<i>Panel E: Fair or Poor Health (Mean: 0.328)</i>								
Receives Transfer	0.090*** (0.016)	0.109*** (0.020)	0.016 (0.018)	0.035 (0.037)	0.078 (0.056)	-0.054*** (0.015)	0.014 (0.014)	0.098*** (0.021)
<i>Panel F: Nonwhite or Hispanic (Mean: 0.472)</i>								
Receives Transfer	0.126*** (0.024)	0.069*** (0.025)	0.257*** (0.029)	0.028 (0.038)	-0.053 (0.071)	0.060*** (0.021)	-0.015 (0.022)	-0.009 (0.032)
<i>Panel G: Has Savings (Mean: 0.539)</i>								
Receives Transfer	-0.149*** (0.019)	-0.009 (0.021)	-0.137*** (0.028)	-0.043 (0.036)	0.160** (0.064)	-0.035** (0.018)	-0.023 (0.020)	-0.017 (0.025)

Notes: Each panel row is for a different outcome, and each column is for a different transfer. Means are computed among the bottom quintile of the current-income distribution. All specifications control flexibly for current-income rank using cubic basis splines and condition on simulated eligibility. Standard errors are clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A5: Food Insecurity and SNAP Receipt (PSID)

	All				Simulated Eligible			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Received SNAP in Year	0.063*** (0.010)				0.057*** (0.015)			
Share of Year on SNAP		0.076*** (0.013)				0.062*** (0.018)		
Received SNAP in Month			0.072*** (0.012)	0.040*** (0.012)			0.059*** (0.016)	0.025 (0.018)
Current-Income Spline	Y	Y	Y	Y	Y	Y	Y	Y
Year-Month FE	Y	Y	Y	Y	Y	Y	Y	Y
Person-Year FE	N	N	N	Y	N	N	N	Y

Notes: This table reports estimates of the predictive effect of transfer receipt on a monthly measure of food insecurity, conditional on current-income rank. The dependent variable is a binary indicator for whether the household reported “difficulty in getting enough food to eat in your household” in a given month. Columns 1–4 are estimated on the full sample, whereas Columns 5–8 limit the sample to people that appear SNAP-eligible. Columns 1 and 5 compare food insecurity among SNAP recipients (people who received any SNAP transfer over the last year) to all and simulated-eligible non-recipients. Columns 2 and 6 change the definition of transfer receipt to the share of months in the year on SNAP. Columns 3 and 7 change the data frequency to monthly. Columns 4 and 8 then include for person–year fixed effects. All specifications include month–year fixed effects. Standard errors are clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A6: Selection into Transfers, Adjusted for Misreporting of Transfer Receipt

	Baseline		Adjusted for Misreporting	
	Consumption (1)	Lifetime Income (2)	Consumption (3)	Lifetime Income (4)
<i>Panel A: PSID</i>				
SNAP	-15.2*** (0.8)	-9.0*** (0.9)	-21.9*** (1.1)	-12.4*** (1.2)
Medicaid	-11.5*** (1.0)	-7.2*** (1.0)	-20.3*** (1.6)	-13.5*** (1.7)
<i>Panel B: CEX</i>				
SNAP	-11.2*** (0.3)		-14.3*** (0.4)	
Medicaid	-12.8*** (0.5)		-23.0*** (0.9)	

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. Panel A presents PSID results, and Panel B presents CEX results. Columns 1 and 2 reproduce our baseline estimates. 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 A7: Sensitivity of Self-Targeting on Consumption to Income Mismeasurement (PSID)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Unconstrained		Rank–Rank Slope Constrained to:						
	Self-Targeting (β)	Rank–Rank Slope (γ)	0.5	0.6	0.7	0.8	0.9	1	Preferred
SNAP	-15.02*** (0.82)	0.51*** (0.02)	-15.09*** (0.77)	-14.44*** (0.77)	-13.78*** (0.76)	-13.12*** (0.76)	-12.46*** (0.76)	-11.80*** (0.76)	-13.58*** (0.76)
Medicaid	-11.62*** (0.98)	0.40*** (0.04)	-11.41*** (0.95)	-11.19*** (0.95)	-10.98*** (0.95)	-10.76*** (0.95)	-10.54*** (0.95)	-10.33*** (0.96)	-11.25*** (0.95)
Housing Assistance	-15.59*** (0.72)	0.56*** (0.01)	-16.88*** (0.64)	-14.89*** (0.62)	-12.89*** (0.61)	-10.90*** (0.61)	-8.91*** (0.61)	-6.92*** (0.62)	-10.77*** (0.61)
TANF	-2.59** (1.12)	0.73*** (0.02)	-7.11*** (1.11)	-5.13*** (1.07)	-3.14*** (1.04)	-1.16 (1.03)	0.83 (1.04)	2.81*** (1.05)	2.81*** (1.05)
SSI	-0.40 (1.56)	0.23* (0.13)	0.07 (1.65)	0.25 (1.68)	0.43 (1.72)	0.61 (1.76)	0.79 (1.82)	0.97 (1.88)	-0.22 (1.61)
School Meals	-4.57*** (0.57)	0.45*** (0.02)	-4.57*** (0.57)	-4.57*** (0.56)	-4.57*** (0.56)	-4.57*** (0.56)	-4.58*** (0.57)	-4.58*** (0.58)	-4.57*** (0.56)
WIC	-6.54*** (0.69)	0.52*** (0.03)	-6.64*** (0.67)	-6.20*** (0.68)	-5.76*** (0.69)	-5.32*** (0.70)	-4.87*** (0.72)	-4.43*** (0.75)	-5.56*** (0.69)
LIHEAP	-9.71*** (0.89)	0.41*** (0.03)	-9.29*** (0.86)	-8.81*** (0.87)	-8.33*** (0.88)	-7.86*** (0.89)	-7.38*** (0.92)	-6.90*** (0.94)	-8.86*** (0.87)

Notes: This table reports estimates of the predictive effect of transfer receipt on consumption rank, conditional on current-income rank. We use the specification $\bar{R}_{it} = \beta D_{it} + \gamma R_{it} + u_{it}$. Columns 1 and 2 report the estimated coefficients on receipt and current-income rank. Columns 3–9 restrict γ to assess the sensitivity of β to income mismeasurement. The Preferred value comes from restricting the reliability ratio to 0.7, based on [Abowd and Stinson \(2013\)](#). We then constrain the coefficient on income to the unconstrained regression, divided by 0.7. The Preferred coefficient is capped at 1. Standard errors clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A8: Sensitivity of Self-Targeting on Lifetime Income to Income Mismeasurement (PSID)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Unconstrained		Rank–Rank Slope Constrained to:						
	Self-Targeting (β)	Rank–Rank Slope (γ)	0.5	0.6	0.7	0.8	0.9	1	Preferred
SNAP	-9.03*** (0.94)	0.32*** (0.03)	-7.82*** (0.93)	-7.16*** (0.92)	-6.50*** (0.93)	-5.85*** (0.93)	-5.19*** (0.93)	-4.53*** (0.94)	-8.13*** (0.93)
Medicaid	-7.01*** (1.01)	0.55*** (0.05)	-7.11*** (0.98)	-6.89*** (0.98)	-6.67*** (0.99)	-6.46*** (0.99)	-6.24*** (1.00)	-6.02*** (1.01)	-6.50*** (0.99)
Housing Assistance	-10.84*** (1.26)	0.38*** (0.01)	-8.48*** (1.19)	-6.49*** (1.18)	-4.50*** (1.17)	-2.50** (1.17)	-0.51 (1.18)	1.48 (1.19)	-7.59*** (1.18)
TANF	-5.97*** (2.29)	0.54*** (0.04)	-6.73*** (2.07)	-4.74** (2.05)	-2.76 (2.03)	-0.78 (2.02)	1.21 (2.01)	3.19 (2.02)	-1.39 (2.02)
SSI	-7.16*** (2.57)	0.46*** (0.10)	-7.08*** (2.62)	-6.90*** (2.61)	-6.72** (2.62)	-6.54** (2.62)	-6.36** (2.63)	-6.18** (2.65)	-6.80*** (2.61)
School Meals	-0.89 (1.13)	0.49*** (0.03)	-0.89 (1.13)	-0.89 (1.13)	-0.89 (1.12)	-0.89 (1.12)	-0.89 (1.12)	-0.89 (1.12)	-0.89 (1.12)
WIC	-1.13 (1.02)	0.53*** (0.04)	-1.26 (1.06)	-0.82 (1.06)	-0.38 (1.06)	0.06 (1.07)	0.50 (1.08)	0.94 (1.09)	-0.13 (1.07)
LIHEAP	-6.89*** (1.18)	0.47*** (0.04)	-6.73*** (1.20)	-6.26*** (1.20)	-5.78*** (1.21)	-5.30*** (1.21)	-4.82*** (1.23)	-4.34*** (1.24)	-5.93*** (1.20)

Notes: This table reports estimates of the predictive effect of transfer receipt on lifetime-income rank, conditional on current-income rank, using $\bar{R}_{it} = \beta D_{it} + \gamma R_{it} + u_{it}$. Columns 1 and 2 report the estimated coefficients on receipt and current-income rank. Columns 3–9 restrict γ to assess the sensitivity of β to income mismeasurement. The Preferred value comes from restricting the reliability ratio to 0.7, based on [Abowd and Stinson \(2013\)](#). We then constrain the coefficient on income to the unconstrained regression, divided by 0.7. The Preferred coefficient is capped at 1. Standard errors clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A9: Sensitivity of Self-Targeting on Consumption to Income Mismeasurement (CEX)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Unconstrained		Rank–Rank Slope Constrained to:						
	Self-Targeting (β)	Rank–Rank Slope (γ)	0.5	0.6	0.7	0.8	0.9	1	Preferred
SNAP	-12.03*** (0.27)	0.53*** (0.01)	-12.11*** (0.26)	-11.78*** (0.26)	-11.46*** (0.27)	-11.14*** (0.27)	-10.81*** (0.27)	-10.49*** (0.28)	-11.30*** (0.27)
Medicaid	-13.66*** (0.52)	0.61*** (0.02)	-14.72*** (0.49)	-13.77*** (0.48)	-12.82*** (0.48)	-11.87*** (0.48)	-10.91*** (0.49)	-9.96*** (0.50)	-11.17*** (0.49)
Housing Assistance	-14.25*** (0.25)	0.63*** (0.00)	-16.06*** (0.24)	-14.70*** (0.24)	-13.34*** (0.24)	-11.98*** (0.24)	-10.61*** (0.24)	-9.25*** (0.25)	-10.56*** (0.24)
TANF	-8.55*** (0.43)	0.68*** (0.01)	-12.78*** (0.41)	-10.45*** (0.40)	-8.11*** (0.40)	-5.77*** (0.40)	-3.43*** (0.40)	-1.09*** (0.41)	-1.72*** (0.41)
SSI	-8.16*** (0.54)	0.73*** (0.01)	-10.98*** (0.51)	-9.75*** (0.50)	-8.52*** (0.50)	-7.28*** (0.50)	-6.05*** (0.50)	-4.82*** (0.50)	-4.82*** (0.50)

Notes: This table reports estimates of the predictive effect of transfer receipt on lifetime-income rank, conditional on current-income rank. We use here the following linear specification: $\bar{R}_{it} = \beta D_{it} + \gamma R_{it} + u_{it}$, where \bar{R}_{it} is consumption, D_{it} is transfer receipt, and R_{it} is current-income rank. Columns 1 and 2 respectively report the estimated coefficients on receipt and current-income rank. Columns 3–9 restrict γ to assess the sensitivity of β to income mismeasurement. The Preferred value comes from restricting the reliability ratio to 0.7, based on [Abowd and Stinson \(2013\)](#). We then constrain the coefficient on income to the unconstrained regression, divided by 0.7. The Preferred coefficient is capped at 1. Standard errors are clustered by household. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A10: Well-Measured Consumption and Transfer Receipt (PSID)

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)	School Meals (6)	WIC (7)	LIHEAP (8)
<i>Panel A: Rent and Owner's Equivalent Rent (1997–2019)</i>								
Receives Transfer	-0.496*** (0.036)	-0.391*** (0.035)	-0.820*** (0.037)	-0.378*** (0.078)	0.074 (0.105)	-0.224*** (0.030)	-0.190*** (0.028)	-0.272*** (0.043)
<i>Panel B: Vehicle Lease Cost and Equivalent Lease Cost (1999–2019)</i>								
Receives Transfer	-0.310*** (0.020)	-0.329*** (0.024)	-0.178*** (0.025)	-0.026 (0.038)	-0.012 (0.067)	-0.254*** (0.015)	-0.082*** (0.019)	-0.155*** (0.026)
<i>Panel C: Food at Home Expenditure (1999–2019)</i>								
Receives Transfer	-0.600*** (0.025)	-0.328*** (0.029)	-0.239*** (0.031)	-0.465*** (0.064)	-0.065 (0.103)	-0.052** (0.025)	-0.282*** (0.029)	-0.298*** (0.034)
<i>Panel D: Utility Expenditure (1999–2019)</i>								
Receives Transfer	-0.094*** (0.030)	-0.170*** (0.034)	-0.338*** (0.048)	-0.169** (0.067)	0.145** (0.069)	-0.155*** (0.024)	-0.113*** (0.026)	0.018 (0.032)
<i>Panel E: Gasoline Expenditure (1999–2019)</i>								
Receives Transfer	-0.116*** (0.032)	-0.199*** (0.038)	-0.128*** (0.039)	-0.097 (0.072)	-0.289** (0.113)	-0.097*** (0.026)	-0.054* (0.029)	-0.177*** (0.041)

Notes: This table reports estimates of the predictive effect of transfer receipt on (log) levels of reported consumption, conditional on current-income rank and being simulated-eligible for the transfer. 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 A11: Well-Measured Consumption and Transfer Receipt (CEX)

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)
<i>Panel A: Rent and Owner's Equivalent Rent</i>					
Receives Transfer	-0.554*** (0.016)	-0.439*** (0.026)	-0.659*** (0.021)	-0.314*** (0.032)	-0.197*** (0.027)
<i>Panel B: Vehicle Lease Cost and Equivalent Lease Cost</i>					
Receives Transfer	-0.136*** (0.016)	-0.182*** (0.022)	-0.308*** (0.019)	-0.270*** (0.040)	-0.109*** (0.030)
<i>Panel C: Food at Home Expenditure</i>					
Receives Transfer	-0.113*** (0.009)	-0.167*** (0.014)	-0.051*** (0.010)	0.026 (0.017)	-0.073*** (0.015)
<i>Panel D: Utility Expenditure</i>					
Receives Transfer	-0.134*** (0.012)	-0.258*** (0.018)	-0.228*** (0.014)	-0.235*** (0.026)	-0.104*** (0.018)
<i>Panel E: Gasoline Expenditure</i>					
Receives Transfer	-0.512*** (0.020)	-0.413*** (0.029)	-0.535*** (0.021)	-0.596*** (0.050)	-0.315*** (0.032)

Notes: This table reports estimates of the predictive effect of transfer receipt on (log) levels of reported consumption, conditional on current-income rank and being simulated-eligible for the transfer. The data source is the Consumer Expenditure Survey, covering the years 1997 through 2019 for all outcomes. 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 A12: Durable-Goods Ownership and Transfer Receipt (PSID)

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)	School Meals (6)	WIC (7)	LIHEAP (8)
<i>Panel A: HH Owns Primary Residence (1997–2019)</i>								
Receives Transfer	-0.129*** (0.016)	-0.074*** (0.018)	-0.382*** (0.012)	-0.158*** (0.027)	0.091 (0.060)	0.035* (0.019)	0.007 (0.017)	-0.040 (0.026)
<i>Panel B: Number of Rooms in Home (1997–2019)</i>								
Receives Transfer	-0.411*** (0.067)	-0.238*** (0.075)	-0.711*** (0.066)	-0.151 (0.152)	0.111 (0.167)	-0.165*** (0.063)	-0.430*** (0.073)	0.132* (0.079)
<i>Panel C: Central Air Conditioning at Home (1997–2009)</i>								
Receives Transfer	-0.014 (0.021)	0.014 (0.025)	0.015 (0.029)	-0.165*** (0.032)	-0.026 (0.073)	-0.005 (0.019)	-0.052** (0.026)	-0.025 (0.027)
<i>Panel D: HH Owns a Car (1999–2019)</i>								
Receives Transfer	-0.083*** (0.020)	-0.024 (0.021)	-0.190*** (0.023)	-0.080** (0.037)	0.026 (0.055)	0.098*** (0.017)	0.036* (0.020)	-0.019 (0.021)
<i>Panel D: HH Owns a Computer (2003–2019)</i>								
Receives Transfer	-0.088*** (0.020)	0.025 (0.024)	-0.080*** (0.024)	-0.069 (0.047)	0.093* (0.054)	0.068*** (0.018)	-0.042* (0.025)	-0.040* (0.023)
<i>Panel E: HH Owns a Smartphone (2015–2019)</i>								
Receives Transfer	-0.005 (0.024)	0.064** (0.028)	-0.013 (0.031)	-0.018 (0.073)	-0.033 (0.089)	0.088*** (0.019)	0.031* (0.017)	-0.060 (0.038)

Notes: This table reports estimates of the predictive effect of transfer receipt on measures of household durable-goods ownership, conditional on current-income rank and being simulated-eligible for the transfer. 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 A13: Durable-Goods Ownership and Transfer Receipt (CEX)

	SNAP (1)	Medicaid (2)	Housing Assistance (3)	TANF (4)	SSI (5)
<i>Panel A: HH Owns Primary Residence</i>					
Receives Transfer	-0.131*** (0.007)	-0.153*** (0.011)	-0.294*** (0.005)	-0.203*** (0.013)	-0.133*** (0.013)
<i>Panel B: Number of Rooms in Home</i>					
Receives Transfer	-0.383*** (0.026)	-0.479*** (0.043)	-0.712*** (0.024)	-0.354*** (0.056)	-0.285*** (0.048)
<i>Panel C: Central Air Conditioning at Home</i>					
Receives Transfer	-0.110*** (0.007)	-0.067*** (0.011)	0.033*** (0.008)	-0.168*** (0.016)	-0.082*** (0.013)
<i>Panel D: HH Owns a Car</i>					
Receives Transfer	-0.132*** (0.007)	-0.093*** (0.011)	-0.210*** (0.008)	-0.276*** (0.017)	-0.125*** (0.013)
<i>Panel D: HH Owns a Computer</i>					
Receives Transfer	-0.162*** (0.007)	-0.008 (0.011)	-0.011 (0.008)	-0.098*** (0.016)	-0.042*** (0.012)

Notes: This table reports estimates of the predictive effect of transfer receipt on measures of household durable-goods ownership, conditional on current-income rank and being simulated-eligible for the transfer. The data source is the Consumer Expenditure Survey, covering the years 1997 through 2019 for all outcomes. 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 A14: Simulated Eligibility and Receipt Rates

Program	P(Receive Simulated Eligible) (1)	P(Receive Not Sim. Elig.) (2)	P(Sim. Elig. Receive) (3)	P(Sim. Elig. Not Receive) (4)
SNAP	0.43	0.05	0.60	0.09
Medicaid	0.52	0.11	0.32	0.05
Housing Assistance	0.13	0.02	0.73	0.27
TANF	0.12	0.01	0.45	0.03
SSI	0.62	0.05	0.08	0.00
School Meals	0.46	0.05	0.59	0.08
WIC	0.47	0.02	0.48	0.02
LIHEAP	0.18	0.02	0.55	0.10

Notes: This table reports conditional receipt and eligibility rates by program, using simulated eligibility measures constructed from observed household characteristics. For each program, we estimate the share of recipients who are eligible, the share of non-recipients who are eligible, the share of eligibles who receive the transfer, and the share of non-eligibles who do. Estimates are based on weighted linear regressions with no covariates. See Appendix B for details on the construction of simulated eligibility.

Table A15: Explaining Differences in Take-Up Rates

	PSID	Official
Receipt Rate	10.9%	7% – 17%
Sim. Elig. Rate	14.4%	14% – 17%
Take-Up		
Ratio of Rates	75.7%	48% – 84%
Among Sim. Elig.	43.4%	n.a.

Notes: This table reports household rates of receipt, simulated eligibility, and take-up for SNAP. PSID estimates reflect an average over PSID samples for 1997–2019. Official estimates are ranges over annual averages for the same years in [U.S. Department of Agriculture \(2022b\)](#).

Table A16: Review of Estimated Take-Up Rates

Program	PSID	Benchmarks	References
SNAP	43.4%	48% – 84%	Newman and Scherpf (2013), U.S. Department of Agriculture (2022b)
Medicaid	54.0%	52% – 70%	Sommers et al. (2012)
Housing Assistance	10.0%	10%, 18.2%	Olsen (2003), Congressional Research Service (2015)
TANF	12.3%	28.4%	Congressional Research Service (2015)
SSI	66.5%	57.5%, 66.6%	McGarry and Schoeni (2015), Congressional Research Service (2015)
WIC	40.9%	39.1% – 65.3%	Guan et al. (2023), McBride et al. (2023), Congressional Research Service (2015), U.S. Department of Agriculture (2022a)
LIHEAP	17.0%	22.2%	Congressional Research Service (2015)

Notes: This table reports estimates of take-up rates from other research.

Table A17: Marginal Value of Public Funds (MVPF), Voluntary and Automatic Transfers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Dollar-Wt. Avg.	SNAP	Medicaid	Housing Assistance	TANF	SSI	School Meals	WIC	LIHEAP
Social Benefits									
Within-Inc. Targeting	0.31	0.44	0.32	0.25	0.08	-0.04	0.09	0.18	0.15
Across-Inc. Targeting	0.08	0.12	0.05	0.10	0.17	-0.01	-0.02	0.08	0.05
Social Costs									
Ordeal	0.16	0.17	0.21	0.06	0.04	0.25	0.18	0.19	0.07
Labor Supply	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.00
Take-Up Rate	0.40	0.43	0.52	0.15	0.10	0.63	0.46	0.48	0.19
MVPF, Automatic	1.27	1.26	1.08	1.57	2.53	0.97	0.94	1.13	1.21
MVPF, Voluntary	1.60	1.64	1.22	2.35	2.51	0.66	0.81	1.09	1.46

Notes: This table presents the Marginal Value of Public Funds (MVPF) for various transfer programs, distinguishing between automatic and voluntary transfers. The values provide all inputs to the MVPFs: within-income and across-income targeting benefits, social costs of ordeals, fiscal externalities from labor supply effects, and take-up rates. Social benefits, social costs, and MVPFs are all in units of dollars per transfer dollar.

B Data Appendix

This appendix first explains measurement details for consumption, lifetime income, and simulated eligibility. Next, it discusses the estimation of take-up rates.

B.1 Consumption Ranks (PSID and CEX)

PSID. 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.³⁵

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 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. House-

³⁵For 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.

holds 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.³⁶

CEX. We adjust the expenditure variables (totexppq and totexpcq) by replacing purchases of vehicles and homes with imputed consumption flows. We also remove mandatory contributions to pensions from the expenditure totals, following [Chetty and Szeidl \(2007\)](#).

For homes, we subtract off outlays on own dwellings (owndwepq and owndwecq) and replace it with self-reported rental equivalents (rnteqvx). For vehicles, we subtract off outlays on cars and trucks, both new and used (cartkncq, cartkucq, cartknpq, cartkupq) as well as vehicle finance charges (vehfincq and vehfinpq). We replace these with predicted values of lease-cost equivalents from a hedonic regression. We estimate this a Poisson regression and includes as covariates the vehicle model, its age (model year and calendar year), its mileage, and its “type” (e.g., convertible, see vehtype). Vehicle age enters as year fixed effects and, specific to each model, a linear slope in age. We use reported lease costs for leased vehicles.

Several variables used in our eligibility simulators for the PSID are not available in the CEX. These are information about citizenship and immigration history, the disability status of children, and data about asset values (home, car, and other). For the purpose of eligibility simulation, we assume all people are citizens, to avoid impacts on transfer eligibility. Similarly, for vehicle values, we assume a value of \$2,000 for any used car and \$5,000 for any new car, and also we assign zero value to any reported owned home.

B.2 Lifetime Income Ranks (PSID)

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:

$$E[y_{it} | X_{it}] = \exp(\alpha_i \lambda_a + X'_{it} \beta_a), \quad (12)$$

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

³⁶For missing values, we impute using the cross-sectional relationship between fitted purchase prices and fitted lease values from these Poisson regressions.

outcome y_{it} is individual income. For individuals with zero income in all observed years, we impute a constant annual income of \$100.³⁷ 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 $\hat{\alpha}_i$ toward a conditional expectation fit from several time-invariant individual characteristics.³⁸ 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 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:

$$E[y_{it} | X_{it}] = \exp(\hat{\alpha}_i \lambda_a + X'_{it} \beta_a), \quad (13)$$

importantly treating $\{\hat{\alpha}_i\}$ as data rather than as parameters. We then obtain coefficient estimates $\{\hat{\lambda}_a\}$, and with these in hand, we return to step 1 and iterate until convergence of $\{\hat{\alpha}_i, \hat{\lambda}_a, \hat{\beta}_a\}$. 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 $\hat{y}_{it} = \exp(\hat{\alpha}_i^* \hat{\lambda}_a + X'_{it} \hat{\beta}_a)$, where $\hat{\alpha}_i^*$ are the shrunken estimates of the individual fixed effects. Lifetime income are then

$$\bar{y}_i = \sum_t \hat{y}_{it}, \quad (14)$$

where the summation over t is for the years $\{T_i, \dots, \bar{T}_i\}$ 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

³⁷In 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.

³⁸Recent 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.

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 $\Pr(y \leq \bar{y}_i | c_i = c)$, where \bar{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 $\Pr(y \leq y_{it} | c_i = c)$, 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

$$\bar{y}_i^h = \sum_t e(\hat{y}_{it}^h) = \sum_t e(\hat{y}_{it} + \hat{y}_{j(i,t),t}) \quad (15)$$

where t is again summed over the years in which i is between ages 18 and 65. The function $e(\cdot)$ equalizes 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 (PSID and CEX)

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.³⁹ 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

³⁹See <https://www.ers.usda.gov/data-products/snap-policy-data-sets/>.

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.⁴⁰

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. We include the value of household vehicles, exceeding \$4,650, in the asset test where the SNAP Policy Database indicates the state applies no exclusion rules. However, we do not incorporate specific state-level vehicle exclusion rules due to limited data coverage. 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 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.

Special eligibility rules also apply to non-citizen immigrants according to their date of arrival and current immigration status, pursuant to the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) and subsequent state and federal laws. We code these from [Zimmermann and Tumlin \(1999\)](#) and capture subsequent updates in an April 2022 report by the Food Research and Access Center, “State Food Assistance Programs: Addressing Gaps in SNAP Eligibility for Immigrants.”

Medicaid. Medicaid eligibility is determined by income and asset tests that vary by state and with household characteristics. After 2013, interstate variation in eligibility follows largely from the Affordable Care Act’s Medicaid expansion. Before 2013, interstate variation is driven by state use of Medicaid waiver programs. In most states, SSI recipients are categorically Medicaid-eligible; we also 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

⁴⁰For further details, see Center on Budget and Policy Priorities, “A Quick Guide to SNAP Eligibility and Benefits.”

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.

Special eligibility rules also apply to non-citizen immigrants according to their date of arrival and current immigration status, pursuant to the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) and subsequent state and federal laws. We code these using [Zimmermann and Tumlin \(1999\)](#), as well as periodic reports of the National Immigration Law Center and the U.S. Department of Health and Human Services (Assistant Secretary for Planning and Evaluation). We do not incorporate work requirements or time limits on SNAP for able-bodied adults without dependents (ABAWDs) or state waivers for these eligibility rules.

Housing Assistance. Eligibility for housing assistance (Section 8 and public housing) is determined by income and household size. Income is measured relative to Area Median Income (AMI) at the level of metropolitan area or non-metropolitan county, further incorporating an adjustment for household size. 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.

Special eligibility rules also apply to non-citizen immigrants according to their date of arrival and current immigration status. Approximating PRWORA, we code a household as ineligible if it contains no citizens or “qualified” non-citizen immigrants (permanent residents, asylees, refugees, and Temporary Protected Status).

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.

Special eligibility rules also apply to non-citizen immigrants according to their date of arrival and current immigration status, due to PRWORA and subsequent state and federal laws. We code these using [Zimmermann and Tumlin \(1999\)](#), as well as periodic reports of the National Immigration Law Center.

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, either in proportion to the state median income or the HHS federal poverty guideline, which we convert into dollars. Income thresholds 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, which is the median limit for states reporting this value. 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.

Special eligibility rules also apply to non-citizen immigrants according to their date of arrival and current immigration status. Approximating PRWORA, we code a person as ineligible if they are neither a citizen nor a "qualified" non-citizen immigrant (permanent residents, asylees, refugees, and Temporary Protected Status).

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.⁴¹

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.⁴² We implement three tests: a gross income test, a net income test, and an asset test. The gross and net income tests vary at the state level.

Tests. From TRIM3, we obtain a household-size-by-state-by-year level gross and net income tests, the thresholds for which are constructed via dollar values times a scalar factor. We convert annual family income excluding TANF into a monthly amount and test whether this value is less than the gross or net test threshold. In some states, the net test also applies an overall earnings disregard, which we obtain from TRIM3. Some states apply only a gross or a net test; we use the appropriate tests as indicated by TRIM3. In state-years where TRIM3 does not record a gross or net test, we assume they have the median gross eligibility threshold for that year.

TRIM3 indicates that rules in 2002 were based on a TANF "standard of need" (a different dollar threshold), which we apply.

We test whether the household's liquid savings exceeds the state-by-year asset tests. This ignores state-specific inclusions of stocks, retirement accounts or other assets. We fully exempt the value of the vehicle and illiquid assets, although rules differ in some states.

Non-Citizens. Special eligibility rules also apply to non-citizen immigrants according to their date of arrival and current immigration status, pursuant to the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA) and subsequent state and federal laws. We

⁴¹See 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.

⁴²<https://boreas.urban.org/documentation/TANF/Main.php>.

code these using [Zimmermann and Tumlin \(1999\)](#) and a 2014 report for the U.S. Department of Health and Human Services (Assistant Secretary for Planning and Evaluation).

Other Assumptions. We restrict TANF eligibility to households with children who do not report TANF benefits in more than two PSID observations before that year. The official TANF limit is 60 months, so this time-limits assumption is accurate if one observation in the PSID (which is collected every-other-year) indicates the household was on TANF for a continuous 24 months. A limitation is that this test is differentially binding in our data for the first few years of PSID observations than subsequent years.

We censor household size at seven. In 1997, we do not have asset data available in the PSID, so we do not apply the asset test in that year. We do not apply specific earnings deductions for work expenses or child care. We do not include work histories or statuses in eligibility computations.

TRIM3 does not include years prior to 2002, so we assume earlier years are the same as 2002.

For net and gross income tests, less than five percent of state-years use tests based on TRIM3's "dollar amount #4," the federal poverty level, or other standards. In these states, we instead apply the modal test in the data.

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.
- *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 Take-Up and Receipt Rates

In this section, we examine why take-up rates in Table 1 are smaller than official estimates. Here, we define take-up rates to be the share of households that take up, among households we simulate to be eligible. Receipt rates are the share of households that take up, among all households.

Summary Statistics. Appendix Table A14 reports sample-average estimates of take-up rates as well as three other quantities: (1) receipt rates among the simulated-ineligible, (2) rates of simulated eligibility among recipients, and (3) rates of simulated eligibility among nonrecipients.

Take-Up Rates in SNAP. Appendix Table A15 compares our estimates of take-up rates to USDA official estimates. Ours are clearly lower than the USDA estimates.

One explanation for the discrepancy is that, without conditioning on simulated eligibility, a lower share of households report that they receive SNAP in the PSID than are known to receive SNAP in administrative data. The implied rate of underreporting of SNAP benefits is roughly consistent with prior research (e.g., Meyer et al., 2009). Underreporting of receipt can thus explain some of the difference in take-up rates. We see this as encouraging: A substantial literature has developed around this form of measurement error in take-up rates, and we use corrections for underreporting developed in Mittag (2019) as a robustness check (see Section 3).

A second explanation for low take-up rates in the PSID relates to how eligibility enters official estimates. Appendix Table A15 shows two methods of computing SNAP take-up rates: (1) the ratio of total receipt to total simulated eligibility and (2) the receipt rate among the simulated eligible. The USDA takes the first approach, which is accurate only if every recipient is simulated ineligible. We avoid that assumption by taking the second approach in our main analyses.

By applying both approaches in the PSID, we can offer novel evidence on this potential contributor to the discrepancy in estimates of take-up rates. Appendix Table A15 shows that, when we follow the first approach, we move toward the USDA take-up rates. The second approach, however, yields a lower take-up rate due to what we call “misalignment”: Many recipient households in administrative data are likely to be simulated as ineligible if they were to appear in the survey data. Such households would enter the numerator but not the denominator of the USDA estimate, and this appears to be a key factor in the discrepancy. For this reason, official estimates of take-up rates

—which divide the administrative data in the numerator by simulated eligibility estimates, without adjusting for false negatives in eligibility—likely overstate true rates.

The most striking example of this upward bias is when the first approach yields estimates exceeding 100-percent take-up. For instance, in fiscal year 2018, the USDA reported take-up statistics with such issues in four states: DE, IL, OR, WA.⁴³ Furthermore, in all of these states in that year, there were specific demographic cells with official reported take-up rates of less than 100 percent, which is mathematically inconsistent with 100-percent take-up overall. In our view, these challenges suggest that survey-data measures of take-up are not obviously dominated by approaches leveraging administrative data.

For further external evidence on this “misalignment” problem, see [Newman and Scherpf \(2013\)](#) and [McBride et al. \(2023\)](#). Using data linked from surveys to administrative records, they observe receipt and simulate eligibility for the same people, studying SNAP and WIC respectively. Both find take-up rates well below official USDA estimates.

Benchmarking Other Programs. Appendix Table A16 reports a selection of estimated take-up rates in other research. These estimates come from a variety of data sources (administrative and survey) and time periods. Overall, our estimates of take-up rates are lower than administrative estimates, but they appear generally consistent with other survey-based estimates of self-reported receipt among simulated eligibles. For these programs, the most plausible explanation for low take-up rates is survey underreporting of transfer receipt, as in, e.g., [Meyer et al. \(2009\)](#).

B.6 Empirical Calibration

We now provide additional details on the empirical calibration of the model.

Utility Parameter ψ . We calibrate $\psi = w/l^{1/\eta}$. This internal calibration arises from the household’s labor–leisure first-order condition:

$$u_l = -w \cdot u_c \tag{16}$$

for hours l , consumption c , and hourly wage w . Obtaining expressions for u_l and u_c from GHH preferences, we obtain $\psi = w/l^{1/\eta}$.

Other Adjustments. Analysis uses household-level variables. We form utility using equivalized consumption and average hours worked across adults. We winsorize consumption and lifetime income at the fifth percentile, as well as social marginal utility at the 95th percentile, to diminish outliers that arise from isoelastic utility. We compute the hourly wage faced by the household by taking equivalized household wages in the PSID and dividing by total hours worked across adults.

⁴³See <https://www.fns.usda.gov/usamap/2018>.

Estimation. Estimation requires aggregating to income quantiles to compute the integrals in the labor-supply effect at the income-group level. We form 100 income-rank quantiles z , based on household equivalized income. We linearly interpolate the tax schedule across quantiles. We form the elasticity of taxable income at the quantile level. For households with zero income, we replace their ETI as the minimum of households with positive income. To estimate the labor-supply term, we convert the keep rate, expressed in the text in units of dollars, to income-rank units. We smooth inputs to the labor-supply term by averaging across quantiles: {1–5, 6–10, 11–15, ..., 26–30, 31–50, 50–100}; higher quantiles typically have small numbers of households because we restrict to the simulated-eligibles.

To estimate the labor-supply term, we apply a change of variables. We compute derivatives of any function f using the chain rule: $\frac{df}{dr} = \frac{df}{dz} \frac{dz}{dr}$, so $\frac{df}{dr} / \frac{dz}{dr} = \frac{df}{dz}$. This gives that the labor-supply term is:

$$\underbrace{\int_z \frac{\bar{M}'(z) z \bar{\varepsilon}_\tau(z)}{1 - T'(z)} \frac{d}{dz} \left(S(z) \bar{M}(z) - T(z) \right) dH(z)}_{\text{labor supply effect}} \approx \frac{1}{N} \sum_{r=1}^{100} \frac{1}{\frac{dz}{dr}} \frac{\frac{d\bar{M}}{dr} z(r) \varepsilon}{1 - \frac{dT}{dz}} \times \frac{1}{\frac{dz}{dr}} \left(\frac{d}{dr} [S(z) \bar{M}(z)] - \frac{dT}{dr} \right) \omega(r) \quad (17)$$

$$= \frac{1}{N} \sum_{r=1}^{100} \frac{\frac{d\bar{M}}{dr} z(r) \varepsilon}{\frac{dz}{dr} - \frac{dT}{dr}} \times \frac{1}{\frac{dz}{dr}} \left(\frac{d}{dr} [S(z) \bar{M}(z)] - \frac{dT}{dr} \right) \omega(r) \quad (18)$$

where $\omega(r)$ is the weighted number of households in each rank cell who are eligible, and N is the total number of households. We approximate derivatives with first differences.

C Further Results

C.1 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, Appendix Figure A31 shows little detectable change in either kind of self-targeting over time.

C.2 Measurement Error Simulation

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.⁴⁴ We simulate measurement error as follows:

1. Obtain the take-up rate among the bottom half of households ranked by equivalized consumption, $\hat{D} \in [0, 1]$.
2. Assign the top $x\%$ in consumption ranks to have some constant $c \in [0, 1]$ times the take-up rate of the bottom half: $c\hat{D}$.
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 half 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 ($c = 1$).

We find measurement error would need to be very severe to overturn our results. Appendix Figure A32 shows the estimated coefficient $\hat{\beta}$ 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 half's take-up rate. The blue lines show the estimates if the true take-up rate is the same as the bottom half'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 $\beta = 0$. If true take-up is equal, then we can no longer reject $\beta = 0$ 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.3 Bias of Estimator

Let $Y_i \in \mathbb{R}$ be the outcome variable, $D_i \in \{0, 1\}$ indicate transfer receipt, $E_i \in \{0, 1\}$ indicate true eligibility, and $\tilde{E}_i \in \{0, 1\}$ indicate simulated eligibility. We assume true eligibility E_i is unobservable, but that $D_i = 1$ implies $E_i = 1$, so that there is no reported transfer receipt among the truly ineligible. Without loss of generality, we can then write out the conditional expectation function as

$$E[Y_i | D_i, E_i, \tilde{E}_i] = \alpha + \beta D_i + \gamma D_i \tilde{E}_i + \delta \tilde{E}_i + \rho E_i + \eta E_i \tilde{E}_i.$$

⁴⁴We do not condition on simulated eligibility in these specifications, to isolate the magnitude of take-up measurement error without controlling for a potentially contaminated confound.

We wish to estimate

$$\Delta = E[Y_i | D_i = 1, E_i = 1] - E[Y_i | D_i = 0, E_i = 1],$$

but since E_i is unobservable, we can only estimate

$$\tilde{\Delta} = E[Y_i | D_i = 1, \tilde{E}_i = 1] - E[Y_i | D_i = 0, \tilde{E}_i = 1].$$

Let us first evaluate both conditional expectations to be differenced:

$$\begin{aligned} E[Y_i | D_i = 1, \tilde{E}_i = 1] &= \alpha + \beta + \gamma + \delta + (\rho + \eta) \Pr(E_i = 1 | D_i = 1, \tilde{E}_i = 1) \\ E[Y_i | D_i = 0, \tilde{E}_i = 1] &= \alpha + \delta + (\rho + \eta) \Pr(E_i = 1 | D_i = 0, \tilde{E}_i = 1), \end{aligned}$$

and thus

$$\begin{aligned} E[Y_i | D_i = 1, \tilde{E}_i = 1] - E[Y_i | D_i = 0, \tilde{E}_i = 1] &= \beta + \gamma + (\eta + \rho) [\Pr(E_i = 1 | D_i = 1, \tilde{E}_i = 1) - \Pr(E_i = 1 | D_i = 0, \tilde{E}_i = 1)] \\ &= \beta + \gamma + (\eta + \rho) [1 - \Pr(E_i = 1 | D_i = 0, \tilde{E}_i = 1)] \\ &= \beta + \gamma + (\eta + \rho) \Pr(E_i = 0 | D_i = 0, \tilde{E}_i = 1). \end{aligned}$$

Returning to what we want to estimate, let us again evaluate both conditional expectations to be differenced:

$$\begin{aligned} E[Y_i | D_i = 1, E_i = 1] &= \alpha + \beta + \rho + (\gamma + \delta + \eta) \Pr(\tilde{E}_i = 1 | D_i = 1, E_i = 1) \\ E[Y_i | D_i = 0, E_i = 1] &= \alpha + \rho + (\delta + \eta) \Pr(\tilde{E}_i = 1 | D_i = 0, E_i = 1), \end{aligned}$$

and thus

$$\begin{aligned} \Delta &= E[Y_i | D_i = 1, E_i = 1] - E[Y_i | D_i = 0, E_i = 1] \\ &= \beta + \gamma \Pr(\tilde{E}_i = 1 | D_i = 1, E_i = 1) \\ &\quad + (\delta + \eta) (\Pr(\tilde{E}_i = 1 | D_i = 1, E_i = 1) - \Pr(\tilde{E}_i = 1 | D_i = 0, E_i = 1)). \end{aligned}$$

and thus the bias equals

$$\begin{aligned} \Delta - \tilde{\Delta} &= -\gamma \Pr(\tilde{E}_i = 0 | D_i = 1, E_i = 1) + (\eta + \rho) \Pr(E_i = 0 | D_i = 0, \tilde{E}_i = 1) \\ &\quad - (\eta + \delta) (\Pr(\tilde{E}_i = 1 | D_i = 1, E_i = 1) - \Pr(\tilde{E}_i = 1 | D_i = 0, E_i = 1)). \end{aligned}$$

All three terms have an economic interpretation.

1. *Differential self-targeting among the simulated-eligible.* Self-targeting behavior may differ between the simulated-eligible and the simulated-ineligible, leading to bias. The magnitude of this bias increases with the false-negative rate, that is, the probability of misclassifying truly eligible recipients as simulated-ineligible. This probability is set to zero in Appendix Figure A23 by reclassifying all recipients as simulated-eligible.
2. *Imbalance on true eligibility.* Individuals who are truly ineligible but classified as simulated-eligible may have systematically different consumption levels than true eligibles. If the comparison between recipients and nonrecipients is unbalanced in terms of unobserved true eligibility, it can also be unbalanced in terms of consumption. The magnitude of this bias increases with the false-positive rate, meaning the probability of misclassifying true ineligibles as simulated-eligible. Appendix Figure A24 addresses this issue by analyzing self-targeting demographic groups with a high likelihood of true eligibility.
3. *Imbalance on simulated eligibility.* People classified as simulated-eligible may have different average outcome ranks compared to those classified as simulated-ineligible. A bias emerges to the extent that transfer receipt predicts simulated eligibility holding fixed true eligibility. This imbalance can be mitigated by controlling for the inputs to simulated eligibility, as illustrated in Appendix Figure A28. This strategy by construction eliminates, by construction, the explanatory power of simulated eligibility over consumption.

C.4 Marginal Value of Public Funds

We have also estimated the Marginal Values of Public Funds (MVPF of marginal increases in voluntary and automatic transfers (Hendren and Sprung-Keyser, 2020)). Our empirical implementation follows Section 4. The theoretical derivation of the MVPFs is in Appendix D.

Appendix Table A17 presents the voluntary-transfer and automatic-transfer MVPFs, along with a componentwise summary of their inputs. The main takeaway is an echo of Table 3: Overall, we find the MVPFs of increases in voluntary transfers exceed those of automatic transfers, but with considerable heterogeneity across programs.

Taking SNAP as an example, we find an automatic-transfer MVPF of 1.25 as compared to a voluntary-transfer MVPF of 1.65. Put another way, the marginal dollar spent on SNAP, inclusive of fiscal externalities, raises social welfare by an additional 30 cents if it is used to increase SNAP benefits as if it were provided automatically to all SNAP-eligibles.

We now highlight one aspect of the MVPF analysis that is novel relative to our welfare comparison. Considered alone, neither reform is distribution-neutral with respect to income, as each is an

increase in a progressive transfer that is financed by lump-sum taxation. Consequently, both reforms redistribute *across* income levels, not just within them. Perhaps surprisingly, the welfare benefits of the within-income redistribution often greatly exceed those from across-income redistribution. Such results underscore the importance of a consumption-based analysis of transfers.

D Theory Appendix

D.1 Proof of Proposition 1

Proof. We begin by analyzing the household's problem. Let η_t and λ_t denote, respectively, the Lagrange multipliers on the household's period- t borrowing and budget constraints. The first-order conditions of the household's problem (Equation 8) are

$$[c_t] \quad U'(c_t - v(z_t; \theta_t); \theta_t) - \eta_t = 0 \quad (19)$$

$$[z_t] \quad U'(c_t - v(z_t; \theta_t); \theta_t) v'(z_t; \theta_t) + \eta_t [1 - T'(z_t) + S'(z_t)M(z_t; \theta_t)] = 0 \quad (20)$$

$$[a_{t+1}] \quad (1 + \rho)^{-1} E[V'(a_{t+1}; \theta_{t+1}) | a_{t+1}, \theta_t] - \eta_t - \lambda_t = 0. \quad (21)$$

We note that, for the first-order condition with respect to z_t , there is no term $S(z)M'(z; \theta)$ coming from a reduction in the take-up probability when the household chooses a higher z_t . This is by the Leibniz rule:

$$\int_0^{S(z)} [S(z) - \kappa] \mu(\kappa | \theta) d\kappa \quad (22)$$

$$= [S(z) - S(z)] \cdot S'(z) - [S(z) - 0] \cdot 0 + \int_0^{S(z)} S'(z) \mu(\kappa | \theta) d\kappa \quad (23)$$

$$= S'(z) \int_0^{S(z)} \mu(\kappa | \theta) d\kappa \quad (24)$$

$$= S'(z)M(z; \theta). \quad (25)$$

Then, combining the first-order conditions with respect to c_t and z_t , we obtain

$$1 - T'(z_t) + S'(z_t)M(z_t; \theta_t) = -v'(z_t; \theta_t). \quad (26)$$

This equation highlights the key property of GHH preferences in our context, which is that the $U'(\cdot)$ terms cancel out. Consequently, there are no income effects on the labor supply.

Combining the first-order conditions for c_t and a_{t+1} , we derive the Euler equation:

$$U'(c_t - v(z_t; \theta_t); \theta_t) = (1 + \rho)^{-1} E[V'(a_{t+1}; \theta_{t+1}) | a_{t+1}, \theta_t] - \lambda_t. \quad (27)$$

Next, we calculate the elasticities of labor supply with respect to tax and transfer changes in terms of primitives. Following [Jacquet and Lehmann \(2014\)](#), we apply perturbations to the tax and transfer system about z_0 of the form $\hat{T} = T + \tau(z - z_0) - \nu$ and $\hat{S} = S + \varsigma(s - s_0) - \vartheta$. This reform increases the marginal tax rate or marginal transfer rate by τ or ς and decreases the tax level or transfer level by ν or ϑ .

The key change is to Equation 26, which becomes

$$\mathcal{F}_1 = 1 - T'(z_t) - \tau + [S'(z_t) - \varsigma] \Pr(S(z_t) + \varsigma(z_t - z_0) \leq \kappa_t | z_t; \theta_t) - v'(z_t; \theta_t). \quad (28)$$

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

$$\mathcal{F}_{1,z}|_{\tau=\varsigma=0} = -T''(z_t) + S''(z_t)M(z_t; \theta_t) + [S'(z_t)]^2 M(z_t; \theta_t) - v''(z_t; \theta_t) \quad (29)$$

$$\mathcal{F}_{1,\tau}|_{\tau=\varsigma=0} = -1 \quad (30)$$

$$\mathcal{F}_{1,\varsigma}|_{\tau=\varsigma=0} = M(z_t; \theta_t). \quad (31)$$

We note that $z_t \rightarrow z_0$ as $\tau, \varsigma, \vartheta \rightarrow 0$. By comparison, there is no impact of the reform on the Euler equation (Equation 27):

$$\mathcal{F}_2 = U'(c_t - v(z_t; \theta_t); \theta_t) - (1 + \rho)^{-1} \mathbb{E}[V'(a_{t+1}; \theta_{t+1}) | a_{t+1}, \theta_t] + \rho_t \quad (32)$$

As a result, we have that

$$\mathcal{F}_{2,a}|_{\tau=\varsigma=0} = -U''(c_t - v(z_t; \theta_t); \theta_t) - (1 + \rho)^{-1} \mathbb{E}[V''(a_{t+1}; \theta_{t+1}) | a_{t+1}, \theta_t] + \frac{\partial \rho_t}{\partial a_{t+1}} \quad (33)$$

$$\mathcal{F}_{2,\tau}|_{\tau=\varsigma=0} = 0 \quad (34)$$

$$\mathcal{F}_{2,\varsigma}|_{\tau=\varsigma=0} = 0. \quad (35)$$

Hence, by the implicit function theorem, we have that

$$\frac{\partial z}{\partial \tau} = \frac{-1}{\mathcal{F}_{1,z}(z; \theta)}, \quad \frac{\partial z}{\partial \varsigma} = \frac{M(z; \theta)}{\mathcal{F}_{1,z}(z; \theta)}, \quad \frac{\partial a}{\partial \tau} = \frac{\partial a}{\partial \varsigma} = 0. \quad (36)$$

For later use, we define the compensated labor supply elasticity as:

$$\varepsilon_z(z; \theta) = \frac{1 - T'(z)}{z} \frac{\partial z(\theta)}{\partial(1 - \tau)} \Big|_{\tau=0} \quad (37)$$

We now wish to be more precise about the specific perturbation reform. The reform is local to an income z_0 , and thus $\hat{T} = T + \tau(z - z_0) - \nu$ and $\hat{S} = S + \varsigma(s - s_0) - \vartheta$. We set the level shift in $S(z)$

to $\vartheta = ds$, the change in the marginal tax rate to $\tau = \frac{d}{dz} \bar{M}(z) ds$, and the decrease in everyone's level of taxes to $\nu = E[S(z)\bar{m}(z)] + \bar{M}(z) ds$ due to the reduced voluntary transfer expenditure.⁴⁵ Here, the term $\bar{m}(z)$ is the density of households with income z who are indifferent to take-up. Finally, any changes in revenue due to changes in labor supply are redistributed as a lump sum.

We proceed in deriving Equation 10 term by term.

First Term (Self-Targeting). The total direct tax change accruing to pre-reform income level z is an increase of $\bar{M}(z) ds$, which has a welfare effect of $E[\alpha(\theta)V'_a(a; \theta) | z = z(\theta)] \bar{M}(z) ds$ at z .

This welfare effect is weighed against an increase in the transfer by ds at all incomes. Again applying the Leibniz rule, this has a welfare effect of

$$ds \int_z E_{\kappa \leq S(z)} [\alpha(\theta)V'_a(a; \theta) | z] \bar{M}(z) dH(z). \quad (38)$$

Combining these two welfare effects, we arrive at the first term:

$$dW_1 = ds \int_z \bar{M}(z) (E_{\kappa \leq S(z)} [\alpha(\theta)V'_a(a; \theta) | z] - E[\alpha(\theta)V'_a(a; \theta) | z]) dH(z). \quad (39)$$

Then, by the law of total expectation, we have that $E[\alpha(\theta) | z] = \bar{M}(z) E_{\kappa \leq S(z)} [\alpha(\theta) | z] + (1 - \bar{M}(z)) E_{\kappa > S(z)} [\alpha(\theta) | z]$. Substituting this into Equation 38 and some algebra yields the following expression:

$$-ds \int_z \bar{M}(z) (1 - \bar{M}(z)) (E_{\kappa \leq S(z)} [\alpha(\theta)V'_a(a; \theta) | z] - E_{\kappa > S(z)} [\alpha(\theta)V'_a(a; \theta) | z]) dH(z),$$

and then we can use the definition of the regression coefficient $\beta(z) = E_{\kappa \leq S(z)} [\alpha(\theta) | z] - E_{\kappa > S(z)} [\alpha(\theta) | z]$ to obtain

$$- \int_z \frac{\bar{M}(z)(1 - \bar{M}(z))}{\int_z \bar{M}(z)(1 - \bar{M}(z)) dH(z)} \beta(z) dH(z),$$

and then by the variance-weighting property of least squares, we reach the expression in the main text:

$$dW_1 = -\beta \sigma_M^2 ds, \quad (40)$$

where the within-income variance of take-up is $\int_z \bar{M}(z)(1 - \bar{M}(z)) dH(z)$.

Second Term (Take-Up Costs). Due to the fiscal savings from marginal recipients, there is also

⁴⁵The derivative $\frac{d}{dz} \bar{M}(z)$ is a total derivative from the planner's perspective. That is, it shifts across people at different incomes z . It includes changes the $\bar{M}(z)$ in z due to both the the schedule $S(z)$ varying in z and the distribution of $\kappa | w$ varying in z .

a lump-sum transfer to all households. As the transfer benefit $S(z)$ declines by ds for all z , the total fiscal savings across all incomes z is $E[S(z)\bar{m}(z)]$, where $\bar{m}(z)$ is the density of $\bar{M}(z)$ near indifference to take-up.

The resultant welfare gain from the lump-sum transfer is $E[S(z)\bar{m}(z)] ds$, recalling the normalization of the population-average welfare weight to one. We now move from the density $\bar{m}(z)$ to a take-up elasticity by

$$\varepsilon_b(z; \theta) = S(z) \frac{\bar{m}(z)}{\bar{M}(z)}, \quad (41)$$

since $d\bar{M}(z)/dS(z) = \bar{m}(z)$. We can manipulate this expression to obtain $\bar{M}(z)\varepsilon_b(z; \theta) = S(z)\bar{m}(z)$. Using this as a substitution, the welfare gain is

$$dW_2 = dsE[\bar{M}(z)\varepsilon_b(z; \theta)] = ds\bar{M}(z)\bar{\varepsilon}_b(z), \quad (42)$$

where this weighted-average take-up elasticity is

$$\bar{\varepsilon}_b = \int_{z \times \Theta} \frac{\bar{M}(z)\varepsilon_b(z; \theta)}{\int_{z \times \Theta} \bar{M}(z) d\mu(z; \theta)} d\mu(z; \theta) \quad (43)$$

and $\bar{M}(z)$ reflects the overall take-up rate at income z .

Third Term (Labor Supply). The government has reduced tax burdens by $\bar{M}(z)ds$ at each z , which results in changes in marginal tax rates of $\bar{M}'(z)ds$ at each z . By assuming $\bar{M}'(z) < 0$, the implication is that the marginal rates will increase in this reform. Thus, labor supply contracts by $ds\bar{M}'(z)\frac{\partial z(\theta)}{\partial \tau}$ at each income z . The fiscal cost per $d\tau$ is $\frac{d}{dz}(T(z) - S(z)\bar{M}(z))$. We assume this fiscal externality is offset via lump-sum taxes on all households, and thus a welfare weight of unity is applied. We thus obtain:

$$dW_3 = ds \int_z \bar{M}'(z) \frac{\partial z(\theta)}{\partial \tau} \cdot \frac{d}{dz} (T(z) - S(z)\bar{M}(z)) dH(z), \quad (44)$$

and applying in the definition of the elasticity $\varepsilon_\tau(z; \theta) = -(1 - T'(z))/z \cdot \partial z(\theta)/\partial \tau$ yields the third term in the proposition:

$$dW_3 = ds \int_z \frac{\bar{M}'(z)z\bar{\varepsilon}_\tau(z)}{1 - T'(z)} \frac{d}{dz} (T(z) - S(z)\bar{M}(z)) dH(z), \quad (45)$$

This completes the proof.

□

D.2 Proof of Proposition 2

In the main text, we treat a dollar of automatic transfer as equivalent to a cash dollar, allowing us to model changes in the automatic transfer as occurring through the tax system. Here, we explicitly distinguish between the tax system, an automatic transfer, and a voluntary transfer. We then derive an analog of Proposition 1 in this enriched setting.

The government must now choose a tax schedule $T(z)$, a voluntary transfer $S_V(z)$ and an automatic transfer $S_A(z)$. The voluntary transfer is what we labeled as $S(z)$ in the main paper: households must pay a cost κ to take up. Taxes and the automatic transfer are received automatically by households at income z (without a cost being paid), except a dollar of the latter is not valued equally to the a dollar of the former. Let λ represent the marginal utility for a dollar of S_A or S_V relative to a dollar of cash.

In the household's problem, the period budget constraint (Equation 9) becomes

$$c_t + a_{t+1} = z_t - T(z_t) + R_t a_t + \lambda S_A(z_t) + \lambda \int_0^{S(z_t)} [S(z_t) - \kappa] \mu(\kappa|w_t) d\kappa, \quad (46)$$

with all other equations left unchanged. We emphasize that the government's period budget constraint remains

$$\int_{\Theta} [T(z(\theta)) - S_A(z(\theta)) - \mathbb{1}_S S_V(z(\theta))] d\mu(a, \theta) = 0. \quad (47)$$

Even though a dollar of S_A or S_V is only valued at λ dollars by the household, that is, it costs the government a dollar to provide.

Reform. We define the reform analogously to the primary reform in Section 4. The voluntary transfer amount is cut by ds_V at all incomes. With the savings, at each income z , automatic transfers are increased by $s_A(z) = \bar{M}(z) ds_V$, so that people at each income level are compensated on average for the voluntary transfer cut. The slope of S_A rates (analogous to marginal tax rates) thus change by $S'_A(z) = \frac{d}{dz} \bar{M}(z) ds_V$ at z . Fiscal savings from marginal transfer recipients are redistributed as a lump sum automatic transfer. The revenue cost of any labor supply response is then paid for via lump-sum taxes (i.e. as cash, not in-kind).

We calculate the welfare effects of this reform analogously to Proposition 1.

Proposition 2. *The welfare effect of the reform is*

$$\begin{aligned} \frac{1}{\lambda} \frac{dW}{ds} = & \underbrace{-\beta\sigma_M^2}_{\text{lost value of self-targeting}} + \underbrace{\bar{M}\bar{\varepsilon}_b}_{\text{fiscal savings from marginals}} \\ & + \underbrace{\int \frac{\bar{M}'(z)z\bar{\varepsilon}_\tau(z)}{1-T'(z)} \left(\frac{d}{dz}(S_V(z)\bar{M}(z)) + S'_A(z) - T'(z) \right) dH(z)}_{\text{labor-supply effect}}. \end{aligned} \quad (48)$$

where all other terms are as in Proposition 1.

While Proposition 1 moves money from from S to T , Proposition 2 moves money from S_V to S_A . In both cases, the dollar being moved has the same ex-post marginal utility (1 in S and T , λ in S_V and S_A). Moving a dollar from the those who take up the voluntary transfer to everyone decreases marginal utility by λ times the lost value of self-targeting from Proposition 1. Similarly, all the dollars saved from the marginals not taking up S_V are redistributed in-kind through S_A and so the utility value is λ times the fiscal savings from marginals term in Proposition 1.

As for the labor-supply effects, the first term is the change in the government budget due to labor supply changes that is then redistributed as an automatic transfer to everyone. But unlike Proposition 1, where marginal tax rates changed, here the slope of the S_A schedule changes, and so the relevant elasticity is ε_{S_A} not ε_z . However, per the household's problem, $\frac{dz}{ds_A} = \lambda \frac{dz}{d\tau}$.

The intuition is that the costs and benefits of this reform are all scaled by λ relative to the reform in Proposition 1. Utility benefits/costs are λ lower since they are paid in in-kind dollars rather than cash dollars, and impacts on the budget constraint, although intrinsically denominated in cash, are convertible to in-kind dollars since labor supply response to a dollar change in S_A is λ times the labor supply response to a dollar change in T .

Proof. The proof is very similar to the proof of Proposition 1. The first term, the lost value of self-targeting, is λ multiplied by the term in Proposition 1. The same is true for the second term, since the fiscal savings are redistributed in-kind. The key differences are in the labor supply term.

Since the slope of the automatic transfer schedule has risen (as $\bar{M}'(z) < 0$), labor supply contracts by $ds\bar{M}'(z) \frac{\partial z}{\partial s_A}$ at each income z . The fiscal cost per $d\tau$ is $T'(z) - \frac{d}{dz}(S(z)\bar{M}(z)) - S'_A(z)$. However, examining the first-order condition of the household, we have that $\frac{\partial z}{\partial s_A} = \lambda \frac{\partial z}{\partial \tau}$. Thus, $ds\bar{M}'(z) \frac{\partial z}{\partial s_A} = ds\bar{M}'(z) \lambda \frac{\partial z}{\partial \tau}$ and plugging in the definition of the elasticity: $\varepsilon_z(\theta) = -\frac{\partial z(\theta)}{\partial \tau} \frac{1-T'(z)}{z}$ and by assumption all fiscal costs are paid lump sum, that is by the average welfare weight $E[\alpha(\theta)] = 1$, yields the first labor supply term after rearrangements similar to Proposition 1. \square

D.3 Proof of Proposition 3

Proposition 3. *The Marginal Value of Public Funds for increases in the voluntary and automatic transfers are respectively*

$$\text{MVPF}_v = \frac{1 + (\beta_{\text{across}}\sigma_{M,\text{across}}^2 + \beta_{\text{within}}\sigma_{M,\text{within}}^2)/\bar{M}}{1 + \bar{\varepsilon}_b}, \quad (49)$$

$$\text{MVPF}_v = \frac{1 + \beta_{\text{across}}\sigma_{M,\text{across}}^2/\bar{M}}{1 + \text{FE}_a/\bar{M}}, \quad (50)$$

where the fiscal externality of the automatic transfer is

$$\text{FE}_a = \int_z \frac{\bar{M}'(z)z\bar{\varepsilon}_\tau(z)}{1 - T'(z)} \frac{d}{dz} \left(S(z)\bar{M}(z) - T(z) \right) dH(z). \quad (51)$$

Proof. We begin by considering the automatic reform. The mechanical cost of cutting taxes by $\bar{M}(z)$ at each z is simply the population-average take-up rate $\bar{M} = \int_z M(z) dH(z)$. There is also a fiscal externality of adjustments in labor supply due to changes in marginal tax rates, which is identical to the final term in Proposition 1. The welfare impact of the automatic reform on inframarginal recipients is $E_z [M(z)\bar{\alpha}(z)]$, where we use the shorthand that $\bar{\alpha}(z) = E_\kappa [\alpha(\theta)V'_a(a; \theta) | z]$. We derive the numerator as follows:

$$E_z [\bar{M}(z)\bar{\alpha}(z)] = E_z [\bar{M}(z)] E_z [\bar{\alpha}\bar{M}(z)] + \text{Cov}_z (\bar{M}(z), \bar{\alpha}(z)) \quad (52)$$

$$= (\bar{M})(1) + \beta_{\text{across}}\sigma_{M,\text{across}}^2 \quad (53)$$

$$= \bar{M}(1 + \beta_{\text{across}}\sigma_{M,\text{across}}^2/\bar{M}). \quad (54)$$

We note that the second line follows from the definition of \bar{M} and the normalization of the population-average welfare weight to one. We then rewrite the covariance as the product of variance and the least-squares coefficient β_{across} from a cross-sectional regression of $\bar{\alpha}(z)$ on $\bar{M}(z)$ over values of z . We will divide both the numerators and the denominators of the MVPF by \bar{M} to obtain the final expression, which we think of as expressed in units of dollars spent rather than dollars “authorized” under complete take-up.

We turn now to the voluntary reform. The voluntary transfer schedule is increased by a dollar at all incomes z . The mechanical cost of this is \bar{M} . The fiscal externality of the marginal recipients is $E_z [S(z)\bar{m}(z)]$, which, as in Proposition 1, is equal to $E_z [\bar{M}(z)\bar{\varepsilon}_b(z)]$. Noting that will divide numerator and denominator by \bar{M} , this yields the same weighted-average take-up elasticity $\bar{\varepsilon}_b$ as in Proposition 1.

The benefits of the voluntary reform at are $\int_z E_{\kappa \leq S(z)} [\alpha(\theta)V'_a(a; \theta) | z] \bar{M}(z) dH(z)$. Recall the

regression coefficient is

$$\beta(z) = \mathbb{E}_{\kappa \leq S(z)} [\alpha(\theta) V'_a(a; \theta) | z] - \mathbb{E}_{\kappa > S(z)} [\alpha(\theta) V'_a(a; \theta) | z].$$

By the law of total expectation

$$\bar{\alpha}(z) = M(z) \mathbb{E}_{\kappa \leq S(z)} [\alpha(\theta) V'_a(a; \theta) | z] + (1 - \bar{M}(z)) \mathbb{E}_{\kappa > S(z)} [\alpha(\theta) V'_a(a; \theta) | z].$$

Combining these expressions yields

$$(1 - \bar{M}(z)) (\mathbb{E}_{\kappa \leq S(z)} [\alpha(\theta) V'_a(a; \theta) | z] - \beta(z)) + \bar{M}(z) \mathbb{E}_{\kappa \leq S(z)} [\alpha(\theta) V'_a(a; \theta) | z] = \bar{\alpha}(z).$$

Hence,

$$\mathbb{E}_{\kappa \leq S(z)} [\alpha(\theta) V'_a(a; \theta) | z = z(\theta)] = \bar{\alpha}(z) + \beta(z)(1 - \bar{M}(z)).$$

Substituting this into the benefits of the voluntary reform leads to

$$\int_z \mathbb{E}_{\kappa \leq S(z)} [\alpha(\theta) V'_a(a; \theta) | z] \bar{M}(z) dH(z) = \mathbb{E}_z [\bar{M}(z) \bar{\alpha}(z) + \beta(z) \bar{M}(z)(1 - \bar{M}(z))].$$

We then proceed by decomposing the first term as we did for the automatic reform. The second term uses the definition that $\sigma_{M, \text{within}}(z) = \bar{M}(z)(1 - \bar{M}(z))$, as in Proposition 1. This yields the numerator for MVPF_v and concludes the proof. □

D.4 Additional Proofs

We now establish conditions under which there is a social benefit from self-targeting. These proofs formally establish claims in Section 4 about the signs of terms in Proposition 1.

Proposition 4. *If the transfer $S(z)$ is positive, and if self-targeting is advantageous, then the first term in Equation 10 is negative, implying a welfare cost of self-targeting forgone due to a marginal shift toward automatic transfers.*

Proof. Recall that Definition 1 makes the marginal social benefit of transfers, $E[\alpha(\theta) V'_a(a; \theta) | \kappa, z]$, “co-monotone” in take-up costs κ for each income level z . By Schmidt (2003), co-monotonicity of random variables X and Y implies $E[XY] \geq E[X]E[Y]$. Rewriting gives

$$E[\alpha(\theta) V'_a(a; \theta) \mathbf{1}\{\kappa \leq t\} | z] \geq E[\alpha(\theta) V'_a(a; \theta) | z] \Pr(\kappa \leq t | z),$$

and dividing both sides by $\Pr(\kappa \leq t)$ completes the proof:

$$\mathbb{E}[\alpha(\theta)V'_a(a; \theta) \mid \kappa \leq t, z] \geq \mathbb{E}[\alpha(\theta)V'_a(a; \theta) \mid z].$$

□

We now establish conditions under which the policy reform we consider in the main text would raise or lower labor supply.

Proposition 5. *Suppose (1) the tax system is optimal and (2) take-up decreases in income (i.e., $\bar{M}'(z) < 0$). Then the labor supply effect in Equation 10 is negative.*

Proof. We derive a necessary 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):

$$\begin{aligned} d\tau \int_{x \geq z} (\mathbb{E}[\alpha(\theta)V'_a(a; \theta)] - \mathbb{E}[\alpha(\theta)V'_a(a; \theta) \mid z = x]) dH(x) \\ = E_{x \geq z} [\mathbb{E}[\alpha(\theta)] - \mathbb{E}[\alpha(\theta)V'_a(a; \theta) \mid z = x]] (1 - H(z)). \end{aligned} \quad (55)$$

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

$$d\tau \frac{\partial z}{\partial \tau} \cdot T'(z) \cdot h(z) \cdot \mathbb{E}[\alpha(\theta)V'_a(a; \theta)] = d\tau \left(-\epsilon^z \cdot z \cdot \frac{T'(z)}{1 - T'(z)} h(z) \mathbb{E}[\alpha(\theta)V'_a(a; \theta)] \right). \quad (56)$$

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

$$d\tau \frac{\partial z}{\partial \tau} \frac{d}{dz} [-S(z)\bar{M}(z)] h(z) \mathbb{E}[\alpha(\theta)V'_a(a; \theta)] \quad (57)$$

$$= d\tau \frac{\epsilon^z \cdot z}{1 - T'(z)} \frac{d}{dz} [S(z)\bar{M}(z)] h(z) \mathbb{E}[\alpha(\theta)V'_a(a; \theta)]. \quad (58)$$

A necessary condition for the optimality of the tax system is that the sum of these welfare effects is weakly negative. In particular, to convert to utility units, suppose the net fiscal externality from the change in marginal tax rates 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. Consequently,

$$E_{x \geq z} [E[\alpha(\theta)V'_a(a; \theta)] - E[\alpha(\theta)V'_a(a; \theta)|z = x](1 - H(z))] \quad (59)$$

$$\leq d\tau E[\alpha(\theta)] \left(T'(z) - \frac{d}{dz} [S(z)\bar{M}(z)] \right) \left(\frac{h(z)\epsilon^z z}{1 - T'(z)} \right). \quad (60)$$

The left hand side is positive, and therefore so must the right hand side be positive. Immediately we have that the labor supply term is negative, as hypothesized. \square

References for Appendices

Abowd, John M and Martha H Stinson, “Estimating Measurement Error in Annual Job Earnings: A Comparison of Survey and Administrative Data,” *Review of Economics and Statistics*, 2013, 95 (5), 1451–1467.

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.

Chetty, Raj and Adam Szeidl, “Consumption Commitments and Risk Preferences,” *The Quarterly Journal of Economics*, 2007, 122 (2), 831–877.

Congressional Research Service, “Need-Tested Benefits: Estimated Eligibility and Benefit Receipt by Families and Individuals,” Report R44327 2015.

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.

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

Guan, Alice, Akansha Batra, Hilary Seligman, and Rita Hamad, “Understanding the Predictors of Low Take-up of the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC): A Nationwide Longitudinal Study,” *Maternal and Child Health Journal*, 2023, 27 (10), 1795–1810.

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

- Hendren, Nathaniel and Ben Sprung-Keyser**, “A Unified Welfare Analysis of Government Policies,” *Quarterly Journal of Economics*, 2020, 135 (3), 1209–1318.
- Jacquet, Laurence and Etienne Lehmann**, “Optimal Nonlinear Income Taxation with Multi-dimensional Types: The Case with Heterogeneous Behavioral Responses,” THEMA Working Papers 2014-01, Université de Cergy-Pontoise 2014.
- McBride, Linden, Thomas B Foster, Renuka Bhaskar, Mark Prell, Maria Perez-Patron, Erik Vickstrom, Brian Knop, and Michaela Dillon**, “Integrating Administrative and Survey Data to Estimate WIC Eligibility and Access,” *Journal of Survey Statistics and Methodology*, 2023, 11 (3), 668–687.
- McGarry, Kathleen M and Robert F Schoeni**, “Understanding Participation in SSI,” Working Paper 319, Michigan Retirement Research Center 2015.
- 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.
- Newman, Constance and Erik Scherpf**, “Supplemental Nutrition Assistance Program (SNAP) Access at the State and County Levels: Evidence from Texas SNAP Administrative Records and the American Community Survey,” 2013.
- Olsen, Edgar O**, “Housing Programs for Low-Income Households,” in “Means-Tested Transfer Programs in the United States,” University of Chicago Press, 2003, pp. 365–442.
- 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,” Working Paper, Professoren des Inst. für Math. Stochastik 2003.
- Sommers, Ben, Rick Kronick, Kenneth Finegold, Rosa Po, Karyn Schwartz, and Sherry Glied**, “Understanding Participation Rates in Medicaid: Implications for the Affordable Care Act,” ASPE Issue Brief, Office of the Assistant Secretary for Planning and Evaluation, U.S. Department of Health and Human Services 2012.

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

U.S. Department of Agriculture, “National- and State-Level Estimates of WIC Eligibility and WIC Program Reach in 2019,” Nutrition Assistance Program Report Series, Food and Nutrition Service, Office of Policy Support 2022.

—, “Trends in Supplemental Nutrition Assistance Program Participation Rates: Fiscal Years 2016 to Fiscal Year 2020,” Technical Report, Food and Nutrition Service, Office of Policy Support 2022.

Zimmermann, Wendy and Karen C Tumlin, “Patchwork Policies: State Assistance for Immigrants Under Welfare Reform,” Occasional Paper 24, Urban Institute 1999.