Multiple Program Participation in the Safety Net: Incidence, Impediments, and Implications*

Neil A. Cholli† Cornell University

Derek Wu‡ University of Virginia

This Draft: March 26, 2026

Abstract

Multiple program participation is a defining feature of the U.S. safety net, with half of recipients enrolling in two or more programs. Yet most research examines programs in isolation, missing the intensity of safety net attachment. We show the implications of a multiple-program framework on take-up, targeting, and welfare analyses, using administrative data and a reform that streamlined SNAP, Medicaid, and TANF applications. Multiple program participation increased by 5%, with 20–56% of gains driven by single-to-multiple-program transitions among the most disadvantaged—a margin missed by single-program frameworks. Evaluating with multiple instead of single programs alters welfare calculations by up to 64%.

*We thank Jeremy Barofsky, Marianne Bitler, Jason Cook, Manasi Deshpande, Chloe East, Erik Hembre, Tim Layton, Lee Lockwood, Michael Lovenheim, Doug Miller, Matt Notowidigdo, Zhuan Pei, Lucie Schmidt, and seminar and conference attendees at Cornell University, Florida State University, the 2026 AEA/ASSA Annual Meetings, the 2025 NTA Annual Meeting, the 2025 APPAM Fall Research Conference, the NBER Children & Families Spring 2025 Meeting, the 6th World Labor Conference, the Virginia Longitudinal Data System Research Forum, and the Virtual Economics of Poverty and Policy Seminar for helpful comments and discussions, as well as Aline Jesus Rafi, Andrew Sell, Will Goldschmidt, and Tamra Arant for help with data acquisition and understanding institutional details. We also thank Dylan Craig and Heidi Clesner for excellent research assistance. We are grateful to the Russell Sage Foundation and the Washington Center for Equitable Growth for financial support. Cholli thanks the Klarman Fellowship at Cornell University and Wu thanks the Bankard Fund for Political Economy at the University of Virginia for support. This research was approved by the University of Virginia's Institutional Review Board (Protocol #7503). All conclusions and analyses do not reflect the views of the Virginia Criminal Sentencing Commission, the Virginia Department of Juvenile Justice, the Virginia Department of Social Services, or the Virginia Employment Commission.

†Email: nac85@cornell.edu

‡Email: derek.wu@virginia.edu

1 Introduction

A defining feature of the contemporary U.S. social safety net is its tapestry of means-tested transfer programs. Many of the most important programs—ranging from in-kind assistance for food (Supplemental Nutrition Assistance Program or SNAP) and medical bills (Medicaid) to cash assistance (Temporary Assistance for Needy Families or TANF)—have overlapping eligibility criteria, such that multiple program participation today is the norm rather than the exception. Thus, a modern-day understanding of safety net participation requires going beyond the receipt of *any given* program to considering the intensity of safety net attachment through the receipt of *multiple* programs. Yet, the extant literature usually studies individual programs in isolation. This risks missing how policies reshape the nature of safety net attachment—and, consequently, how their impacts should be understood and evaluated. This matters especially for households with children, who are disproportionately eligible for multiple programs and central to the social returns of the safety net (Aizer et al. 2022).

This paper examines the implications of a multiple-program framework on take-up, targeting, and welfare evaluation. We start by analyzing the breadth and nature of multiple program receipt using administrative data. We then consider how a multiple-program framework shapes our assessment of a major reform that simplified enrollment across various programs. By estimating the reform's causal impacts on multiple program receipt, we go beyond take-up of any given program to assess changes in the intensity of safety net attachment. We introduce a partial identification approach to disentangle extensive-margin responses (entry into the safety net) and intensive-margin responses (transitions from single to multiple programs). This allows us to distinguish how responders along these distinct margins differ in socioeconomic status, expanding the scope of traditional analyses of targeting based on a single (i.e., the extensive) margin. Finally, we quantify the degree to which considering multiple programs rather than single programs reappraises welfare evaluations of the reform. Collectively, our results demonstrate that adopting a multiple-program framework can reveal new and important insights for evaluating social policies.

Although many programs serve similar populations, our basic understanding of multiple program receipt remains limited. Prior studies have largely relied on national surveys like the Survey of Income and Program Participation (Blank and Ruggles 1996; Edelstein et al. 2014; Moffitt 2016; Jackson and Fanelli 2023; Macartney and Ghertner 2023), which are known to underreport incomes (Meyer et al. 2015; Meyer and Wu 2018). In our context, 30–60% of joint SNAP and Medicaid recipients in Virginia are missed in conventional surveys. We overcome this challenge by leveraging linked administrative microdata across SNAP, Medicaid, and TANF from Virginia to track the universe of joint program recipients at the

monthly level. We therefore begin by documenting several descriptive facts on the prevalence and nature of multiple program participation with our high-quality data.

First, we find that half of all program recipients participate in two or more programs. The vast majority (80%) of multiple program recipients reside in households with children, who are among the most common beneficiaries of the safety net (Currie 2006b) and thus our population of focus in all ensuing analyses. Moreover, transitions into and out of programs often occur jointly, highlighting how safety net receipt can be best understood in terms of program "bundles." However, we find striking evidence of incomplete take-up of multiple programs—even among eligible individuals already attached to the safety net. Between 10–60% of various subgroups do not participate in additional programs despite being inferred to be eligible based on receipt of another program. Incomplete take-up along this intensive margin can explain a large fraction of overall non-participation: nearly 50% of eligible children not enrolled in Medicaid already receive another program.

Incomplete take-up of multiple programs among existing safety net recipients could reflect two competing hypotheses. On the one hand, households may select into fewer programs since they assign little to no value to other programs that they are eligible for. On the other hand, administrative burdens—such as application hassle costs, information frictions, and stigma—may impede access to additional benefits they highly value (Herd and Moynihan 2018). Descriptively, we find that multiple program participation rates decline with distance from county field offices that serve program applicants, suggesting the presence of travel costs and other geographic barriers. Yet, there is scarce causal evidence on how burdens might hinder full participation in all qualifying programs and the types of individuals they screen out. Standard theoretical models (e.g., Nichols and Zeckhauser 1982; Moffitt 1983; Bertrand et al. 2004; Kleven and Kopczuk 2011) and many empirical analyses (e.g., Alatas et al. 2016; Deshpande and Li 2019; Rafkin et al. 2025) of the targeting properties of burdens are largely based on selection into individual programs or any safety net attachment in general. Nevertheless, introducing multiple margins of selection—based on the intensity of safety net attachment—can fundamentally alter and broaden our view of how burdens shape the take-up and targeting efficiency of the safety net *as a whole*.

Motivated by the descriptive facts, we employ a multiple-program framework to examine the take-up, targeting, and welfare effects of a policy reform in Virginia called CommonHelp. Before the reform, program applicants were required to hand-deliver or mail lengthy paperwork to their local field office. CommonHelp reduced these burdens by streamlining application and recertification processes onto a centralized online platform. As of 2024, 31 states have implemented similar integrated online platforms for SNAP, Medicaid, and TANF (Code for America 2024), reflecting growing national interest in improving access to the so-

cial safety net (National Governors Association 2022). Our emphasis on burdens for *multiple* programs is thus timely and promising for informing state policies around the country.

CommonHelp's statewide rollout in 2012 provided no natural control group, complicating causal identification of effects on program receipt. To address this, we introduce an empirical framework that combines two complementary research designs to identify the reform's total effect on program participation. First, we use a difference-in-differences approach that compares changes in enrollment across zip codes varying in distance to their assigned field offices, with treatment effects increasing in distance. This design isolates the effects of reducing "geographic" burdens more pronounced in distant areas, such as travel costs. While similar spatial difference-in-differences strategies have been used in prior work (e.g., Linden and Rockoff 2008; Currie and Walker 2011; Marcus 2021), such a design captures only *relative* differences between near and far zip codes and misses the *absolute* effects experienced even in nearby areas. We therefore complement this approach with a variant of a regression discontinuity design around the time of treatment, focusing on zip codes close to their office and differencing out confounding seasonality effects. This design isolates the effects of reducing "fixed" burdens (unrelated to distance and thus relevant for all applicants), such as office congestion and the stigma associated with in-person visits.

We find that CommonHelp significantly increased safety net participation, especially in multiple programs. Combining results from geographic and fixed burdens, we estimate that CommonHelp brought in 64,000 additional multiple program recipients (a 4.6% increase) on average over a three-year period. Crucially, focusing on geographic burdens alone would miss 54% of the growth in multiple program receipt, underscoring the importance of utilizing both research designs to estimate total effects in our framework. These results reveal that a considerable share of the population reap positive value from accessing multiple programs but would have been deterred by burdens from the original enrollment process.

Our results also demonstrate how policy reforms can affect not only *whether* individuals participate in the safety net but also *how intensively* they engage with it. CommonHelp's impact on multiple program participation *exceeded* its impact on caseloads of any program by roughly 6,000 recipients. Correspondingly, single program recipient counts *fell* by about 13,000, consistent with transitions from single to multiple programs. Focusing solely on any program receipt would conceal how policies can operate through transitions between bundles.

To unpack these aggregate enrollment effects, we develop a potential outcomes framework and employ a partial identification approach to disentangle the margins of selection into multiple program receipt. We invoke a monotonicity condition to bound the shares of recipients induced into specific program bundles along the extensive margin (from no programs to single or multiple programs) and intensive margin (from single to multiple programs). Our

evidence suggests that 20–56% of the increase in program receipt occurred along the intensive margin, indicating that CommonHelp reached those who would have otherwise been only partially attached to the safety net. Based on back-of-the-envelope calculations, CommonHelp closed the intensive-margin take-up gap by at least 20–42%—a sizable effect given that incomplete take-up among existing recipients can explain half of overall non-participation for some groups. Critically, these impacts along the intensive margin would be missed had individual program receipt been studied in isolation.

We next investigate *who* the reform reached by characterizing responders along the extensive and intensive margins. To capture socioeconomic disadvantage—a proxy for marginal utility of program receipt—we link administrative data on earnings and criminal justice involvement. By connecting observable measures of disadvantage to unobservable counterfactual choice behavior, our framework adds a novel dimension to targeting analyses. Strikingly, intensive-margin responses were concentrated among individuals with lower socioeconomic status (i.e., lower labor market attachment and greater criminal justice involvement). These responders, who would have self-selected into part of the safety net absent the reform, were more disadvantaged not only than extensive-margin responders but also than those who would have enrolled in multiple programs *without* the reform. This implies that burdens were *inefficiently* screening out existing recipients with a high marginal utility of additional benefits. This points to a counterintuitive policy insight: beyond expanding coverage to new populations, efforts to target the most disadvantaged should consider improving program access among *current* participants. This departs from conventional targeting wisdom that solely focuses on extensive-margin attachment.

Finally, we combine our enrollment estimates with institutional parameters and estimates from prior studies to examine the welfare impacts of the policy reform. Specifically, we evaluate how CommonHelp's Marginal Value of Public Funds (MVPF)—defined as the ratio of a policy's willingness-to-pay to its net fiscal cost (Hendren and Sprung-Keyser 2020)—would be misstated if one fails to account for impacts on multiple programs. We compare the complete MVPF that accounts for CommonHelp's effects on all programs against partial MVPFs that would be calculated if researchers only observed data for a single program. Since MVPF levels may be sensitive to modeling assumptions, we instead focus on differences between single- and multiple-program MVPFs to isolate the role of considering multiple programs. We find that single-program MVPFs—and even weighted averages of them—can differ substantially from the complete MVPF, overstating the actual value by up to 25% or understating it by up to 64% depending on modeling assumptions. This suggests that welfare analyses of reforms can critically hinge on whether one adopts a multi-program perspective, given the interconnected nature of the modern-day safety net.

This paper contributes to an expansive literature that has largely focused on the take-up and targeting of individual programs (see Currie 2006a and Ko and Moffitt 2024 for reviews). While these studies provide valuable insights, they can overlook the broader context in which programs overlap (Heflin et al. 2022; Wu and Meyer 2023; Schmidt et al. 2024). Our multiple-program framework introduces a new dimension to take-up and targeting analyses by considering the intensity of safety net attachment. Multiple margins of selection expand the scope of traditional models of take-up and targeting that have exclusively focused on the extensive margin. This builds on a longstanding literature considering multiple pathways into safety net receipt due to social, economic, or health shocks or underlying risk (e.g., Plotnick 1983; Blank and Ruggles 1996; Blundell and Pistaferri 2003; Deshpande and Lockwood 2022; Wu and Zhang 2025). Relative to single-program analyses, a multiple-program framework yields a substantially different picture of who is reached and how reforms affect welfare.[1] This complements a burgeoning literature on spillovers between different programs.[2]

Methodologically, we draw on the econometric literature related to multiple treatments, which generate various "response types" with different fallback options (e.g. Kirkeboen et al. 2016; Heckman and Pinto 2018; Mountjoy 2022; Pinto 2022). Our partial identification approach to bounding response types complements empirical and methodological efforts that bound behavioral responses in the context of structural models (Kline and Tartari 2016) and experimental or instrumental variable frameworks (Borusyak 2015; Mountjoy 2024). Separately, our empirical framework for identifying CommonHelp's effects on program receipt provides a flexible approach for evaluating any policy change rolled out simultaneously to an entire population (a setting without an obvious control group), so long as there are heterogeneous effects among different subpopulations (e.g., by distance to field office).

Finally, this paper advances the literature on administrative burdens (Herd and Moynihan 2018; Herd et al. 2023). Prior work has examined burdens in individual programs such as SNAP (e.g., Finkelstein and Notowidigdo 2019; Gray 2019; Homonoff and Somerville 2021; Unrath 2024), health insurance (e.g., Herd et al. 2013; Heinrich et al. 2021; Ericson et al. 2023; Arbogast et al. 2024; Shepard and Wagner 2025), the Earned Income Tax Credit (e.g., Bhargava and Manoli 2015; Linos et al. 2022), and unemployment benefits (e.g., Ebenstein and Stange 2010; Castell et al. 2025). We are among the first to examine

1. A smaller strand of the literature uses structural, simulation-based methods to study the effects of multiple program participation on marginal tax rates and on downstream outcomes like labor supply (Keane and Moffitt 1998; Flood et al. 2004; Chan 2013).

2. Prior work has studied how Unemployment Insurance receipt (e.g. Lindner 2016; Mueller et al. 2016; Leung and O'Leary 2020; East and Simon 2024) and expansions in the Earned Income Tax Credit (e.g., Grogger 2003; Hoynes and Patel 2018; Bastian and Jones 2021), health insurance (e.g., Baicker et al. 2014; Burns and Dague 2017; Schmidt et al. 2019; Heflin et al. 2023), and SNAP (e.g., Han 2020) have spillover effects on the receipt of other programs.

how barriers affect take-up and targeting across *multiple* programs. Our targeting estimates along the intensive margin suggest that burdens may have compounding effects for multiple-program eligibles who are particularly disadvantaged and may be less equipped to overcome them (Christensen et al. 2020). Our research design around "geographic burdens" also builds on past work demonstrating the importance of geographic proximity to field offices (Rossin-Slater 2013; Deshpande and Li 2019; Hicks 2024; Bitler et al. 2025).

Our paper proceeds as follows. Section 2 describes the institutional background and our data sources. Section 3 presents descriptive facts on multiple program participation, which help motivate the subsequent analyses. Sections 4 and 5 develop our empirical framework and report CommonHelp's causal impacts on program enrollment. Sections 6, 7, and 8 describe the implications of considering multiple programs for take-up, targeting, and welfare evaluation, respectively. Section 9 concludes.

2 Background and Data Sources

2.1 Background on Programs and Enrollment Procedures

We focus on three key means-tested programs in the U.S. social safety net: the Supplemental Nutrition Assistance Program (SNAP), Medicaid, and the Temporary Assistance for Needy Families (TANF). In 2023, these programs collectively disbursed roughly $1 trillion in benefits (Centers for Medicare & Medicaid Services 2024; USDA Food and Nutrition Service 2024). All three programs are federally funded, while eligibility rules and enrollment procedures can vary across states. Appendix Table A.1 summarizes the eligibility criteria of these programs in Virginia as of 2012, the year when CommonHelp was launched.

SNAP (formerly Food Stamps) provides in-kind benefits for food and beverage purchases and is often considered a "cash-like" benefit. Unlike other programs, SNAP eligibility is determined almost exclusively by household income and assets, making it broadly available regardless of age, presence of children, or medical condition. Thus, even able-bodied adults without dependents (ABAWDs) who meet income and asset limits are eligible for SNAP.[3] However, households with children generally qualify for higher benefit amounts (due to larger case sizes) and face fewer work-related requirements.

Medicaid is a public health insurance program that covers medical expenses for low-income children, pregnant women, certain parents, seniors, and people with disabilities under federal guidelines. Virginia also administers a supplementary healthcare program called the

3. While ABAWDs typically face work requirements, these were temporarily waived in Virginia during the Great Recession and reinstated in October 2013.

Family Access to Medical Insurance Security Plan (FAMIS), which extends coverage to children under age 19 whose family income falls under 205% of the Federal Poverty Line (FPL). Throughout this paper, we use "Medicaid" as an umbrella term for both programs. Importantly, Virginia did not expand Medicaid to low-income ABAWDs until 2019.

In contrast to the other two in-kind transfers, TANF provides cash assistance to very low-income families with children and imposes work requirements for participating parents. In Virginia, TANF operates through the Virginia Initiative for Education and Work program, which mandates employment-related activities and enforces a 24-month limit on continuous receipt (within a 60-month lifetime cap). Given its stringent income limits—below 41% FPL in Virginia—TANF reaches far fewer individuals than SNAP and Medicaid and typically serves single-parent families. At the same time, TANF recipients are "categorically eligible" for both SNAP and Medicaid and often participate in all three programs. Taken together, the design of Virginia's safety net during our analysis period targeted multiple forms of assistance toward households with children and, to a lesser degree, the elderly and disabled.

Across all states, enrolling in SNAP, Medicaid, and TANF generally involves an initial application process and then recertifying eligibility at periodic intervals. Yet, households often face barriers to enrollment. For example, information about program eligibility may be hard to come by (e.g., Daponte et al. 1999; Aizer 2007). Even if families learn they are eligible, they might face burdens associated with the application process—including navigating complex paperwork, providing supplemental documents, physically traveling to an office to complete face-to-face interviews, and facing social stigma (e.g., Moffitt 1983; Currie and Grogger 2001; Herd et al. 2013; Celhay et al. 2022; Giannella et al. 2024).

Prior to 2012, applicants to SNAP, Medicaid, and TANF in Virginia had to complete a 20-page initial application form as well as separate recertification forms for each program to continue receiving benefits. Notably, unlike many states at the time, Virginia's initial application already bundled all three programs on a single form. Since baseline barriers to multiple-program enrollment were likely lower in Virginia than in other settings, the potential gains from further simplifying access may be even larger where application processes are more fragmented. All applications had to be completed on paper and submitted in person or by mail/fax at an assigned field office. Critically, field office assignments are based on *county of residence*, and each county typically has a single office. In practice, many applicants relied heavily on in-person guidance from caseworkers at their field office to determine which programs they were eligible for and to navigate the complex application process (Cook and East 2025). As a result, these application processes may have been particularly burdensome for those living far from their assigned county office.

2.2 Data

We leverage longitudinal administrative microdata from Virginia spanning 2007–2022. The core data come from the Virginia Department of Social Services (VDSS), which administers SNAP, Medicaid, and TANF.[4] The VDSS data provide monthly records on program participation for the universe of recipients. Individuals maintain consistent anonymized identifiers over time, allowing us to track them longitudinally even as they move in and out of program receipt. Since we can link individuals across monthly SNAP, Medicaid, and TANF records, we construct detailed histories of both single and multiple program receipt.

A key advantage of the VDSS data is that they accurately capture all program recipients in Virginia. In contrast, data from household surveys—commonly used in the literature—substantially under-report program participation. As shown in Appendix Figure A.3, only 60–70% of SNAP and Medicaid recipients in the VDSS records are captured in the American Community Survey (ACS), and just 40–60% in the Current Population Survey (CPS) Annual Social and Economic Supplement.[5] This undercount is severe for multiple program recipients, with the ACS capturing half of those receiving both SNAP and Medicaid and the CPS capturing 40%. These comparisons underscore the value of using high-quality administrative data to study multiple program receipt.

We use program-specific case numbers to group together individuals who are members of the same program case (e.g., a single mother and child on TANF). To construct "households," we further group individuals appearing in overlapping cases within the same month. This innovation allows us to approximate households even in the absence of a master relationship file or census, although it is possible that some household members are excluded from a case and non-cohabiting members are included in a case. This approach enables us to link child-only cases (e.g., for TANF or Medicaid) to their parents so long as all such individuals are part of the same case for another program. VDSS records also include demographic characteristics such as age and gender, which permit us to infer household characteristics such as single- versus two-parent status, household size, and presence of children. Importantly, VDSS records also report the zip code and county of residence of each program recipient. Since zip codes sometimes lie in multiple counties and are served by multiple field offices, our geographic unit of analysis will therefore be a zip-county pair (which we sometimes refer to as a "zip code" for simplicity).

To supplement the VDSS records, we draw from several other data sources. First, for cer-

4. VDSS also administers child support, child care assistance, foster care, and refugee assistance. With the exception of child support, these programs are relatively small (see Appendix Figure A.1). We exclude child support from our analysis because it is ultimately a private transfer rather than a public benefit.

5. Our comparisons to survey data focus on the ACS and CPS rather than the SIPP, as we are conditioning on recipients in Virginia and the SIPP is not representative at the state level.

tain analyses aggregated at the zip-county level, we incorporate the Department of Housing and Urban Development (HUD) United States Postal Service (USPS) ZIP Code Crosswalk to define a complete frame of zip-county pairs (see Appendix B.1 for details). This enables us to bring in zip codes for which there may be no program recipients in a given month. We also collect historical addresses of field offices between 2007–2022 using archived versions of the VDSS website. These institutional data allow us to calculate the Haversine distance between each zip-county pair's centroid and its assigned field office. Because administrative burdens likely vary with proximity to field office, geographic distance will serve as a key source of identifying variation. We also bring in zip-level characteristics from the Decennial Census and American Community Survey (ACS), which are reported at the Census-generated zip code tabulation area (ZCTA) level and can be merged using zip-to-ZCTA crosswalks. These data provide zip- and time-varying covariates for our descriptive and causal analyses.

Finally, we link the VDSS microdata to individual-level administrative records from the Virginia Employment Commission (VEC), the Virginia Criminal Sentencing Commission (VCSC), and the Department of Juvenile Justice (DJJ). VEC records contain longitudinal quarterly earnings data, which can proxy for work ability and measure work history. VCSC and DJJ records detail individual interactions with the criminal justice system, including the timing and type of offenses and incarceration sentence. These rich data will be used to characterize the composition of program recipients in our targeting analyses.

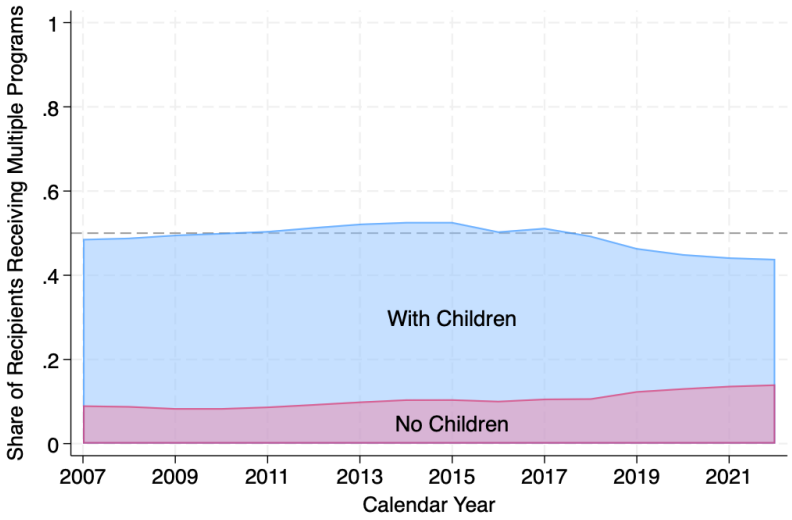
3 Four Facts about Multiple Program Receipt

In this section, we document novel descriptive facts on the breadth and nature of multiple program participation. These results provide motivating evidence on how recipients engage with multiple programs and lay the groundwork for our subsequent analyses.

Fact 1. *Roughly half of all program recipients enroll in multiple programs, with the vast majority concentrated in households with children.*

Figure 1 shows that approximately 50% of program recipients in Virginia participate in two or more programs in a given year, suggesting that multiple program participation is the norm rather than the exception. This share has remained largely steady over time, despite overall program participation in Virginia increasing from 15% of the population in 2007 to 26% in 2021 (Appendix Figure A.4). Among the possible multiple-program bundles of SNAP, TANF, and Medicaid, only two are taken up at non-trivial rates: SNAP/Medicaid and SNAP/Medicaid/TANF. This reflects the fact that TANF recipients overwhelmingly enroll in all three programs.

Figure 1: Share of Program Recipients Receiving Multiple Programs



Notes: This figure illustrates the share of program recipients (receiving either SNAP, Medicaid, or TANF) participating in two or more programs, subdivided by individuals in households with and without children.

Strikingly, the vast majority (80%) of multiple program recipients are part of households with children. Prior to 2019, the remaining recipients of multiple programs were likely elderly or disabled individuals. Following Virginia's expansion of Medicaid in 2019, the share of multiple program recipients in households without children rose modestly, reflecting the broadening of Medicaid eligibility among ABAWDs. In contrast, about 60% of *single* program recipients are in households with children. While this is a lower share, it remains large enough to raise the possibility that many of these single-program recipients are leaving additional program benefits on the table.

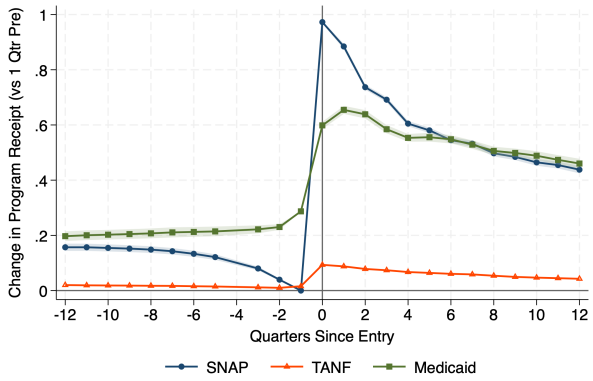
Fact 2. *Individuals often enter and exit multiple programs at the same time.*

While trends in aggregate participation are interesting in their own right, they mask important dynamics present among individual-level participation trajectories. Figure 2 presents event-study plots of program receipt surrounding entry into and exit out of SNAP, TANF, and Medicaid, focusing on those in households with children. While these are not causal estimates, they reveal clear patterns of dynamic interaction across programs, indicating that individuals often enroll in or leave multiple programs simultaneously. These results likely generalize to more programs at a national level given similar patterns found in the SIPP (Cook and East 2025; Wu and Zhang 2025).

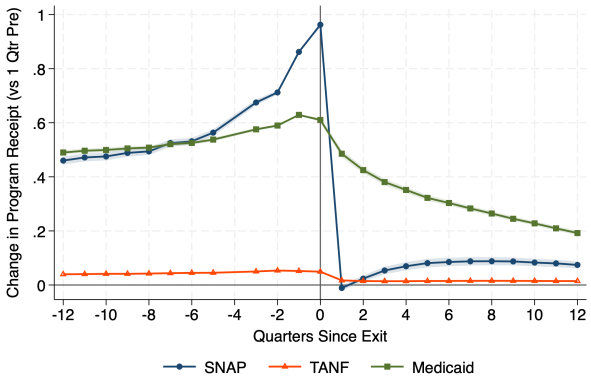
Examining first the entry margin, Panel (a) shows that 60% of SNAP entrants are also enrolled in Medicaid at the point of entry. For Medicaid in Panel (c), the interactions are less

Figure 2: Interactions Around Program Entry and Exit for Households with Children

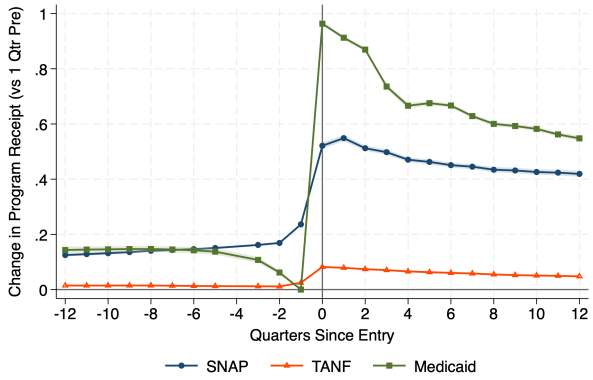
(a) SNAP Entry



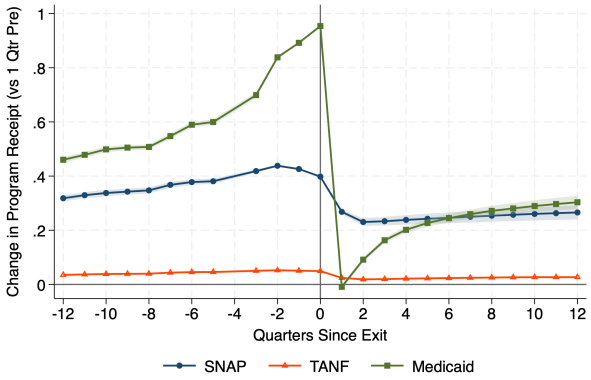
(b) SNAP Exit



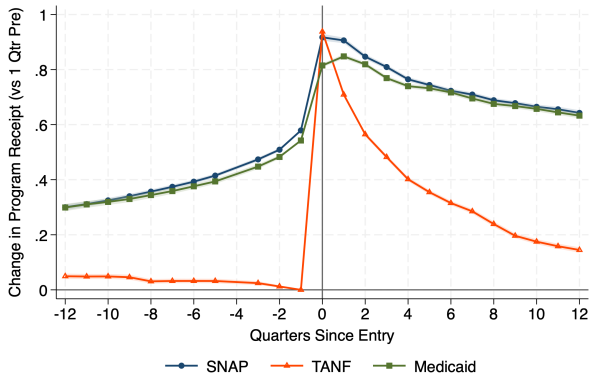
(c) Medicaid Entry



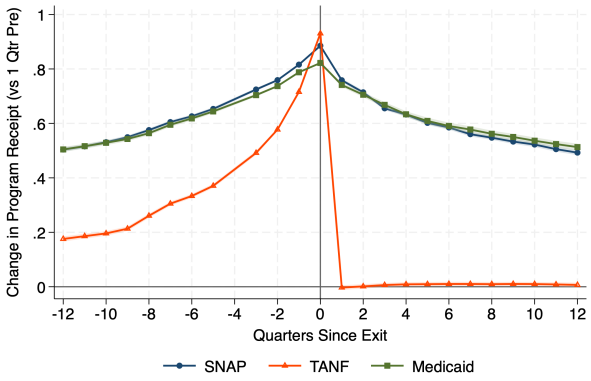
(d) Medicaid Exit



(e) TANF Entry



(f) TANF Exit



Notes: These figures plot event-study estimates of each program receipt on program entrants and exiters, focusing on individuals in households with children and centered around the quarter of entry and exit. We use a 10% random sample of entrants and exiters. Left panels condition on individuals who enter a given program and did not receive the program in the immediate prior quarter. Right panels condition on individuals who exit a given program and did not receive the program in the immediate subsequent quarter. Our sample consists of those entering or exiting between January 2010-December 2019, so that we can observe at least 12 quarters of participation data before and after the event. Shaded regions are 95% confidence intervals.

strong but still pronounced. Program interactions appear strongest for TANF, with Panel (e) revealing that roughly 80% and 90% of TANF recipients are also enrolled in Medicaid and SNAP, respectively—with participation increasing leading up to TANF entry. Notably, across all three programs, a substantial share of individuals (roughly half of TANF entrants and one-third of SNAP and Medicaid entrants) appear to *already receive* another benefit before enrolling in the focal program. The transition matrices in Appendix Figure A.6 corroborate the presence of "intensive-margin" movements: in any given month, 5% of single program recipients transition to a multi-program bundle.

On the exit margin, we continue to see inflection points in the receipt of other programs around exit out of a given program. Those exiting SNAP experience a steadier decline in Medicaid receipt, those exiting Medicaid tend to experience a more rapid decline in SNAP receipt, and both SNAP and Medicaid exiters also tend to exit TANF at the same time. Collectively, these dynamics suggest that enrollment is better understood not through isolated programs but through bundles of benefits that individuals move into and out of together.

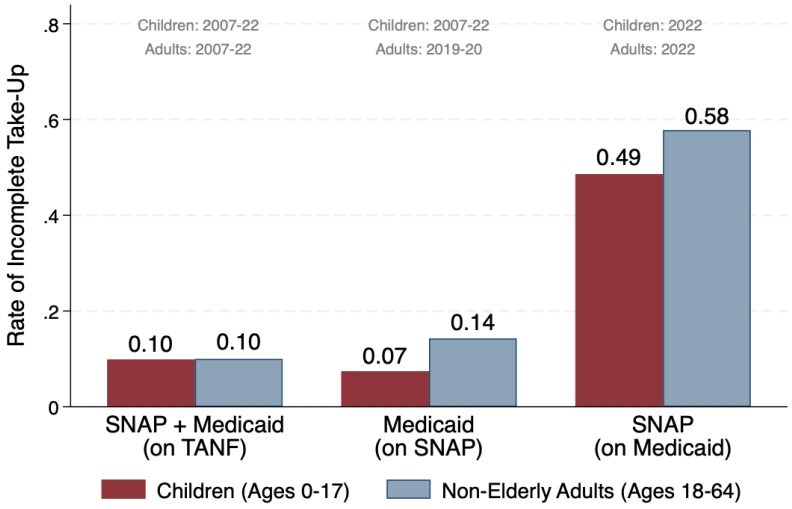
Fact 3. *Many recipients already attached to the safety net do not take up all of the programs for which they are likely eligible.*

A pervasive challenge in estimating program take-up is determining the size of the eligible population. Existing studies often rely on population-wide surveys like the ACS or CPS to impute eligibility, but such approaches face known limitations such as measurement error and misaligned reference periods. We take a different approach by examining take-up among subpopulations inferred to be eligible based on receipt of *another* program. This "revealed-eligibility" approach has been used in prior studies (e.g., Schmidt et al. 2024) and avoids the measurement challenges of estimating the full eligible population.

Using this approach, we find persistent evidence of incomplete take-up of multiple programs.[6] Examining first TANF recipients (who are categorically eligible for all three programs), the left set of bars in Figure 3 show that 10% of child and adult recipients do not participate in both of the other programs. Among children receiving SNAP (middle bars), 7% do not enroll in Medicaid despite being eligible. And for adults on SNAP, 14% do not enroll in Medicaid even after all low-income adults became eligible after the 2019 expansion. Most notably, after income thresholds were raised and asset tests were eliminated for SNAP in 2021, 49% of children and 58% of adults on Medicaid still remained unenrolled in SNAP

6. We use annual rather than monthly data to avoid conflating incomplete take-up with delays in transitioning to steady-state program bundles. This likely yields a lower bound on under-enrollment. Appendix Figure A.7 presents results requiring at least three or six months of participation in the conditioning program and finds similar patterns, suggesting that our main results are not driven by short-term timing differences.

Figure 3: Incomplete Take-Up of Multiple Programs (Conditional on Receiving a Program)



Notes: This figure plots conditional participation rates of different programs among children and adults who are inferred to be eligible based on their participation in other programs. The gray text notes the time periods used for the data, which differ due to changes in eligibility criteria. Individuals on TANF (both left bars) are categorically eligible for both SNAP and Medicaid and children on SNAP (red middle bar) are eligible for Medicaid throughout the time period. On January 1st, 2019, Medicaid expansion increased the income eligibility threshold for adults above SNAP (gray middle bar); on July 1st, 2021, SNAP expanded its income eligibility threshold beyond Medicaid (right bars).

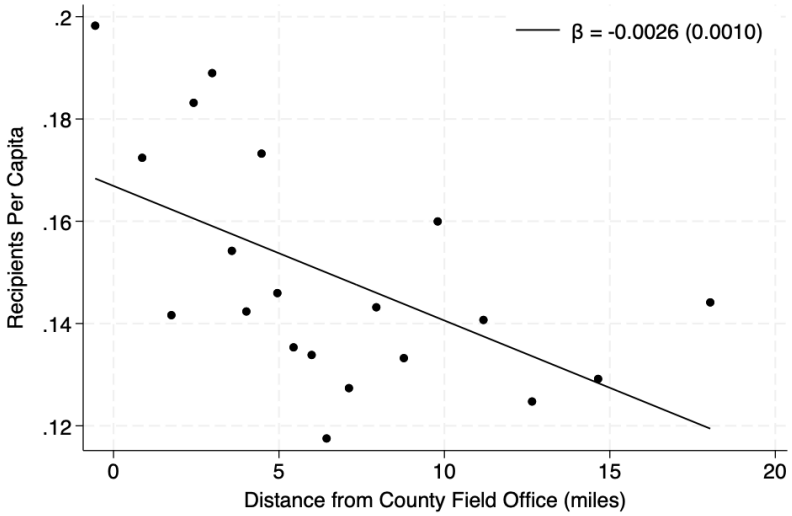
despite being eligible based on gross income and assets (right bars).[7] If these Medicaid recipients were eligible for only 20% of average SNAP benefits paid out, then this implies they left over half a billion dollars on the table. These magnitudes are especially striking given that all three programs are integrated in a common application.

Importantly, incomplete take-up along the intensive margin accounts for a substantial share of overall non-participation. In 2012, 43% of Medicaid-eligible children not enrolled in the program were estimated to already receive SNAP.[8] Appendix Figure A.8 provides further indirect evidence of this incomplete take-up, as 60% of single program recipients

7. Our data end in 2022, so these calculations are strictly based on a single year after Virginia expanded SNAP eligibility in 2021 via Broad-Based Categorical Eligibility. Thus, they may not reflect steady-state levels of incomplete take-up, although aggregate data suggest SNAP take-up did not markedly increase in 2023 and 2024. These rates also do not account for Medicaid recipients who would receive zero SNAP benefits because their net income (after deductions) exceeds the statutory limit, but little is known about how binding the net income test is among higher-income Medicaid recipients.

8. To calculate this share, we first use the overall Medicaid take-up rate for children in Virginia in 2012—88% according to Centers for Medicare & Medicaid Services (2013)—to back out the total number of Medicaid-eligible children from the observed number of child Medicaid recipients in the administrative data. Using our microdata, we then estimate the number of eligible children not enrolled in Medicaid who were receiving SNAP in 2012. Finally, we divide this number by the total number of Medicaid-eligible children not enrolled in the program.

Figure 4: Multiple Program Receipt Rates Decrease in Distance to County Field Office



Notes: This figure illustrates the association between zip-level multiple program participation rates and zip code distance to their assigned county field office using monthly data between January 2007 and December 2022. Participation rates are expressed in per capita terms using population estimates of zip code tabulation areas (ZCTAs) from the 5-year ACS, scaled to the zip-county level using the HUD-USPS crosswalk. Binscatters are weighted by any program receipt and control for ZCTA-level poverty rate, zip-level population share of Blacks and Hispanics, and county fixed effects (drawn from 5-year ACS estimates).

with children have no wages at all, suggesting many could qualify for additional benefits.

Fact 4. *Those living farther from their county field office participate less in multiple programs than those living closer, highlighting the spatial distribution of administrative burdens.*

As discussed in Section 2.1, prospective applicants living farther from their assigned field office may face greater administrative burdens in accessing programs. Figure 4 provides suggestive evidence for this hypothesis by illustrating the association between distance to office and per capita participation in multiple programs at the zip code level from 2007 to 2022. This binscatter controls for a variety of zip-level covariates to account for potential confounding factors that may be correlated with both distance and take-up. For comparison, Appendix Figure A.9 shows binscatters without any controls.

We find a strong negative gradient between distance to a local office and multiple program participation. Every additional mile is associated with a 0.26 percentage point decline in per capita receipt of multiple programs (p-value < 0.05), while the gradient for single program receipt is smaller and not statistically significant (Appendix Figure A.9). Given that programs share overlapping eligibility criteria and are bundled in the same application, the results suggest that more distant residents eligible for multiple programs face larger application frictions

that prevent enrollment. These patterns persist when examining intensive-margin take-up rates and are particularly pronounced for children (Appendix Figure A.10).

These results align with prior work showing that distance to administrative offices affects take-up in programs such as SNAP (Bitler et al. 2025), Social Security Disability Insurance (Deshpande and Li 2019), the Special Supplemental Nutrition Program for Women, Infants, and Children (Rossin-Slater 2013), and social assistance in Canada (Hicks 2024). Our paper differs in that we use this spatial variation as a source of identification to estimate the impacts of *reducing* administrative burdens tied to distance—rather than studying the direct effect of distance itself—on program participation and targeting.

Collectively, these facts reveal that multiple program receipt is common, with entry and exit often occurring jointly. Yet, incomplete take-up persists even among those already attached to the safety net and is more pronounced in zip codes farther from offices. To assess whether administrative burdens contribute to this incomplete take-up, we use a multiple-program framework to analyze a reform that simplified program enrollment processes.

4 Policy Reform and Empirical Framework

4.1 CommonHelp and the Reduction in Administrative Burdens

Motivated to improve safety net access and caseworker efficiency, Virginia launched the CommonHelp portal statewide in October 2012 at an up-front cost of $10.7 million (Pittman 2012). The platform centralized and digitized key steps of the enrollment process for various programs (including SNAP, Medicaid, and TANF), offering an alternative to the paper application process described in Section 2.1. Rather than eliminating in-person caseworker assistance, CommonHelp operated as a parallel online pathway. Consequently, unlike other reforms that increased burdens by removing human support (Wu and Meyer 2023), CommonHelp offered applicants more flexibility and autonomy in navigating the application process.

As summarized in Table 1, CommonHelp reduced burdens at *each* stage of the enrollment process. The first step—determining which programs to apply for—previously depended on applicant knowledge or caseworker guidance, which could result in households overlooking benefits they qualified for. CommonHelp instead introduced a 24/7 online screening tool that automatically recommended all eligible programs based on basic demographic and income inputs. This lowered information frictions, especially for those eligible for multiple programs, and may have reduced stigma by offering a more private alternative to in-person assistance.

In the next step—completing the application form—CommonHelp provided a step-by-

Table 1: Reductions in Administrative Burdens from the CommonHelp Reform

Step of Enrollment Process	Before October 2012: Paper Form Application	After October 2012: CommonHelp Portal	Admin. Burdens Likely Affected
1. Determine eligibility	Caseworker assistance (business hours only); own knowledge	Quick eligibility screening tool in online portal (24/7)	Information frictions, stigma
2. Complete application	Single 20-pg. form for all programs, but elect desired ones; technical language	Step-by-step screens customized based on elected programs; 7th-grade reading level	Hassle & cognitive burdens
3. Submit application	In-person/mail to assigned field office (1 per county)	Online submission	Geographic access
4. Provide supplemental documents	Submit in person/mail	Upload documents to portal (after January 2015)	Geographic access
5. Recertify	Separate paper forms in-person/mail	Online submission	Geographic access

step online interface tailored to the applicant's selected programs and written at a 7th-grade reading level.[9] This eased cognitive burdens and reduced the likelihood of errors compared to filling out the traditional 20-page paper form. Applications could then be submitted online, eliminating the need to physically deliver, mail, or fax the application to one's assigned county field office. Program recertifications were also simplified: users could complete online renewals via CommonHelp instead of separate paper forms for each program.

In January 2015, CommonHelp upgraded the online platform to allow applicants to upload digital copies of supplemental documents (e.g., proof of identity, residence, and earnings) rather than submit physical copies to their field office. This allowed households to complete the entire application process remotely, easing geographic barriers to accessing programs. In short, CommonHelp reduced administrative burdens throughout the application process, especially for those eligible for multiple programs and living farther from their offices.

Importantly, Virginia's adoption of CommonHelp was not an isolated case. As of 2024, 31 states have implemented integrated online application platforms for SNAP, Medicaid, and TANF, and 14 others have integrated at least two programs (Appendix Figure A.11a). The number of states pursuing such integration has grown steadily over time, with Virginia among the early adopters (Appendix Figure A.11b). While our analysis focuses on Virginia, the policy change we study has broader relevance across the country.

9. This reading level assessment is based on Code for America's Benefits Enrollment Field Guide.

4.2 Empirical Framework and Identification Strategies

A central challenge of identifying CommonHelp's impact is that it was launched nearly simultaneously across Virginia without an obvious control group. To address this, we develop a framework that exploits variation in how localities with differing baseline exposure to administrative burdens were affected by the reform. Specifically, those living farther from their field offices faced higher travel costs and reduced access to caseworker assistance, which we refer to as "geographic burdens."[10] These households likely experienced larger gains since CommonHelp reduced the need for in-person visits. At the same time, all households—including those close to their offices—faced baseline frictions from complicated paperwork, office congestion, and stigma. CommonHelp may have also reduced these "fixed burdens." Our empirical strategy combines two complementary research designs that assess reductions in each type of burden to recover CommonHelp's total effects.

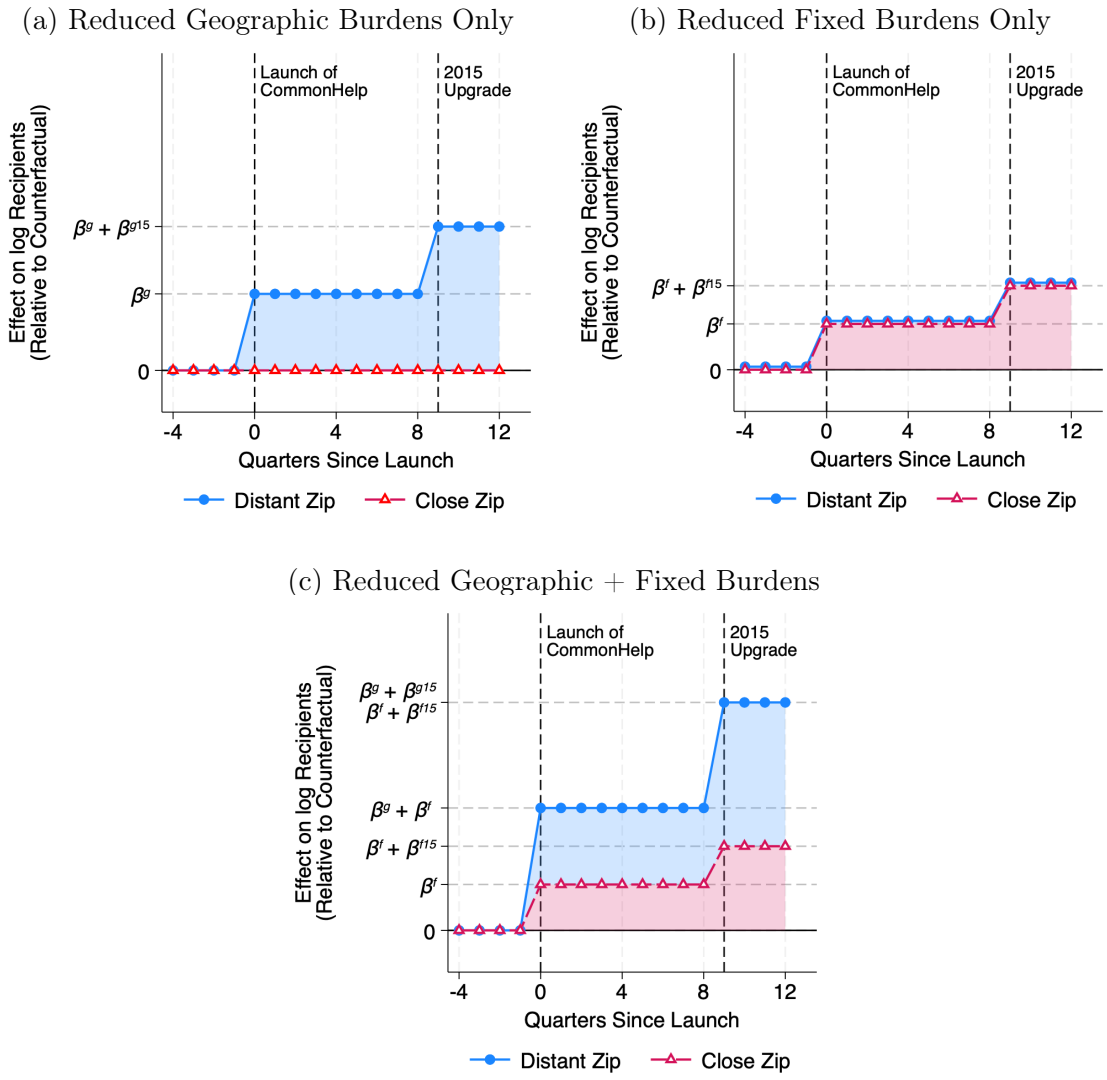
As a visual aid, Figure 5 presents stylized examples of how CommonHelp could have affected the growth of caseloads in zip codes near versus far from their field office. For simplicity, we focus here on static effects, assuming immediate awareness of CommonHelp and no learning costs in utilizing the platform.[11] Panel (a) considers a scenario in which CommonHelp reduced only geographic burdens, holding fixed burdens constant. In this case, participation would rise only in distant zip codes (blue lines), jumping by β^g after the 2012 launch and further by β^{g15} after the 2015 upgrade, while remaining flat in close zip codes (red lines). Assuming a stable difference in counterfactual caseloads between distant and close zip codes over time, then a standard difference-in-differences (DD) strategy would identify the effects of CommonHelp's launch and upgrade, represented by the shaded blue region. If treatment effects evolve over time (Appendix Figure A.12a) due to learning costs or snowballing peer effects in program participation (Dahl et al. 2014), then an event-study DD specification would appropriately capture these dynamics.

Panel (b) illustrates an alternative scenario in which CommonHelp only reduced fixed burdens affecting all zip codes while holding geographic burdens constant. This would lead receipt to increase (in both distant and close zip codes) by β^f and β^{f15} around CommonHelp's launch and upgrade, respectively—represented by the shaded red region below the red line. To identify the effects from reduced fixed burdens, we pursue a variant of a regression discontinuity (RD)-in-time design among zip codes close to their field office. This strategy has a similar spirit to an interrupted time series design; however, by identifying impacts only

10. Geographic burdens may also capture costs or barriers associated with the composition of populations living in more distant areas, such as differences in the salience of social stigma.

11. Appendix Figure A.12 shows analogous figures allowing effects to evolve dynamically over time, but the intuition is largely unchanged.

Figure 5: Expected Impacts on Program Receipt under Different Assumptions



Notes: This figure illustrates hypothetical scenarios of CommonHelp's impact on program receipt. The blue and red lines show CommonHelp's causal effect on receipt in zip codes with different field office distances over time. The blue and red shaded areas represent the effects identifiable from difference-in-differences and regression discontinuity designs, respectively. See main text for details.

at CommonHelp's launch and upgrade dates, we do not require stringent assumptions for modeling and extrapolating time trends (Baicker and Svoronos 2019).[12] We restrict to close zip codes to isolate fixed burdens, since farther zip codes experience a combination of both

12. Processing times for programs like SNAP, TANF, and Medicaid often have rapid turnarounds. In Virginia, SNAP mandates that applications be processed within 30 days, with expedited cases addressed within seven days. Similarly, TANF applications are required to be processed within 30 days, and Medicaid applications are generally within 45 days (Virginia Department of Social Services 2017). In addition, benefits in Virginia are provided retroactively to the month of application for eligible households (Virginia Department of Social Services 2024). Thus, even if applications were submitted and processed over several weeks following the launch of CommonHelp, the effective receipt of benefits would be backdated to the application date.

fixed and geographic burden reductions. If the fixed burden effects instead increase over time (Appendix Figure A.12b), then the RD design identifies a *lower bound* on caseload effects.

Panel (c) brings together the prior two panels by recognizing that CommonHelp likely reduced *both* geographic and fixed burdens, allowing us to identify the total effects of CommonHelp. The DD strategy would incorporate the relative effect in distant zip codes (the blue area), while the RD-type strategy would capture the component of the effect that is constant across all zip codes (the red area). Effectively, the latter strategy recovers the "intercept effect" purged in treatment-dosage DD designs with no untreated units (Callaway et al. 2024). In sum, the identified parameters from the DD and RD-type designs (β^g, β^{g15} and β^f, β^{f15}, respectively) provide the relevant ingredients for quantifying CommonHelp's total effect on recipient growth. Specifically, if the DD coefficients are expressed as linear per-mile growth parameters, $\bar{\beta}^g$ and $\bar{\beta}^f$ pool over the launch and upgrade effects over time, and D_z denotes the distance from zip code z to its assigned field office, then CommonHelp's impact on overall caseload growth in z is:

$$g_z \equiv \bar{\beta}^g \cdot D_z + \bar{\beta}^f. \tag{1}$$

Multiplying zip-level recipient counts by g_z and summing over all zips delivers (a lower bound of) the total number of recipients brought in by CommonHelp. We now turn to describing the DD and RD-type designs underlying this empirical framework.

4.3 Identifying Effects of Reducing Geographic Burdens

To estimate the effects of reducing geographic burdens, we implement a difference-in-differences (DD) design that compares changes in program participation between zip codes farther versus closer to their assigned field offices before and after the launch of CommonHelp in October 2012. Unlike conventional DD designs that rely on a binary treatment definition (e.g., "far" versus "close"), our specification treats distance as a continuous "dosage" variable to exploit the full range of variation across zip codes. Letting z index zip-county pairs, t index calendar year-month, and $q(t)$ denote t's quarter, we estimate the following event-study specification:

$$\log(Y_{zt}) = \sum_{\substack{-11 \leq k \leq 12, \\ k \neq -1}} \beta_k^g \cdot (D_z \times \mathbb{1}\{q(t) = k\}) + \alpha D_z + \delta_{c(z)} + \lambda_t + \varepsilon_{zt} \tag{2}$$

Here, Y_{zt} is the number of recipients in zip z and year-month t for a given combination of programs, D_z measures distance, and $\mathbb{1}\{q(t) = k\}$ indicates that year-month t belongs to event quarter k. County- and month-level fixed effects are denoted by $\delta_{c(z)}$ and λ_t. Our

analysis focuses on a window of 11 quarters before and 12 quarters after CommonHelp's launch—i.e., January 2010 through December 2015.[13]

Our primary outcomes of interest are log counts of individuals receiving (i) any program, (ii) multiple programs, and (iii) a single program.[14] In light of Fact 1, we focus on enrollment impacts among individuals in households with children, who comprise the vast majority of those eligible for and receiving multiple programs. We use logarithmized recipient counts (setting a lower bound of zero) because recipient counts in zip codes with differing population sizes do not increase by the same absolute level over time. To avoid estimation error from small zip codes, we restrict our sample to zip codes with an average of at least five monthly recipients pre-CommonHelp and then weight by these enrollment counts.[15] Standard errors are clustered at the county level to account for serial correlation within counties over time.

Our coefficients of interest in Equation (2) are $\{\beta_k^g \,|\, k \geq 0\}$, which measure CommonHelp's marginal effect (per mile) on program receipt growth k quarters after launch relative to $k = -1$, the omitted quarter. Based on our two-way fixed effects specification, β_k^g reflects a positively-weighted average across zip codes of reducing geographic burdens from an additional mile (Callaway et al. 2024). Since CommonHelp was rolled out nearly simultaneously statewide, we avoid complications arising in staggered DD designs (Goodman-Bacon 2021).

Our identifying assumption is that, absent the reform, more and less distant zip codes would have exhibited parallel trends in program participation.[16] This assumption may be violated if, for example, differential exposure to business cycles or lingering effects from the Great Recession affected farther versus closer zip codes differently. We probe these threats to identification in a battery of robustness checks. Moreover, the parameters $\{\beta_k^g \,|\, k < 0\}$ constitute our pre-trends test to help validate our parallel trends assumption. If there are no differential trends prior to CommonHelp, then we expect $\beta_k^g = 0$ for $k < 0$.

For reporting impacts in tables, we consider a simplified pre-post analog of Equation (2):

$$\log(Y_{zt}) = \beta^g \cdot (D_z \times Post_t) + \beta^{g15} \cdot (D_z \times Post15_t) + \alpha D_z + \delta_{c(z)} + \lambda_t + \boldsymbol{X}_{zt}\boldsymbol{\gamma} + \varepsilon_{zt}, \quad (3)$$

where $Post_t \equiv \mathbb{1}\{t \geq \text{October 2012}\}$ and $Post15_t \equiv \mathbb{1}\{t \geq \text{January 2015}\}$ indicate the post-launch and post-upgrade periods, respectively. The coefficients β^g and β^{g15} map to

13. We focus on 11 rather than 12 quarters prior to CommonHelp's launch to avoid overlap with the peak of the Great Recession (2007–2009) and because zip-level covariates are only available beginning in 2010.

14. Due to their rarity, we exclude individuals receiving TANF only or TANF with only one other program (see Appendix Figure A.4). Thus, we consider only five possible bundles: (i) no program, (ii) SNAP only, (iii) Medicaid only, (iv) SNAP and Medicaid, and (v) SNAP, Medicaid, and TANF.

15. The majority of the zip codes we drop are non-residential zip codes (based on the HUD-USPS files), and they collectively comprise less than 0.1% of all program recipients.

16. This is a "strong" parallel trends assumption (Callaway et al. 2024). This assumes that treatment effects do not vary in distance due to unobserved factors correlated with distance.

the labeled parameters in Panel (a) of Figure 5 and summarize the overall effects of CommonHelp's launch and upgrade on recipient growth from reducing geographic burdens. To improve precision, we control for time-varying zip-level characteristics $\boldsymbol{X}_{zt}$, such as poverty rate, unemployment rate, and demographics. Appendix B.2 provides additional details on these covariates and on the DD specification.

4.4 Identifying Effects of Reducing Fixed Burdens

To estimate the impacts from reducing fixed burdens, we implement a difference-in-regression discontinuity (DRD) design, focusing on zip codes in close proximity to their field office. This approach compares discontinuities in program receipt around CommonHelp's launch and upgrade ("treatment windows") to those around the same calendar months in other years ("control windows"). By differencing these discontinuities, we isolate the causal effect of CommonHelp from seasonal and reporting fluctuations in enrollment, which violate smoothness conditions required in standard RD designs (see Appendix B.2 for details).

Formally, let $Treat_t$ and $Treat15_t$ indicate the time windows around CommonHelp's 2012 launch and 2015 upgrade, respectively. Similarly, let $Control_t$ and $Control15_t$ indicate time windows around the same cutoff month as the launch (October) and upgrade (January) but in different years. Letting our running variable t be normalized to zero at each cutoff and restricting the sample to zip codes within $\bar{d}$ miles of their field office, we estimate:

$$\begin{aligned}\log(Y_{zt}) = &\left(\beta_1 \cdot \mathbb{1}\{t \geq 0\} + f_{1,pre}(t) \cdot \mathbb{1}\{t < 0\} + f_{1,post}(t) \cdot \mathbb{1}\{t \geq 0\}\right) \times Treat_t \\ &+ \left(\beta_0 \cdot \mathbb{1}\{t \geq 0\} + f_{0,pre}(t) \cdot \mathbb{1}\{t < 0\} + f_{0,post}(t) \cdot \mathbb{1}\{t \geq 0\}\right) \times Control_t + \epsilon_{zt}\end{aligned} \tag{4}$$

where $\{f_{k,pre}(t), f_{k,post}(t)\}_{k=0,1}$ are smooth functions of log program recipients on either side of the date cutoff. Here, $\beta^f \equiv \beta_1 - \beta_0$ captures the impact of CommonHelp's launch on caseload growth from reducing fixed burdens. The corresponding estimate for the 2015 upgrade, β^{f15}, can be obtained by modifying Equation (4) in the obvious way. These map precisely to the β^f, β^{f15} parameters labeled in Panel (b) of Figure 5.

We specify $f_0(\cdot)$ and $f_1(\cdot)$ as local linear functions of t with varying slopes on either side of the cutoff, and limit attention to a 6-month window around CommonHelp's launch and upgrade. To estimate seasonality effects, we pool all six possible control windows between 2009–2015 to maximize precision. We also restrict the sample to zip-counties within $\bar{d} = 1$ mile of their field office, which is roughly 10% of the median distance to field office and includes only 47 zip codes (4.8% of our sample). Zip-counties are again weighted by mean enrollment prior to CommonHelp, and standard errors are clustered at the county level.

5 Impacts on Safety Net Enrollment

This section presents our main findings on how CommonHelp affected program enrollment, focusing on individuals in households with children. We begin with visual evidence that underlies each empirical strategy and validates the key identifying assumptions. We then summarize regression estimates of the effects of reducing geographic and fixed burdens, and combine these to quantify CommonHelp's total impact on safety net participation. Finally, we lay out a range of robustness checks to assess the validity and stability of our results.

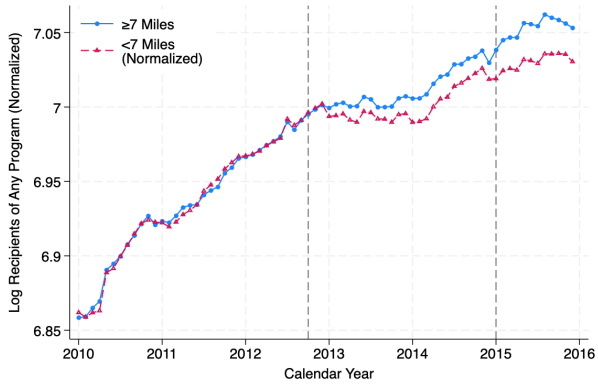
5.1 Visual Evidence of Enrollment Effects

Effects from Reduced Geographic Burdens (DD). We begin by presenting visual evidence that CommonHelp impacted program participation in zip codes farther from their field office. The left panels of Figure 6 illustrate raw trends in monthly log program recipients between zips more versus less than 7 miles from their office (which roughly splits our sample in half). Panel (a) presents trends for any program recipients and Panel (c) focuses on multiple program recipients; Appendix Figure A.13 shows trends for single program recipients. Log counts are normalized to be equal in September 2012 (the month prior to CommonHelp's launch) to facilitate comparisons. Leading up to CommonHelp, farther and closer zip codes possess strikingly concordant trends, even in months with sharp fluctuations. However, these trends diverge immediately after CommonHelp's launch, signaling that differences in trends can be attributed to CommonHelp rather than another time-varying shock. Differences between farther and closer zips continue to widen after the 2015 platform upgrade.

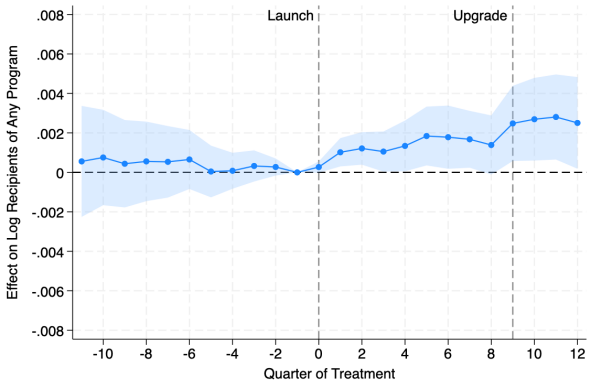
The right panels of Figure 6 show corresponding event-study estimates from the continuous DD specification, which exploits the full range of zip-to-office distances. Panel (b) shows that CommonHelp increased participation by 0.1% per mile immediately after launch and by 0.2% per mile after eight quarters. Effect estimates further spike to nearly 0.4% per mile following the 2015 upgrade. Though these per-mile estimates are modest in magnitude, they represent increases in caseloads on the order of tens of thousands of new recipients (as we quantify in Section 5.2). Reassuringly, the pre-trend is flat and possesses precise null estimates, reinforcing the plausibility of the parallel trends assumption. Panel (d) focuses on on multiple program receipt, which grew by 0.2% per mile immediately after CommonHelp's launch and by nearly 0.5% per mile after the upgrade. In contrast, single program participation initially increased but subsequently decreased to pre-treatment levels, suggesting substitution from single to multiple programs (Appendix Figure A.13).

Figure 6: Raw Trends and Event-Study Estimates Before and After CommonHelp's Launch

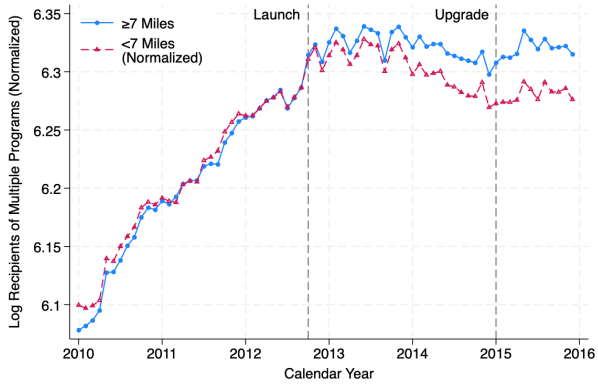
(a) Any Program: Raw Trends



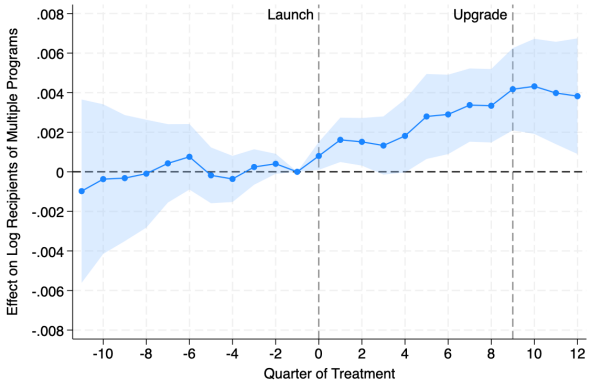
(b) Any Program: Event Study



(c) Multiple Programs: Raw Trends



(d) Multiple Programs: Event Study

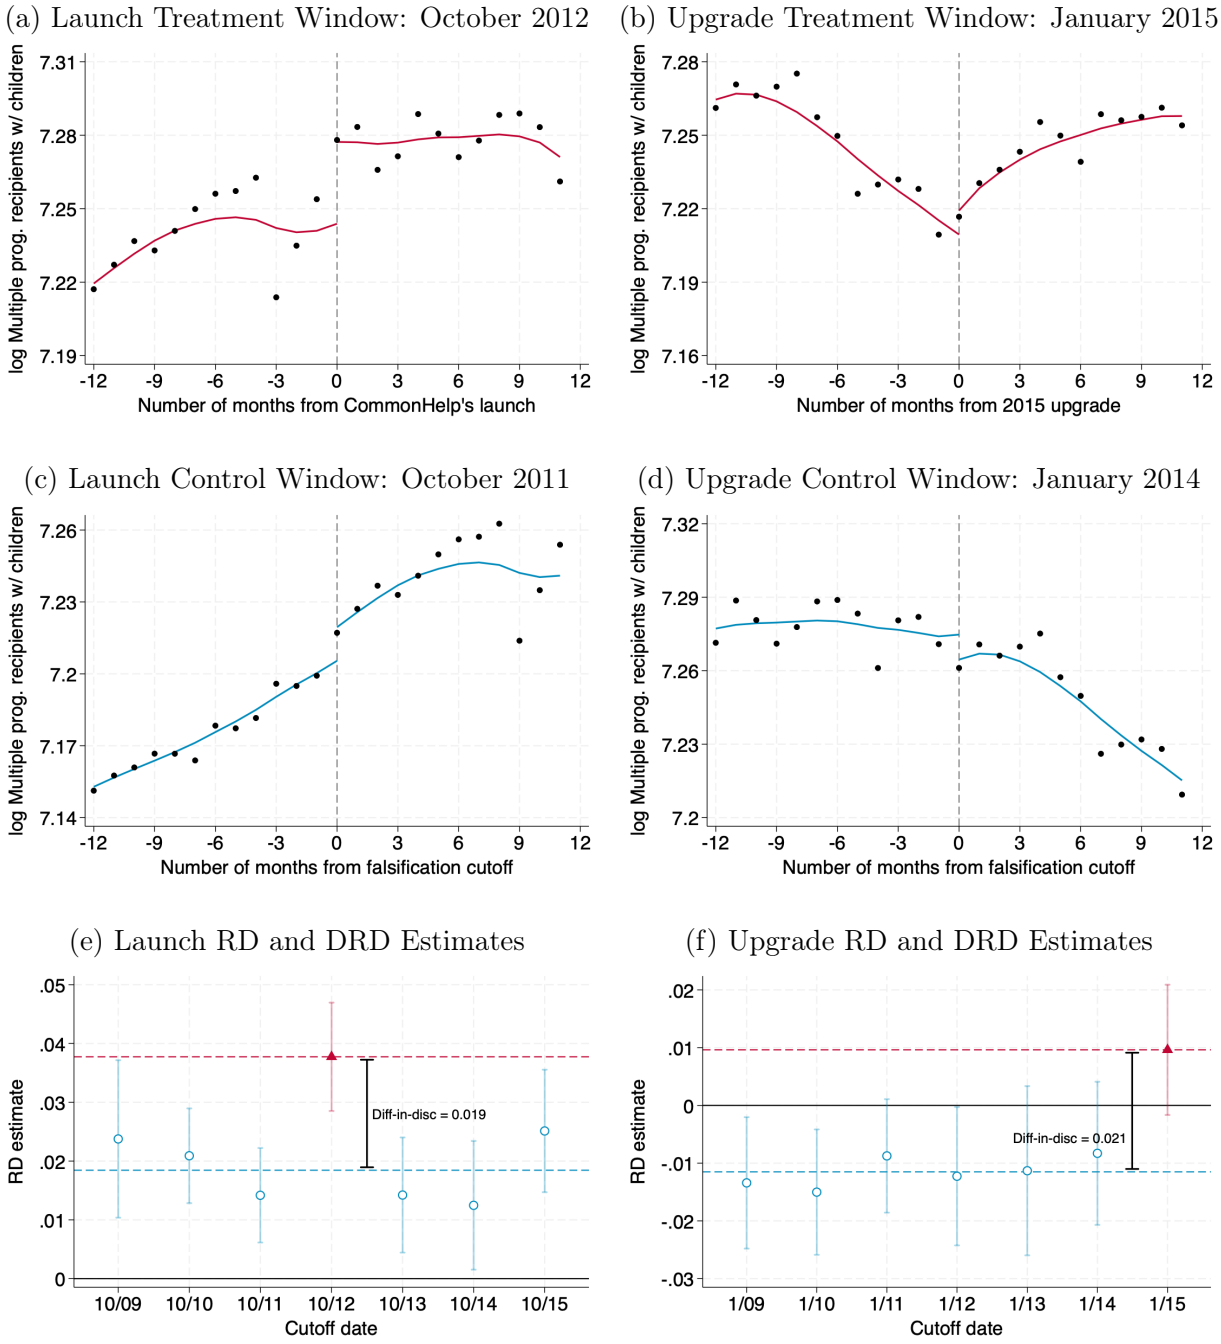


Notes: Left panels plot raw monthly trends in log program enrollment (among those in households with children) for the average zip code based on distance to assigned county field office. The level of zip codes less than 7 miles away is normalized to the level of zip codes more than 7 miles away in September 2012. Right panels illustrate event-study estimates of $\{\beta_k^g\}$ from Equation (2). Shaded regions are 95% confidence intervals based on robust standard errors clustered at the county level. In all panels, zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

Effects from Reduced Fixed Burdens (DRD). We now turn to our second empirical strategy, focusing on zip codes close to their office to isolate the effects of CommonHelp's reductions in fixed burdens. Panels (a) and (b) of Figure 7 illustrate discontinuities in multiple program receipt around the launch and upgrade for zips within 1 mile of their office (with Appendix Figure A.14 providing analogous plots for any program and single program receipt). The black dots are monthly log counts of multiple program recipients pooled across zip codes, while the red lines are nonparametric fits estimated separately on either side of the cutoff. Visually, there is a clear 1–4% jump in multiple program enrollment at both cutoffs.

While the results suggest that CommonHelp's launch and upgrade immediately impacted multiple program receipt in close zip codes, they could be an artifact of seasonality patterns

Figure 7: Visualization of Difference-in-Discontinuity Design for Multiple Program Receipt



Notes: The top and middle panels plot raw monthly enrollment (black dots) and nonparameteric fits (red and blue lines) around different windows. Bottom panels plot RD estimates in each estimation window between 2009–2015 and notes the corresponding DRD estimate. These estimates are based on local linear regressions with a 6-month bandwidth and triangular kernel around the cutoffs. Red triangular markers and red dashed lines correspond to estimates from the treatment windows; blue dashed lines correspond to the average estimates from the control windows. Left (right) panels correspond to windows around the launch (upgrade) cutoff. Sample is restricted to zip codes within $\bar{d} = 1$ mile of their field office. All estimates are weighted by zip-level pre-recipient counts and (inverse) distance from field office.

rather than the policy's true effects. To probe this, Panels (c) and (d) plot enrollment counts one year *prior* to the launch and upgrade dates. Some discontinuities are also present in these "control" windows, indicating that standard RD estimates would confound CommonHelp's impacts with seasonality effects. Panels (e) and (f) report the full set of RD estimates across all possible treatment and control windows between 2009 and 2015. The red triangular marker denotes the RD estimate from the treatment window ($\hat{\beta}_1$ from Equation 4), while the hollow blue circles plot RD estimates from control windows. The control-window estimates are similar in magnitude, suggesting that seasonality is a stable feature in the data. Strikingly, the treatment-window RD is the only estimate between 2009 and 2015 for which we can reject equality to the mean of the control windows, indicating that the observed discontinuity likely reflects a true effect of CommonHelp rather than pure seasonality. The DRD design differences the pooled mean over the control RDs ($\hat{\beta}_0$ from Equation 4) from the treatment RD estimate to purge confounding effects, showing that CommonHelp's reduction in fixed burdens increased multiple program receipt by 1.9-2.1%.

5.2 Summary Effects on Enrollment Growth and Caseloads

Effects on Enrollment Growth. Table 2 summarizes our main DD and DRD estimates of CommonHelp's effects on program enrollment through reductions in geographic and fixed burdens, respectively.[17] Panel A reports time-averaged DD effects corresponding to β^g and β^{g15} from Equation (3), as well as a pooled effect averaged over all twelve quarters after CommonHelp's launch. Overall, CommonHelp increased enrollment in any program by 0.20% per additional mile, driven by a 0.33% per mile increase in multiple program participation (p-value < 0.01). Effects are somewhat stronger following the 2015 platform upgrade. In contrast, single program participation remained essentially unchanged.

Panel B turns to DRD estimates of the effects of reducing fixed burdens in close zip codes. Overall, CommonHelp raised any program receipt by 0.92% in close zip codes, driven entirely by a 2.58% increase in multiple program receipt. Conversely, single program participation *declined* significantly by 1.17%, suggesting that some individuals who previously enrolled in a single program switched to multiple programs after CommonHelp made it easier to do so. Panel C combines the DD and DRD estimates to assess total enrollment effects, following Equation (1), for the average zip code with a caseload-weighted mean distance of $\bar{D} = 6.6$ miles. We estimate that CommonHelp increased total enrollment in any program by 2.21%, driven by a 4.73% increase in multiple program receipt. Meanwhile, single program participation fell by a statistically insignificant 1.05%.

17. Appendix Tables A.2 and A.3 show analogous tables for independent programs and different combinations of programs, respectively.

Table 2: Summary Effects on Program Enrollment Growth

	(1)	(2)	(3)	(4)	(5)
	Program type			Number of observations	
	Any	Multiple	Single	Zip codes	Zip-county-months
A. Effects from Reducing Geographic Burdens (DD: far vs. close zips)					
Overall effect (%/miles)	0.197**	0.329***	0.018	982	85,176
	(0.086)	(0.122)	(0.091)		
Launch effect (%/miles)	0.124	0.239**	0.015	982	85,176
	(0.080)	(0.113)	(0.074)		
Upgrade effect (%/miles)	0.238**	0.294***	0.011	982	85,176
	(0.094)	(0.098)	(0.120)		
Outcome mean (count/zip)	600.6	318.6	271.6	982	85,176
B. Effects from Reducing Fixed Burdens (DRD: close zips only)					
Overall effect (%)	0.923***	2.579***	−1.173**	47	4,200
	(0.276)	(0.596)	(0.473)		
Launch effect (%)	0.864***	1.929***	−0.622	47	4,200
	(0.246)	(0.536)	(0.474)		
Upgrade effect (%)	0.192	2.113***	−1.789***	47	4,200
	(0.288)	(0.423)	(0.507)		
Outcome mean (count/zip)	1,146.7	588.9	541.3	47	4,200
C. Total Average Effects (DD+DRD)					
Overall effect (%)	2.209***	4.727***	−1.053	982	85,176
	(0.657)	(0.957)	(0.758)		
Launch effect (%)	1.673***	3.488***	−0.525	982	85,176
	(0.606)	(0.838)	(0.708)		
Upgrade effect (%)	1.745***	4.028***	−1.716**	982	85,176
	(0.645)	(0.808)	(0.837)		
Outcome mean (count/zip)	600.6	318.6	271.6	982	85,176

Notes: Panel A reports estimates of β^g and β^{g15} from the DD specification (Equation 3). Panel B reports estimates of $\beta^f = \beta_1 - \beta_0$ and $\beta^{f15} = \beta_1^{15} - \beta_0^{15}$ from the DRD specification (Equation 4). Panel C reports total growth estimates g_z (Equation 1) setting $D_z = \bar{D} = 6.6$ miles, the average distance to field office. The first row of each panel reports the pooled effect estimate representing average point-in-time estimates between October 2012 and December 2015. Standard errors in parentheses are clustered at the county level.

Effects on Total Enrollment. Given these growth effects, how many individuals changed their safety net attachment due to CommonHelp? Appendix B.3 shows that *percent* changes in enrollment can be transformed into *counts* of additional recipients via:

$$\Delta Y \equiv \sum_{z \in \mathcal{Z}} \left(1 - \frac{1}{1+g_z}\right) \cdot \bar{Y}_{z,post}, \tag{5}$$

Table 3: Total Effects of CommonHelp on Program Enrollment

	(1)	(2)	(3)	(4)	(5)	(6)
	Any		Multiple		Single	
	ΔY	%	ΔY	%	ΔY	%
1. Total: $g_z = \bar{\beta}^g \cdot D_z + \bar{\beta}^f$	58,424***	2.206	63,885***	4.642	−13,015	−1.048
	(20,953)		(15,503)	[2.360]	(10,115)	[−0.481]
2. Geographic only: $g_z = \bar{\beta}^g \cdot D_z$	34,278**	1.283	29,101***	2.062	1,535	0.125
	(17,876)		(12,090)	[1.075]	(8,144)	[0.057]
3. Fixed only: $g_z = \bar{\beta}^f$	24,764***	0.923	36,208***	2.579	−14,587***	−1.173
	(8,084)		(9,202)	[1.338]	(6,128)	[−0.539]

Notes: Odd-numbered columns report estimates of the effects of CommonHelp on the number of program recipients ΔY (defined in Equation 5) under different choices of the total growth parameter g_z (defined in Equation 1). Standard errors are computed from 250 county-level block bootstrap replications and reported in parentheses. Even-numbered columns express effect estimates in terms of program caseload growth relative to the counterfactual caseload under no CommonHelp. Percentage point changes in the share of recipients that fall into a given category are reported in brackets in Columns (4) and (6). Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

where g_z is the percent change in caseloads for zip z and $\bar{Y}_{z,post}$ is the mean number of recipients in z after CommonHelp's launch. Table 3 reports estimates of ΔY for different choices of g_z in the odd-numbered columns, while even-numbered columns express these as percent changes in total enrollment growth. These estimates pool impacts from geographic and fixed burden reductions averaged over the full post-reform period (October 2012–December 2015).

In Row 1, we estimate that CommonHelp increased participation in any program by more than 58,000 individuals in households with children (a 2.2% increase). However, it is remarkable that roughly 64,000 individuals enrolled in multiple programs (a 4.6% increase)—exceeding the total effect on the number of any-program enrollees. To reconcile these figures, note that single program participation *fell* by approximately 13,000 recipients (−1.0%). Assuming CommonHelp weakly improved safety net access for all prospective applicants, these patterns indicate that CommonHelp not only drew new enrollees into the safety net, but also strengthened the *intensity* of safety net attachment by promoting transitions from a single program to multiple programs. Thus, analyzing any-program attachment alone would miss how CommonHelp reshaped the nature of safety net attachment among existing recipients. We quantify the magnitude of these "intensive-margin" movements in Section 6.

While these estimates reflect average enrollment growth, it is also useful to calculate the cumulative impact at a single point in time. Summing the separate launch and upgrade effects—which approximates a point-in-time estimate at the end of 2015—suggests CommonHelp brought in nearly 70,000 additional recipients of any program (Appendix Table A.4). This was driven again by a sharp rise in multiple program participation (+76,500) and

a substantial decline in single program participation (−32,000). These economically and statistically significant magnitudes indicate that burdens prevented many individuals from accessing benefits to which they assigned positive value.

To test the importance of combining the DD and DRD designs in our empirical framework, we examine the sensitivity of these estimates to including only geographic or only fixed burdens. For instance, Row 2 shows that ignoring fixed burdens would underestimate the overall number of recipients by 41%, as well as understate multiple program gains and overstate growth in single program participation. Row 3 shows that omitting geographic burdens would lead to more severe declines in single program recipients. The magnitudes of these biases underscore the importance of accounting for both types of administrative burdens to accurately evaluate CommonHelp's aggregate effects.

As discussed in Section 4.2, these estimates of any and multiple program enrollment effects are lower bounds if the reduction in fixed burdens grows over time. In contrast, it is difficult to interpret bounds for single program receipt since CommonHelp generated non-monotone responses along this choice margin. In particular, evolving fixed burdens (e.g., lower learning costs) may result in greater intensive-margin responses, causing *lower* single program receipt over time than implied by our DRD analysis. Our baseline results thus assume that the effects of fixed burden reductions are constant over time. Section 5.3 provides sensitivity checks examining alternative assumptions on fixed burden effect dynamics.

5.3 Robustness Checks

This subsection summarizes the battery of robustness checks we conduct to validate our empirical strategies.

Effects from Reduced Geographic Burdens (DD). To probe the validity of our DD strategy, we begin by examining threats to the parallel trends assumption. First, we find null effects after replicating our analysis using recipients of Unemployment Insurance—a program not administered by VDSS and thus unaffected by CommonHelp—as the outcome (Appendix Figure A.15). This suggests that our main results are not driven by differential economic trends or business cycles between farther and closer zip codes. Our estimates are also highly robust to the inclusion and choice of time-varying zip-level covariates (Appendix Figure A.16), which is consistent with the remarkably parallel trends in covariates between farther and closer zip codes (Appendix Figure and A.17). Collectively, these checks support the claim that exposure to local economic shocks is not correlated with distance to field office and is unlikely to confound our effect estimates.

One might also worry that the larger effects we estimate in farther zip codes stem not from geographic burdens per se, but from systematic differences in demographic composition that make these areas more responsive to administrative burdens. Appendix Figure A.18a shows that, prior to CommonHelp, zip codes farther from their field office had smaller populations, lower Black non-Hispanic shares, and higher shares of elderly, female, and low-income residents. To probe this concern, we re-estimate our event-study regressions after residualizing distance by these pre-treatment characteristics, thereby isolating variation in geographic burdens uncorrelated with local composition. As shown in Appendix Figure A.18, the resulting event-study estimates for both any program and multiple program receipt are nearly identical to our original results, suggesting that differences in demographic composition explain little if any of the observed treatment effects.

Finally, we test robustness across alternative samples and outcome definitions. Our estimates remain stable when excluding the most distant zip codes and adjusting measurement of distance of zip codes containing field offices, (Appendix Figure A.16). They are also robust to using inverse hyperbolic sine transformations and Poisson specifications, which guard against interpretation concerns with "log–like" specifications (Chen and Roth 2024). In addition, our findings hold using alternative definitions of distance including travel time by car or public transit drawn from Google API (Appendix Figure A.19). They are also robust to using binarized treatments based on discrete distance thresholds, with Appendix Figure A.20 showing that estimated effects rise roughly linearly with distance. This monotonic pattern is reassuring because it lessens the weighting and heterogeneity concerns that can arise in continuous-dose DiD settings (Callaway et al. 2024). Together, these checks reinforce the credibility and robustness of our DD design and thus the causal interpretation of our findings.

Effects from Reduced Fixed Burdens (DRD). We also conduct various falsification and specification checks to increase confidence in our DRD design. First, we estimate DRD effects around different placebo months before and after the true launch and upgrade dates, with control windows drawn from the same placebo month in other years. We focus on placebo months at least three months away from the true launch or upgrade month to avoid contamination bias in our estimation windows. Reassuringly, all placebo DRD estimates are statistically insignificant or of opposite sign (Appendix Figure A.21), suggesting our main results are not driven by spurious month-to-month noise or confounding calendar year shocks. We further check whether the main discontinuities reflect anomalous seasonality—especially around the launch, given its control windows exhibit positive but small effects—by estimating RD effects around every possible monthly cutoff between January 2009 and December 2015.

Many of the largest placebo RD estimates occur outside of January and October, indicating that large control RDs are not unique to our design. However, the true launch RD exceeds *all but one* placebo estimate, placing it in roughly the 99th percentile (Appendix Figure A.22). The largest RD estimate occurs in September 2012—just one month before the launch date—and is likely influenced by the true launch effect. Thus, it is unlikely that the observed launch effect is driven by chance.

We also evaluate sensitivity to modeling choices. Appendix Figure A.23 confirms that estimates are stable across different bandwidths and highly robust to varying the distance threshold $\bar{d}$ (in miles) that defines which zip codes are included in our sample. Unsurprisingly, estimates are slightly more unstable and imprecise for $\bar{d} \in [0.4, 0.6]$, but the point estimates remain remarkably similar to our preferred estimates at $\bar{d} = 1$. Appendix Table A.5 reports additional checks, including controlling for zip-level covariates and specifying $\{f_{k,pre}(\cdot), f_{k,post}(\cdot)\}_{k=0,1}$ as parametric linear or quadratic fits. To further guard against the possibility that our results are affected by CommonHelp's reduction in *geographic* burdens, one of our checks additionally applies a kernel that upweights zip codes that are *closer* to their field office. Across all specifications, our results remain qualitatively unchanged.

Finally, while our main DRD estimates are based on a highly selected sample—zip codes within 1 mile of their field office—our framework allows us to extend these results to *all* zip codes. This is because the DD design separately identifies the effects of reducing burdens that vary with distance above and beyond the fixed burden effects recovered by the DRD design. As an additional check, Appendix Figure A.24 plots DRD estimates across various distance bins, providing a nonparametric way of estimating the *total* immediate impacts of CommonHelp's launch and upgrade that capture effects from reducing both fixed and geographic burdens. These show near-linear gradients in treatment effects by distance, which align with the estimates implied by the event-study DD analysis. This consistency further validates our collective set of identification strategies.

Total Effects on Caseloads. The estimates of ΔY in Table 3 assume that the enrollment effects from fixed burden reductions remained constant over time—amounting to an extrapolation of the DRD estimates away from the launch and upgrade dates. To assess sensitivity to this assumption, Appendix Figure A.25 reports total effect estimates across a range of alternative assumptions of how effects from reduced fixed burdens evolved over time. To do this, we first estimate how much the effects from reduced *geographic* burdens grew over time based on our event-study estimates, denoted $\hat{g}_{\beta g}$. We then re-estimate ΔN where we assume fixed burdens grow at rate $\psi \cdot \hat{g}_{\beta g}$ for different choices of scalars ψ. When estimating ΔN for any and multiple programs—choice margins with plausibly monotone responses—we allow

$\psi \in [0, 2]$. For single programs—a choice margin with potentially non-monotone responses over time—we also allow the growth rate to shrink by allowing $\psi \in [-2, 2]$. Across a variety of plausible growth rate assumptions, the alternative total effect estimates lie within the 95% confidence interval of our main estimates, indicating that our main results are robust to different assumptions about the dynamics of fixed burdens reductions.

6 Impacts on Intensity of Safety Net Attachment

The total multiple-program enrollment effects in Section 5.2 reflect a composite of two distinct "response types": those who expanded participation by enrolling in additional programs (intensive-margin response) and those who entered the safety net for the first time (extensive-margin response). Quantifying the relative importance of each margin is policy relevant. It sheds light on the pathways through which individuals take up multiple programs when faced with a menu of program choices. The intensive margin is especially relevant given the meaningful rates of incomplete take-up documented along this margin in Fact 3. To unpack the relative contributions of both margins, we develop a potential outcomes framework and employ a partial identification approach that bounds the shares of each response type.[18] To our knowledge, this analysis represents the first attempt in the social safety net literature to quantify intensive- versus extensive-margin responses, as prior work has traditionally focused on single programs where responses can only occur on the extensive margin.

6.1 Potential Outcomes Framework

Let P_{it} denote individual i's program participation status at time t, taking values in the set $\mathcal{P} = \{n, s, m\}$, where n denotes no program receipt, s denotes single program receipt, and m denotes multiple program receipt. For each possible status $p \in \mathcal{P}$, define the indicator variable $P_{p,it} \equiv \mathbb{1}\{P_{it} = p\}$. Let $k \in \{0, 1\}$ indicate the presence of CommonHelp, and define $P_{it}(k)$ as the potential status under state k. Define $P_{p,it}(k)$ for $k \in \{0, 1\}$ analogously. Finally, define $\omega_{(p_0,p_1)} \equiv \Pr(P_{it}(0) = p_0, P_{it}(1) = p_1)$, the share of individuals who choose program status p_1 under CommonHelp but would have chosen p_0 absent CommonHelp. Our aim is to (partially) identify the shares of the three latent groups with positive responses to

18. While a simpler approach would be to estimate effects on groups defined by their program participation prior to CommonHelp, this would yield less informative estimates because (i) substantial churn in program receipt would lead to misclassification error in the groups, leading to attenuation bias, and (ii) participation spells are often short, making it difficult to credibly assess impacts three years later (see Figure 2). By relying on a counterfactual framework that accounts for what individuals would have done absent the reform, we more cleanly identify who was induced to change behavior by CommonHelp.

CommonHelp: $\omega_{(n,s)}$, $\omega_{(n,m)}$, and $\omega_{(s,m)}$. Doing so enables us to parse out whether CommonHelp expanded participation along (1) the *extensive margin*, by drawing in individuals who would have otherwise been completely detached from the safety net ($\omega_{(n,s)}$ and $\omega_{(n,m)}$) or (2) the *intensive margin*, by helping those who would have been partially attached to a single program access additional programs ($\omega_{(s,m)}$).

To make progress, we invoke the following monotonicity assumption: $P_{it}(0) \subseteq P_{it}(1)$, for each individual i and time t (see Appendix Figure A.26 for a visualization). In other words, monotonicity assumes that CommonHelp did not cause *reductions* in the number of programs individuals enroll in. This assumption would be violated if CommonHelp displaced high-touch in-person assistance from field offices that could have facilitated receipt of more programs. Nevertheless, the absence of statistically significant declines in program receipt across our DD and DRD designs suggest that such violations are likely minimal in practice.

Finally, for simplification, we define our baseline population as *recipients* in the post-CommonHelp period. This expresses $\{\omega_{(p_0,p_1)}\}$ as shares of the *recipient population* and excludes non-responders with no program receipt from consideration ($\omega_{(n,n)} = 0$).

6.2 Partial Identification

We start by scaling the change in recipients from Equation (5) by our base population of post-CommonHelp recipients, which yields the percentage-point change in the share of recipients under program status $p \in \mathcal{P}$ after CommonHelp:

$$\tilde{\beta}_p \equiv \frac{\Delta Y_p}{\sum_{z \in \mathcal{Z}} \bar{Y}_{z,post}}, \tag{6}$$

where ΔY_p is the number of new recipients with participation status p and $\sum_{z \in \mathcal{Z}} \bar{Y}_{z,post}$ is the total caseload of recipients after CommonHelp. This allows us to directly map these shares—whose estimates are reported in brackets in Table 3, Columns (2) and (4)—to the underlying shares of responder types $\{\omega_{(p_0,p_1)}\}$. Appendix B.4 shows that CommonHelp's percentage-point effect on a given program status is simply the mass of individuals moving into that program status less the mass moving out of it:

$$\begin{aligned} \tilde{\beta}_s &= \omega_{(n,s)} - \omega_{(s,m)} \\ \tilde{\beta}_m &= \omega_{(n,m)} + \omega_{(s,m)}. \end{aligned} \tag{7}$$

Notice that $\tilde{\beta}_s < 0$ would imply that more individuals shifted from single to multiple program receipt than entered into a single program from no receipt.

At this stage, we have two equations (for $\tilde{\beta}_s$ and $\tilde{\beta}_m$ above) but three unknowns: $\omega_{(n,s)}$,

Table 4: Estimated Ranges of Shares of Intensive vs. Extensive Margin Responses

	(1)	(2)	(3)
	Intensive-Margin, Multiple: $\omega_{(s,m)}$	Extensive-Margin, Multiple: $\omega_{(n,m)}$	Extensive-Margin, Single: $\omega_{(n,s)}$
Response share (p.p.)	[0.481, 2.360]	[0, 1.879]	[0, 1.879]
Assuming $\omega_{(n,m)} = 0$ (p.p.)	2.360	0	1.879
Assuming $\omega_{(n,s)} = 0$ (p.p.)	0.481	1.879	0
% Total responses	[20.372, 55.671]	[0, 79.628]	[0, 44.329]

Notes: This table reports the bounds defined in Equation (8), based on the estimates reported in the even-numbered columns of Table 3, in square brackets. Each row also reports values of each share under different assumptions and expresses estimates as percent shares of total responses.

$\omega_{(n,m)}$, $\omega_{(s,m)}$. As such, we cannot point identify the shares of each responder type, but we can *partially identify* them through bounds.[19] Fortunately, we can potentially tighten these bounds by utilizing the adding-up constraint $\sum_{(p_0,p_1)} \omega_{(p_0,p_1)} = 1$ and the shares of recipient non-responders (e.g., $\omega_{(s,s)}$). Given $\omega_{(p_0,p_1)} \geq 0$ and Equation (7), Appendix B.4 derives sharp bounds on each responder-type share:

$$\begin{aligned} \max\{0,\ \tilde{\beta}_s\} \ &\leq\ \omega_{(n,s)} \ \leq\ \min\{1+\tilde{\beta}_s,\ \tilde{\beta}_s+\tilde{\beta}_m\} \\ 0 \ &\leq\ \omega_{(n,m)} \ \leq\ \min\{\tilde{\beta}_m,\ \tilde{\beta}_s+\tilde{\beta}_m\} \\ \max\{0,\ -\tilde{\beta}_s\} \ &\leq\ \omega_{(s,m)} \ \leq\ \min\{1-\tilde{\beta}_s,\ \tilde{\beta}_m,\ \theta_{(s,s)}\}, \end{aligned} \tag{8}$$

where $\theta_{(s,s)}$ is an expression defined in Appendix B.4. Note that depending on the sign of $\tilde{\beta}_s$, either $\omega_{(n,s)}$ or $\omega_{(s,m)}$ will achieve an informative (i.e., strictly positive) lower bound.

6.3 Results

Table 4 utilizes the percentage-point effect estimates on single and multiple recipient shares from Table 3 to estimate bounds on the underlying responder-type shares $\omega_{(n,s)}$, $\omega_{(n,m)}$, and $\omega_{(s,m)}$. Because we estimate an overall *negative* effect on single program recipient shares, we achieve an informative lower bound on the share of intensive-margin multiple program responders. In other words, we can definitively conclude that CommonHelp's impact came in part from individuals who would have enrolled in one program but, due to the reform, accessed *additional* benefits they would have left on the table.

Based on overall effects pooled over the analysis period, we estimate that the share of intensive-margin responders ranged between 0.48–2.36 percentage points, comprising 20–

19. Appendix B.4 proves that $\tilde{\beta}_s$ and $\tilde{\beta}_m$ exhaust all information about the underlying response shares.

56% of all responses to CommonHelp.[20] Appendix Table A.6 shows that this share grew to 1.16–2.78 percentage points (42–63% of all responses) after the 2015 upgrade. For extensive-margin responders, the upper bounds suggest they could make up a meaningful share of new multiple program recipients, but the zero lower bounds leave open the possibility that most movement toward multiple programs occurred on the intensive margin. These results suggest that platforms like CommonHelp can help reduce incomplete take-up of additional programs among *existing* recipients (Fact 3). Under additional assumptions, we estimate that even the lower bound of CommonHelp's intensive-margin effects closed 20–42% of this take-up gap among eligible single program recipients.[21]

The informative bounds on intensive-margin multiple program responders are driven by the significant reductions in single program receipt following the decrease in fixed burdens uncovered by our DRD analysis. In contrast, the DD results show modest, statistically insignificant increases in single program receipt from reducing geographic burdens. Together, these findings suggest that intensive-margin responses should be more prominent in areas where fixed burdens are more binding—i.e., closer to field offices where there are stronger baseline ties to the safety net. Appendix Figure A.27, Panel (a) verifies this prediction by plotting bounds on $\omega_{(s,m)}$ within one-mile distance bins.

7 Targeting Analysis with Multiple Programs

Having established that CommonHelp bolstered multiple program enrollment—especially along the intensive margin—we now examine the socioeconomic composition of individuals induced into multiple program enrollment. Identifying *who* changes safety net attachment from reforms that reduce burdens to multiple programs is central to understanding the targeting of the safety net and the efficient level of enrollment costs.

To conduct this analysis, we link program recipients to administrative earnings and adult/juvenile crime records to measure pre-determined characteristics. A key strength is that we leverage a long time horizon, with earnings measures calculated over a three-year window prior to CommonHelp's launch and crime measures based on fourteen years prior.

20. These percentage ranges come from dividing the estimates from Table 4, Column (1) under different assumptions of $\omega_{(n,m)}, \omega_{(n,s)}$ by their row sums; e.g., the lower bound of $\omega_{(s,m)}$ based on the overall effects is $0.481 \div (0.481 + 1.879) \times 100\% \approx 20\%$.

21. To approximate the size of the intensive-margin take-up gap, we begin with the number of individuals receiving only a single program in September 2012 (pre-CommonHelp), and divide by an assumed intensive-margin take-up rate ranging from 86–93% (which reflects the estimates in Fact 3 that are representative of steady-state levels). We then multiply the lower bound share of CommonHelp's effect attributable to intensive-margin movement (20%) by the total number of induced multiple program recipients (63,885, from Table 3), and divide by the estimated range of overall intensive-margin non-participants.

This reduces the scope of measurement error from transitory fluctuations, providing more reliable proxies of well-being and self-sufficiency. We aggregate earnings at the household level, where households are based on linkages across program assistance units. Our analysis focuses on a composite "low-socioeconomic status (SES)" indicator, defined as having either zero household earnings between 2009-2011 or any criminal history. We also conduct our analysis using this indicator's sub-components and alternative measures of SES.[22]

We begin by summarizing results from a more standard approach in targeting analyses that examines *overall* changes in the observable characteristics of multiple program recipients. We then extend the scope of targeting analyses by comparing characteristics between intensive- versus extensive-margin responders that underlie the overall estimates. Finally, we discuss the implications of introducing multiple margins of selection from our multiple-program framework on evaluating targeting efficiency.

7.1 Targeting on Observables

We begin from a more conventional perspective on targeting by examining whether CommonHelp drew in multiple program recipients with observable indicators of socioeconomic disadvantage. We adapt our DD and DRD models to be at the individual level, using repeated monthly cross-sections of program recipients in order to assess changes in recipient composition. See Appendix B.5 for details on this empirical analysis.

Appendix Table A.7, Column (3) pools estimates from the adapted DD and DRD models to report the *average* targeting effects—the parameters typically reported in the literature—among all multiple program recipients. On average, CommonHelp's impact on the composition of multiple program recipients was small. The share of low-SES recipients increased by a statistically insignificant 0.48 percentage points (1.4%), with little to no significant changes in employment, earnings, or criminal history. Although the average targeting effects appear small, the underlying DD and DRD estimates (reported in Columns 1 and 2) reveal considerable heterogeneity in effects. Specifically, reductions in fixed burdens tended to increase the share of low-SES multiple program recipients, whereas reductions in geographic burdens had the opposite impact. To visualize this phenomenon, Appendix Figure A.27, Panel (b) plots CommonHelp's impacts on the share of low-SES recipients by distance to field office. The share of low-SES multiple program recipients rose by a statistically significant 2% in the *closest* zip codes, but declined in farther zip codes to statistically insignificant magnitudes.

22. While federal welfare reform imposed a lifetime SNAP ban for drug felony convictions starting in 1996, Virginia lifted this ban in 2005 (well before our analysis period) for all drug felonies except trafficking convictions. Moreover, our criminal justice data include all offenses (misdemeanors and felonies), not just serious convictions that might affect program eligibility. Thus, any mechanical relationship between criminal justice involvement and program enrollment should be minimal in our setting.

These results provide two main takeaways. First, the opposing targeting effects from reducing geographic versus fixed burdens explain the negligible aggregate targeting effects. This highlights how average effects, while typical in the literature, may obscure heterogeneity in how burdens impede access within the population. Second, CommonHelp disproportionately improved the targeting profile of multiple program participants closest to field offices—the same areas where intensive-margin responses were found to be strongest. This suggests an association between responder types and SES, which we now investigate more directly.[23]

7.2 Heterogeneous Targeting Effects by Selection Margin

To link our targeting analysis with responder types, we conduct a heterogeneity analysis that re-estimates bounds on responder-type shares *separately* by SES group. Figure 8, Panel (a) shows that intensive-margin responses were much more common among low-SES individuals. Based on overall effects, intensive-margin response shares ranged between 47–65% for the low-SES group—higher than the 20–56% range for the full sample in Table 4. Meanwhile, for high-SES individuals (defined as the complement of low-SES), we cannot reject that intensive-margin responses were zero. In contrast, Appendix Figure A.29 shows that extensive-margin responses largely lack informative lower bounds for both SES groups.

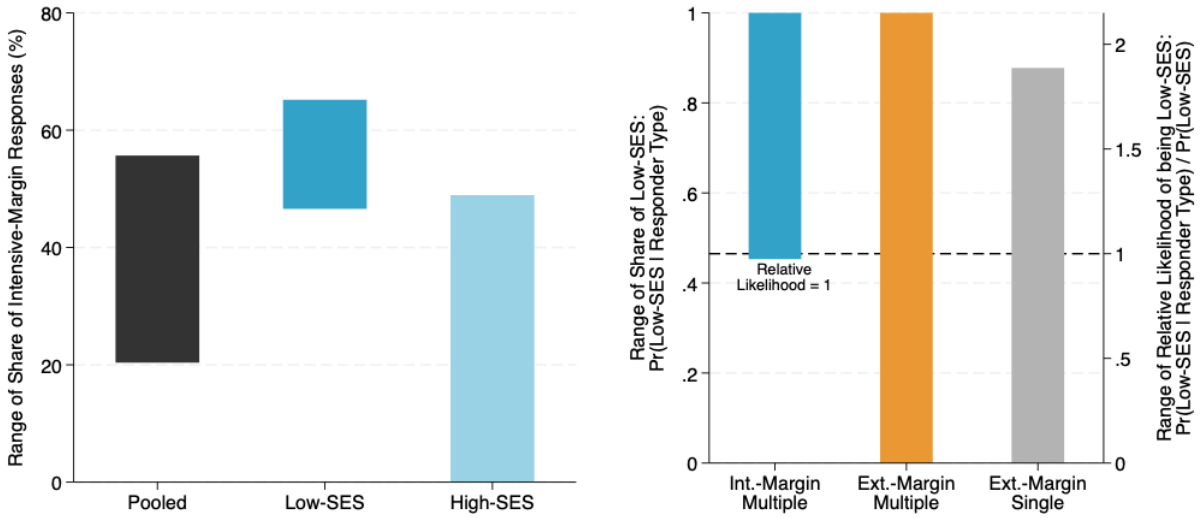
To facilitate a more interpretable analysis, we apply Bayes' rule to convert the range of probabilities that a given SES group responds along different margins into a range for how likely each responder type is low-SES. This approach is similar in spirit to methods used to characterize "compliers" in instrumental variables analyses (Abadie 2003); see Appendix B.4 for details. Panel (b) illustrates the resulting probability ranges by responder type. The figure reveals two striking insights. First, the probability range over which an *intensive*-margin multiple program responder is low-SES possesses large and informative lower bounds (blue bars). For the overall effects, between 45–100% of intensive-margin responders were low-SES. Since low-SES individuals make up 47% of all recipients, this implies a relative likelihood of 0.97–2.15, meaning low-SES individuals were at least as—if not far more—likely to respond along the intensive margin as the average individual. Second, the probability range over which an *extensive*-margin responder is low-SES possesses wide, uninformative bounds (orange and gray bars). In other words, we cannot reject that no low-SES individuals responded along the extensive margin to either single or multiple programs.

23. Appendix Table A.8 provides an alternative, indirect comparison by reporting targeting on observable estimates on *any program* as well, which captures strictly extensive-margin responses. CommonHelp tended to bring additional individuals with higher-SES backgrounds into any program relative to multiple programs.

Figure 8: Targeting on Response Types and Socioeconomic Status

(a) Share of Intensive-Margin Responses, by SES Group

(b) Share of Low-SES Recipients, by Response Type



Notes: Panel (a) plots the range of share of intensive-margin responses based on Equation (8) for the pooled sample (from Table 4) and separately by socioeconomic status (SES) group. Panel (b) plots estimated range of probabilities that different responder types are low-SES. The left y-axis expresses these as raw probabilities; the right y-axis expresses these as relative likelihoods by rescaling by the share of low-SES individuals. Except for the left bar in Panel (a), samples are restricted to the population of program recipients where SES is properly measurable; see notes of Table A.7 for details.

7.3 Implications of Multiple Programs on Targeting Efficiency

That intensive-margin responders are likely more disadvantaged than extensive-margin ones carries important insights for our understanding of the safety net's targeting efficiency. Based on neoclassical theory, burdens in the enrollment process can serve as a screening device that improves the targeting of social programs: individuals who self-select into program participation are more disadvantaged since they are willing to overcome these burdens (Nichols and Zeckhauser 1982; Kleven and Kopczuk 2011). For example, recent work by Alatas et al. (2016), Finkelstein and Notowidigdo (2019), and Rafkin et al. (2025) empirically support this theory for a variety of programs; meanwhile, Deshpande and Li (2019) and Homonoff and Somerville (2021) find that office closures and recertification burdens can worsen targeting.

Our findings suggest that the story becomes more nuanced when evaluating targeting efficiency from a multiple-program framework, which introduces *multiple* margins of self-selection. In particular, intensive-margin responders would have still revealed need, absent CommonHelp, by self-selecting into a single program. In accordance with the neoclassical theory of burdens, they are more disadvantaged than extensive-margin responders who, absent CommonHelp, would have remained fully detached from the safety net. Yet, these

intensive-margin responders are also likely more disadvantaged than those self-selecting into multiple programs even absent the reform (i.e., "multiple-program always-takers"). These patterns suggest that administrative burdens prevented more intensive attachment to the safety net among the most vulnerable. Thus, ordeals may not just screen out relatively more advantaged individuals from *any* safety net attachment, but also cause the most disadvantaged *among existing participants* to leave additional benefits on the table.

These findings suggest a potentially counterintuitive policy implication: the safety net's targeting could be improved by focusing on existing single-program recipients. This is especially practical from a policy implementation perspective, since policymakers can more readily intervene among populations they they already serve. This complements recent evidence from Wu and Meyer (2023), who find that recipients who drop off at recertification are disproportionately disadvantaged due to dynamic selection over time. In our context, recertification may instead act as an opportunity to improve targeting, especially when delivered through a simplified, integrated platform like CommonHelp that allows recipients to realize their full eligibility. Collectively, our results indicate how a multiple-program framework can shed new insights for evaluating targeting efficiency of the safety net as a whole.

8 Welfare Analysis with Multiple Programs

As shown in prior sections, CommonHelp increased participation in multiple programs among those in households with children. Yet, formal evaluations of safety net reforms often focus on a single program and analyze its social costs and benefits in isolation. Given that programs today are deeply interconnected, it is natural to ask how sensitive welfare analyses are to accounting for multiple programs. To explore this, we conduct a Marginal Value of Public Funds (MVPF) analysis of CommonHelp (Hendren and Sprung-Keyser 2020), showing the degree to which the MVPF might be misstated if a researcher has data for only one program at a time and thus fails to account for joint participation in SNAP, Medicaid, and TANF.

Conceptually, the MVPF is defined as the ratio of a policy's willingness-to-pay (WTP) to its net fiscal cost (inclusive of fiscal externalities from behavioral responses). Both WTP and fiscal costs generally rise with the number of programs included in the analysis. For instance, time savings from CommonHelp's simplified application process increase when considering additional program recipients, while benefit disbursements also grow. As a result, whether the MVPF rises or falls with the inclusion of multiple programs depends on how the changes in WTP and fiscal cost compare to each other. Because different modeling assumptions govern the formulation of the WTP and fiscal externalities, we focus on *differences* between MVPFs when considering a single program versus all programs. This allows us to assess the

sensitivity of MVPF calculations to the inclusion of multiple programs, independent of any particular modeling choice. We also compare our estimates to a weighted average of single-program MVPFs—a natural benchmark for researchers relying on program-specific welfare evaluations conducted in isolation. For more details on these calculations, see Appendix B.6.

8.1 MVPF Components

Since we consider combinations of programs, our methodology is complex and incorporates roughly 125 parameters drawn from our estimates, external data, and the broader literature. We briefly describe our methods and alternative modeling assumptions before turning to our results. To aid interpretation, we group MVPF components into three categories: (1) direct WTP, (2) direct fiscal costs, and (3) indirect changes in WTP and fiscal costs.

Direct Willingness-to-Pay (WTP). CommonHelp directly reduced administrative burdens for applicants, generating time savings that we convert into WTP. First, the transition from a paper-based form to an online system reduced the application completion time by approximately 90 minutes per application (Virginia Department of Social Services 2012; Code for America 2024). Second, CommonHelp reduced the need for in-person field office visits to meet with caseworkers and submit the application, which we convert into travel time savings. Note that these WTP estimates are conservative as they exclude transportation expenses, potential wait time reductions, and the WTP for applicants who may have been rejected.

Direct Fiscal Costs. We account for changes in two types of administrative costs: (1) up-front (fixed) costs of $10.7 million associated with the CommonHelp platform's development (Pittman 2012) and (2) changes in per-recipient (variable) administrative costs via reduced agency workloads from streamlining the processing of applications. To assess the latter, we conduct a synthetic control analysis that compares Virginia's per-recipient SNAP administrative costs to those of a "synthetic" control state and assume that the effects for SNAP extend proportionally to other programs (see Appendix Figure A.30). We also estimate the increase in benefit disbursements associated with enrollment growth, relying on amounts imputed from matched public-use administrative data for Virginia.

Indirect Effects. In addition to time savings and direct costs, CommonHelp's MVPF may be shaped by behavioral responses and other indirect effects from program receipt. It is beyond the scope of this paper to estimate these in our data, so we draw on Hendren and Sprung-Keyser (2020) and other studies to assign relevant values. For fiscal externalities, we restrict to estimates based on first-order outcomes, such as labor supply, long-run

earnings, and (in Medicaid's case) healthcare cost savings.[24] Externalities are calculated separately across four age groups: young children (0–5), older children (6–17), adults (18–64), and seniors (65+). When computing "complete MVPFs" (accounting for welfare and fiscal changes across all programs), we consider three modeling assumptions regarding the degree of complementarity in fiscal externalities across programs: (1) externalities are additive, (2) externalities are amplified by 10%, and (3) externalities are amplified by 50%.[25]

Finally, we allow for the possibility that increased program receipt raises the WTP for CommonHelp. Under a neoclassical model with full information, induced recipients are marginal and thus possess a WTP of \$0. Yet, credible evidence suggests that informational frictions limit take-up. We thus consider incorporating WTP estimates from the literature scaled by different degrees of misperception of the value of program benefits: 50%, 90%, and 100% (full misperception). We also consider cases where only the WTP of children are incorporated, assuming that parents applying for programs do not internalize the value of benefits reaped by their children. For "complete MVPFs," we sum the valuations across programs; for "single-program MVPFs," we include only the respective program's WTP.

8.2 MVPF Estimates

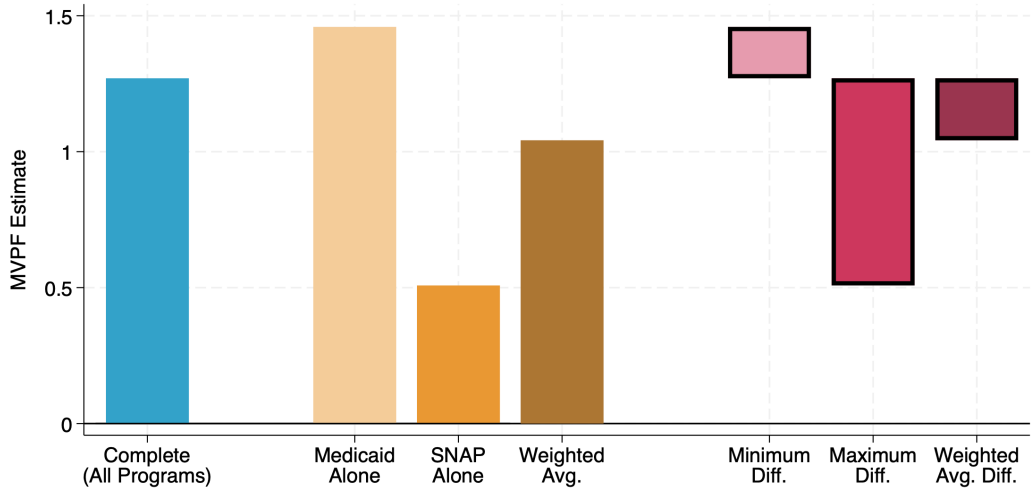
Figure 9, Panel (a) presents an illustrative MVPF calculation under the assumptions that fiscal externalities are additive across programs and that 90% of new recipients misperceive the value of program benefits (based on Finkelstein and Notowidigdo 2019). In this case, the complete MVPF (blue bar) is 1.27, suggesting that technological platforms like CommonHelp may yield positive social returns per dollar of fiscal cost. However, this MVPF differs considerably from estimates that focus on a single program in isolation. The MVPF for Medicaid alone (1.46) would exceed the complete MVPF, indicating that Medicaid's "value-added" to the complete WTP is higher than that of SNAP and TANF. In contrast, the MVPF for SNAP alone is 0.51, suggesting that evaluations of CommonHelp solely through SNAP might lead one to conclude that it was not a fruitful public investment. Even the weighted average of the single-program MVPFs—the closest approximation to a multiple-program MVPF without linked data—yields an MVPF of 1.04, still below the complete MVPF. This example highlights how incorporating joint participation in multiple programs can lead to qualita-

24. Ultimately, fiscal externality and WTP estimates depend on the range of outcomes that can be observed in data. A growing literature has expanded the set of important outcomes considered for specific programs. For example, Bailey et al. (2024) estimate the effects of early-life SNAP receipt on later-life mortality. Unfortunately, comparable estimates are not available for other programs like TANF. To ensure consistency and comparability across programs, we restrict attention to the most salient outcomes studied in the literature.

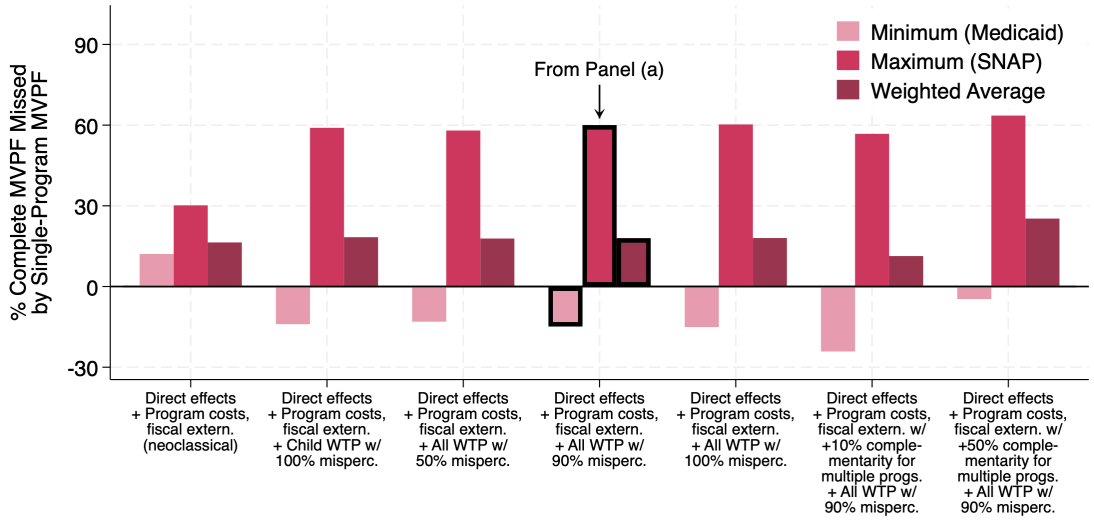
25. The latter two scenarios might arise from sharper behavioral responses to benefit cliffs (Moffitt 1979) or complementarities in human capital investments across program benefits (Rossin-Slater and Wüst 2020).

Figure 9: The MVPF of CommonHelp: Complete vs. Single Programs

(a) Example of MVPF Calculations



(b) Differences between Single-Program and Complete MVPFs, by Modeling Assumption



Notes: Panel (a) compares estimates of the complete MVPF and single-program MVPFs, where the complete MVPF assumes fiscal externalities of programs are additive and that new program recipients misperceive the value of program benefits by 90%. The horizontal solid line takes the average of the MVPF estimates considering Medicaid alone and SNAP alone, and the right-most bar takes differences between the complete MVPF and each of the single-program MVPFs. Panel (b) reports these differences between the complete and single-program MVPF estimates under alternative assumptions; see main text for details.

tively different welfare conclusions depending on the modeling assumptions. We exclude TANF from the main discussion due to its small size (see Appendix Figure A.31).

Appendix Figure A.32 breaks down the sources of differences between the complete and single-program MVPFs. For simplicity, we again focus on the set of assumptions from our il-

lustrative example (Figure 9, Panel (a)). The Medicaid-only MVPF overstates the complete MVPF by 15% because its 17% understatement of the complete WTP (the MVPF's numerator) falls short of its 27% understatement of the complete fiscal costs (the denominator). On the one hand, the smaller proportional decrease in WTP stems from the high social value of Medicaid (Finkelstein et al. 2019), which factors into WTP calculations when there are misperceptions of program benefits. On the other hand, the larger proportional decrease in the denominator stems from Medicaid's fiscal externalities often serving to *reduce* net costs, due to large cost savings for children and limited evidence of labor supply reductions among adults. Conversely, the SNAP-only MVPF understates the complete MVPF by 60%, with its WTP nearly 77% lower and fiscal costs only 42% lower than in the complete case. This is a result of SNAP recipients tending to value benefits less than one-for-one (Hendren and Sprung-Keyser 2020), while SNAP's fiscal externalities tend to *raise* net costs from adults via reduced earnings (Hoynes and Schanzenbach 2012).

Panel (b) of Figure 9 extends the overall comparisons of MVPFs across a range of modeling assumptions. Specifically, it shows minimum and maximum differences between the complete and single-program MVPFs, expressed as a percent of the complete MVPF (i.e., the share of the total value missed if evaluating a single program in isolation). Across all modeling assumptions, the single-program MVPFs substantially differ from the complete MVPF. Under the neoclassical model (the left-most bars), the single-program MVPFs are consistently *smaller* than the complete MVPF. This reveals a key feature of the complete MVPF: it need not be a convex combination of the individual program-level MVPFs if certain components like the WTP for time savings and fixed administrative costs do not scale up with the number of programs analyzed. In this case, the gains from higher WTP from incorporating time savings of all recipients are sufficiently larger than their additional variable administrative costs, leading the complete MVPF to exceed even the Medicaid-only MVPF.

Introducing misperceptions of the value of program benefits (the remaining bars) amplifies the range of differences between single-program and complete MVPF estimates. The understatement of the SNAP-only MVPF grows to around 60%, driven in large part by the omission of the substantial WTP for Medicaid under informational frictions. Conversely, the Medicaid-only MVPF *overestimates* the complete MVPF by 4–25%. Across all specifications, the weighted average of the single-program MVPFs underestimates the complete MVPF by 11–25%. These results show how considering multiple programs that comprise the safety net can substantively change our welfare evaluation of policy reforms like CommonHelp.

9 Conclusion

The coexistence of multiple programs is a fundamental feature of the modern U.S. social safety net. This paper examines the implications of adopting a multiple-program framework on take-up, targeting, and welfare analysis. In doing so, our study builds upon an active literature that has largely focused on studying individual safety net programs in isolation.

Using linked administrative data from Virginia, we show that receipt of multiple programs is popular and often occurs jointly, particularly among households with children. Yet, persistent incomplete take-up remains, especially in areas more distant from field offices. Motivated by these facts, we develop an empirical framework to identify the effects of the CommonHelp platform that streamlined applications across programs. CommonHelp significantly increased multiple program receipt, exceeding the growth in caseloads of any program. These patterns were driven by transitions from single to multiple programs, reflecting the reform's impact on the intensity of safety net attachment. These intensive-margin responders were disproportionately more disadvantaged, suggesting that those only partially attached to the safety net may be among the most vulnerable. These results not only imply that many individuals place positive value on additional programs, but also that the average *level* of that value may be higher among those already connected to the safety net. Finally, accounting for multiple program participation can meaningfully affect welfare analyses: failing to do so could understate the MVPF by up to 64% or overstate it by up to 25%.

These findings carry important policy implications. First, they reveal that employing a multiple-program framework can fundamentally change how we understand the nature of safety net attachment and targeting efficiency. Since the most disadvantaged individuals may already be partially connected to the safety net, policies that expand program access for this group could improve targeting. This insight would be missed by traditional targeting approaches that model safety net attachment as a single margin of selection. Second, our results suggest that integrated online platforms like CommonHelp can successfully reduce both distance-related and distance-invariant barriers for accessing programs. Third, our welfare analysis demonstrates that single-program evaluations can dramatically misstate the social value of policies when the multifaceted nature of the modern-day safety net is ignored.

While this paper focuses on households with children, future work could extend the framework to examine other groups eligible for multiple programs. These include the elderly, disabled, and (recently) childless adults, who may respond differentially to policies like CommonHelp. In addition, our MVPF calculations rely on outside estimates and assumptions on complementarities across programs. In future work, we hope to examine the downstream outcomes of multiple programs to provide ingredients for more direct cost-benefit analyses.

References

Abadie, Alberto. 2003. "Semiparametric Instrumental Variable Estimation of Treatment Response Models." *Journal of Econometrics* 113 (2): 231–263.

Aizer, Anna. 2007. "Public Health Insurance, Program Take-Up, and Child Health." *Review of Economics and Statistics* 89 (3): 400–415.

Aizer, Anna, Hilary Hoynes, and Adriana Lleras-Muney. 2022. "Children and the US Social Safety Net: Balancing Disincentives for Adults and Benefits for Children." *Journal of Economic Perspectives* 36 (2): 149–174.

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

Arbogast, Iris, Anna Chorniy, and Janet Currie. 2024. "Administrative Burdens and Child Medicaid and CHIP Enrollments." *American Journal of Health Economics* 10 (2): 237–271.

Baicker, Katherine, Amy Finkelstein, Jae Song, and Sarah Taubman. 2014. "The Impact of Medicaid on Labor Market Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment." *American Economic Review* 104 (5): 322–328.

Baicker, Katherine, and Theodore Svoronos. 2019. *Testing the Validity of the Single Interrupted Time Series Design.* Technical report. National Bureau of Economic Research.

Bailey, Martha J., Hilary Hoynes, Maya Rossin-Slater, and Reed Walker. 2024. "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program." *Review of Economic Studies* 91 (3): 1291–1330.

Bastian, Jacob E., and Maggie R. Jones. 2021. "Do EITC Expansions Pay for Themselves? Effects on Tax Revenue and Government Transfers." *Journal of Public Economics* 196:104355.

Bertrand, Marianne, Sendhil Mullainathan, and Eldar Shafir. 2004. "A Behavioral-Economics View of Poverty." *American Economic Review* 94 (2): 419–423.

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

Bitler, Marianne P., Jason B. Cook, Chloe N. East, Sonya R. Porter, and Laura Tiehen. 2025. *The Intersection of Place and Need: How Lack of Enrollment Offices Deters Participation in the Safety Net.* Working Paper.

Blank, Rebecca M., and Patricia Ruggles. 1996. "When Do Women Use Aid to Families with Dependent Children and Food Stamps?" *Journal of Human Resources* 31 (1): 57–89.

Blundell, Richard, and Luigi Pistaferri. 2003. "Income Volatility and Household Consumption: The Impact of Food Assistance Programs." *Journal of Human resources,* 1032–1050.

Borusyak, Kirill. 2015. *Bounding the Population Shares Affected by Treatments.* Working Paper.

Burns, Marguerite, and Laura Dague. 2017. "The Effect of Expanding Medicaid Eligibility on Supplemental Security Income Program Participation." *Journal of Public Economics* 149:20–34.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H.C. Sant'Anna. 2024. *Difference-in-Differences with a Continuous Treatment.* Working Paper. National Bureau of Economic Research.

Castell, Laura, Marc Gurgand, Clement Imbert, and Todor Tochev. 2025. "Shifting Administrative Burden to the State: The Case of Medicaid Take-Up." *American Economic Journal: Economic Policy* 17 (4): 1–29.

Celhay, Pablo A., Bruce D. Meyer, and Nikolas Mittag. 2022. *Stigma in Welfare Programs.* Working Paper.

Centers for Medicare & Medicaid Services. 2013. *Connecting Kids to Coverage Campaign Updates: November 8, 2013.* Accessed: 2025-08-07. https://www.insurekidsnow.gov/newsletter/2013/11/08.

———. 2024. *National Health Expenditure Data: Fact Sheet.* Accessed: 2025-03-17. https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/nhe-fact-sheet.

Chan, Marc K. 2013. "A Dynamic Model of Welfare Reform." *Econometrica* 81 (3): 941–1001.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with Zeros? Some Problems and Solutions." *The Quarterly Journal of Economics* 139 (2): 891–936.

Christensen, Julian, Lene Aarøe, Martin Baekgaard, Pamela Herd, and Donald P. Moynihan. 2020. "Human Capital and Administrative Burden: The Role of Cognitive Resources in Citizen-State Interactions." *Public Administration Review* 80 (1): 127–136.

Code for America. 2024. *Benefits Enrollment Field Guide.* Accessed: December 2, 2024. https://codeforamerica.org/explore/benefits-enrollment-field-guide/?att=online&j=VA&app=102&dev=d#explore.

Cook, Jason B., and Chloe N. East. 2025. *The Effect of Means-Tested Transfers on Work: Evidence from Quasi-Randomly Assigned SNAP Caseworkers.* Technical report. National Bureau of Economic Research.

Currie, Janet. 2006a. "Public Policy and the Income Distribution." Chap. The Take-Up of Social Benefits. Russell Sage Foundation.

Currie, Janet. 2006b. *The Invisible Safety Net: Protecting the Nation's Poor Children and Families.* Princeton: Princeton University Press.

Currie, Janet, and Jeffrey Grogger. 2001. "Explaining Recent Declines in Food Stamp Program Participation." *Brookings-Wharton Papers on Urban Affairs,* 203–244.

Currie, Janet, and Reed Walker. 2011. "Traffic Congestion and Infant Health: Evidence from E-ZPass." *American Economic Journal: Applied Economics* 3 (1): 65–90.

Dahl, Gordon B., Katrine V. Løken, and Magne Mogstad. 2014. "Peer Effects in Program Participation." *American Economic Review* 104 (7): 2049–2074.

Daponte, Beth, Seth Sanders, and Lowell Taylor. 1999. "Why Do Low-Income Households Not Use Food Stamps? Evidence from an Experiment." *Journal of Human Resources* 34 (3): 612–628.

Deshpande, Manasi, and Yue Li. 2019. "Who Is Screened Out? Application Costs and the Targeting of Disability Programs." *American Economic Journal: Economic Policy* 11 (4): 213–248.

Deshpande, Manasi, and Lee M. Lockwood. 2022. "Beyond Health: Nonhealth Risk and the Value of Disability Insurance." *Econometrica* 90 (4): 1781–1810.

East, Chloe N., and David Simon. 2024. "The Safety Net and Job Loss: How Much Insurance Do Public Programs Provide?" *Journal of Public Economics* 238 (105171).

Ebenstein, Avraham, and Kevin Stange. 2010. "Does Inconvenience Explain Low Take-Up? Evidence from Unemployment Insurance." *Journal of Policy Analysis and Management* 29 (1): 111–136.

Edelstein, Sara, Michael R. Pergamit, and Caroline Ratcliffe. 2014. *Characteristics of Families Receiving Multiple Public Benefits.* Technical report. Washington, DC: Urban Institute.

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

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127 (6): 2836–2874.

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

Flood, Lennart, Jörgen Hansen, and Roger Wahlberg. 2004. "Household Labor Supply and Welfare Participation in Sweden." *Journal of Human Resources* 39 (4): 1008–1032.

Giannella, Eric, Tatiana Homonoff, Gwen Rino, and Jason Somerville. 2024. "Administrative Burden and Procedural Denials: Experimental Evidence from SNAP." *American Economic Journal: Economic Policy* 16 (4): 316–340.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–277.

Gray, Colin. 2019. "Leaving Benefits on the Table: Evidence from SNAP." *Journal of Public Economics* 179:1–15.

Grogger, Jeffrey. 2003. "The Effects of Time Limits, the EITC, and Other Policy Changes on Welfare Use, Work, and Income among Female-Headed Families." *Review of Economics and Statistics* 85 (2): 394–408.

Han, Jeehoon. 2020. "SNAP Expansions and Participation in Government Safety Net Programs." *Economic Inquiry* 58 (4): 1929–1948.

Heckman, James J., and Rodrigo Pinto. 2018. "Unordered Monotonicity." *Econometrica* 86 (1): 1–35.

Heflin, Colleen, Jun Li, and Dongmei Zuo. 2023. "Changing Patterns of SNAP Take-Up and Participation and the Role of Out-of-Pocket Medical Expenses among Older Adults." *Applied Economic Perspectives and Policy* 45 (1): 336–349.

Heflin, Colleen, Michah W. Rothbart, and Mattie Mackenzie-Liu. 2022. "Below the Tip of the Iceberg: Examining Early Childhood Participation in SNAP and TANF from Birth to Age Six." *Population Research and Policy Review* 41 (2): 729–755.

Heinrich, Carolyn J., Sayil Camacho, Sarah Clark Henderson, Monica Hernandez, and Ela Joshi. 2021. "Consequences of Administrative Burden for Social Safety Nets That Support the Healthy Development of Children." *Journal of Policy Analysis and Management* 41 (1): 11–44.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics* 135 (3): 1209–1318.

Herd, Pamela, Thomas DeLeire, Hope Harvey, and Donald P. Moynihan. 2013. "Shifting Administrative Burden to the State: The Case of Medicaid Take-Up." *Public Administration Review* 73 (S1): S69–S81.

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

Herd, Pamela, and Donald Moynihan. 2018. *Administrative Burden: Policymaking by Other Means.* New York: Russell Sage Foundation.

Hicks, Jeffrey. 2024. *The Effects of Field Office Closures on Social Assistance Take-Up and Targeting.* Working Paper.

Homonoff, Tatiana, and Jason Somerville. 2021. "Program Recertification Costs: Evidence from SNAP." *American Economic Journal: Economic Policy* 13 (4): 271–298.

Hoynes, Hilary W., and Ankur J. Patel. 2018. "Effective Policy for Reducing Poverty and Inequality?: The Earned Income Tax Credit and the Distribution of Income." *Journal of Human Resources* 53 (4): 859–890.

Hoynes, Hilary Williamson, and Diane Whitmore Schanzenbach. 2012. "Work Incentives and the Food Stamp Program." *Journal of Public Economics* 96 (1-2): 151–162.

Jackson, Margot I., and Ester Fanelli. 2023. "Who Uses the Social Safety Net? Trends in Public Benefit Use Among American Households with Children, 1980–2020." *ANNALS of the American Academy of Political and Social Science* 706 (1): 16–36.

Keane, Michael, and Robert Moffitt. 1998. "A Structural Model of Multiple Welfare Program Participation and Labor Supply." *International Economic Review* 39 (3): 553–589.

Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad. 2016. "Field of Study, Earnings, and Self-Selection." *The Quarterly Journal of Economics* 131 (3): 1057–1111.

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

Kline, Patrick, and Melissa Tartari. 2016. "Bounding the Labor Supply Responses to a Randomized Welfare Experiment: A Revealed Preference Approach." *American Economic Review* 106 (4): 972–1014.

Ko, Wonsik, and Robert Moffitt. 2024. "Handbook of Labor, Human Resources and Population Economics." Chap. Take-Up of Social Benefits, 1–42. Springer International Publishing.

Leung, Pauline, and Christopher O'Leary. 2020. "Unemployment Insurance and Means-Tested Program Interactions: Evidence from Administrative Data." *American Economic Journal: Economic Policy* 12 (2): 159–192.

Linden, Leigh, and Jonah E. Rockoff. 2008. "Estimates of the Impact of Crime Risk on Property Values from Megan's Laws." *American Economic Review* 98 (3): 1103–1127.

Lindner, Stephan. 2016. "How Do Unemployment Insurance Benefits Affect the Decision to Apply for Social Security Disability Insurance?" *Journal of Human Resources* 51 (1): 62–94.

Linos, Elizabeth, Allen Prohofsky, Aparna Ramesh, Jesse Rothstein, and Matthew Unrath. 2022. "Can Nudges Increase Take-Up of the EITC? Evidence from Multiple Field Experiments." *American Economic Journal: Economic Policy* 14 (4): 432–452.

Macartney, Suzanne, and Robin Ghertner. 2023. *Participation in the U.S. Social Safety Net: Multiple Programs, 2019.* Technical report. April: Department of Health and Human Services, ASPE Office of Human Services Policy.

Marcus, Michelle. 2021. "Going Beneath the Surface: Petroleum Pollution, Regulation, and Health." *American Economic Journal: Applied Economics* 13 (1): 72–104.

Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan. 2015. "Household Surveys in Crisis." *Journal of Economic Perspectives* 29 (4): 199–226.

Meyer, Bruce D., and Derek Wu. 2018. "The Poverty Reduction of Social Security and Means-Tested Transfers." *Industrial and Labor Relations Review* 71 (5): 1106–1153.

Moffitt, Robert. 1979. "Cumulative Effective Tax Rates and Guarantees in Low-Income Transfer Programs." *The Journal of Human Resources* 14 (1): 122–129.

———. 1983. "An Economic Model of Welfare Stigma." *The American Economic Review* 73 (5): 1023–1035.

———. 2016. "Multiple Program Participation and the SNAP Program." In *SNAP Matters: How Food Stamps Affect Health and Well-Being,* edited by Craig Gundersen, James P. Ziliak, Judith Bartfeld, and Timothy Smeeding. Stanford University Press.

Mountjoy, Jack. 2022. "Community Colleges and Upward Mobility." *American Economic Review* 112 (8): 2580–2630.

———. 2024. *Marginal Returns to Public Universities.* Technical report. National Bureau of Economic Research.

Mueller, Andreas I., Jesse Rothstein, and Till M. von Wachter. 2016. "Unemployment Insurance and Disability Insurance in the Great Recession." *Journal of Labor Economics* 34 (S1): S445–S475.

National Governors Association. 2022. *Advancing Economic Mobility for Low-Income Families: Policy Options for Governors.* Accessed: 2024-10-22. https://www.nga.org/wp-content/uploads/2022/07/NGA-Economic-Mobility-Report-Final.pdf.

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

Pinto, Rodrigo. 2022. *Beyond Intention-to-Treat: Using the Incentives of Moving to Opportunity to Identify Neighborhood Effects.* Working Paper.

Pittman, Elaine. 2012. "Virginia Launches Online Public Assistance System." Published by Government Technology magazine. Accessed: 2025-05-07. https://www.govtech.com/health/Virginia-Launches-Online-Public-Assistance-System.html.

Plotnick, Robert. 1983. "Turnover in the AFDC Population: An Event History Analysis." *Journal of Human Resources,* 65–81.

Rafkin, Charlie, Adam Solomon, and Evan J. Soltas. 2025. *Self-Targeting in U.S. Transfer Programs.* Working Paper.

Rossin-Slater, Maya. 2013. "WIC in Your Neighborhood: New Evidence on the Impacts of Geographic Access to Clinics." *Journal of Public Economics* 102 (C): 51–69.

Rossin-Slater, Maya, and Miriam Wüst. 2020. "What Is the Added Value of Preschool for Poor Children? Long-Term and Intergenerational Impacts and Interactions with an Infant Health Intervention." *American Economic Journal: Applied Economics* 12 (3): 255–286.

Schmidt, Lucie, Lara Shore-Sheppard, and Tara Watson. 2019. *The Impact of Expanding Public Health Insurance on Safety Net Program Participation: Evidence from the ACA Medicaid Expansion.* Working Paper. National Bureau of Economic Research.

Schmidt, Lucie, Lara Shore-Sheppard, and Tara Watson. 2024. "Safety-Net Program Spillovers and Administrative Burden." *National Tax Journal* 77 (1): 199–221.

Shepard, Mark, and Myles Wagner. 2025. "Do Ordeals Work for Selection Markets? Evidence from Health Insurance Auto-Enrollment." *American Economic Review* 115 (3): 772–822.

Unrath, Matthew. 2024. *Targeting, Screening, and Retention: Evidence from the Supplemental Nutrition Assistance Program in California.* Working Paper.

USDA Food and Nutrition Service. 2024. *Supplemental Nutrition Assistance Program (SNAP) – Annual Summary.* Accessed: 2025-03-17. https://fns-prod.azureedge.us/sites/default/files/resource-files/snap-annualsummary-1.pdf.

Virginia Department of Social Services. 2012. *A Quick Guide to Using CommonHelp.* https://www.dss.virginia.gov/files/division/pa/commonhelp_community/guidance_procedures/qrg.pdf. Accessed: 2025-05-07.

———. 2017. *Benefits Program Brochure.* Accessed: 2025-03-17. https://www.dss.virginia.gov/files/division/bp/verification_requirements/brochures/b032-01-0002-20-eng.pdf.

———. 2024. *Supplemental Nutrition Assistance Program (SNAP) Manual – Volume V, Part II: Application Processing.* Accessed: 2025-03-17. https://www.dss.virginia.gov/files/division/bp/fs/manual/Part_ii.pdf.

Wu, Derek, and Bruce D. Meyer. 2023. *Certification and Recertification in Welfare Programs: What Happens When Automation Goes Wrong?* Working Paper.

Wu, Derek, and Jonathan Zhang. 2025. "Sliding into Safety Net Participation: A Unified Analysis Across Multiple Programs." *National Tax Journal* 78 (1): 45–86.

Online Appendix for "Multiple Program Participation in the Safety Net: Incidence, Impediments, and Implications"

Neil A. Cholli and Derek Wu

March 26, 2026

Table of Contents

A Supplemental Figures and Tables **OA-1**

B Technical Appendix **OA-39**

B.1 Creating Program Participation Data at Zip-County Level OA-39

B.2 Additional Details on Identification Strategies OA-39

B.3 Converting Regression Parameters to Percentage Point Effects OA-41

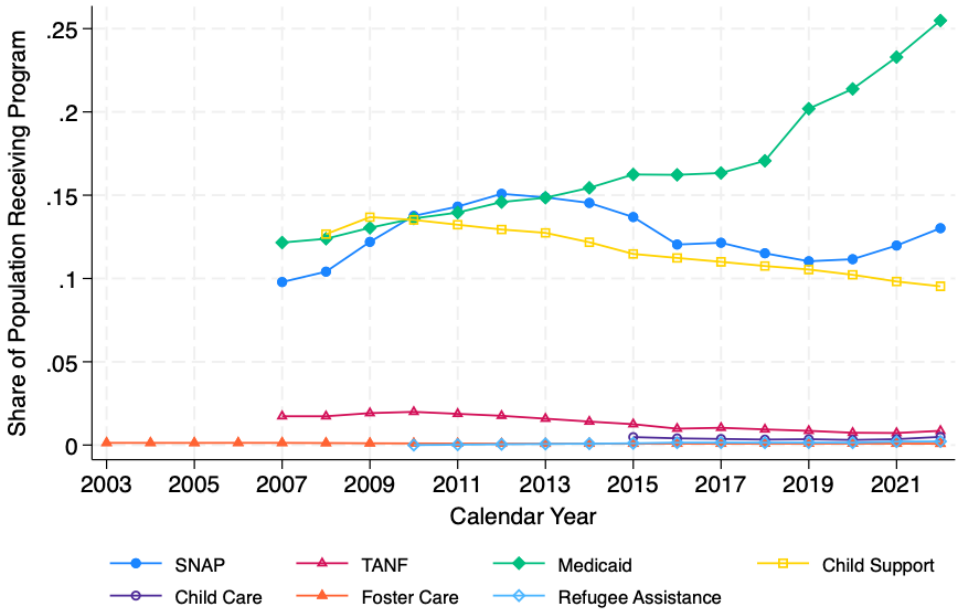
B.4 Partial Identification Analysis OA-41

B.5 Conventional Targeting Analysis OA-45

B.6 Details on Marginal Value of Public Funds Analysis OA-46

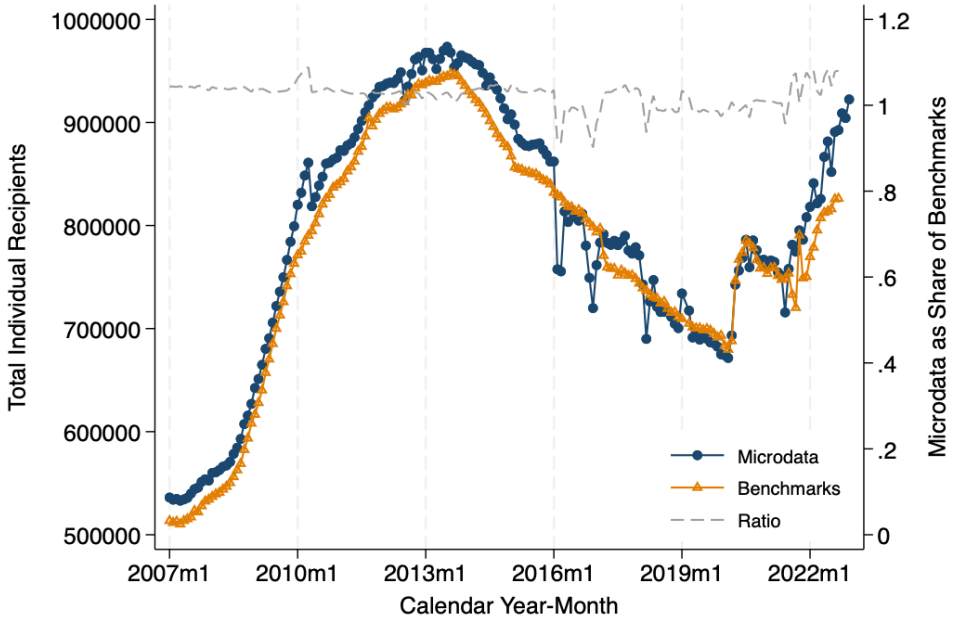
A Supplemental Figures and Tables

Figure A.1: Trends in Receipt of VDSS Programs



Notes: This figure illustrates the share of the population in Virginia receiving each of the seven programs administered by VDSS in a given year. Population counts (for the denominator) are taken from the Census Bureau. Given variation in data availability, we observe: foster care receipt from 2003-2022; SNAP, TANF, and Medicaid receipt from 2007-2022; child support receipt from 2008-2022; refugee assistance receipt from 2010-2022; and child care subsidy receipt from 2015-2022.

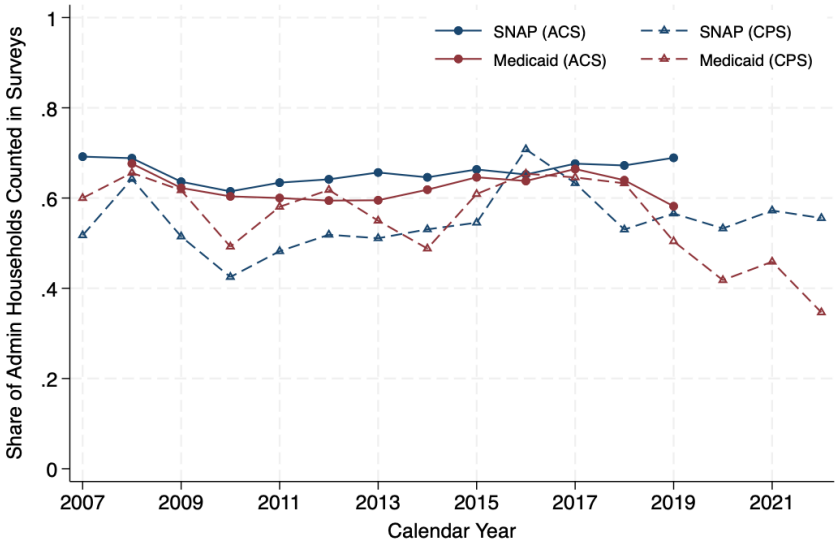
Figure A.2: Comparison of Microdata Aggregates to USDA Totals for SNAP



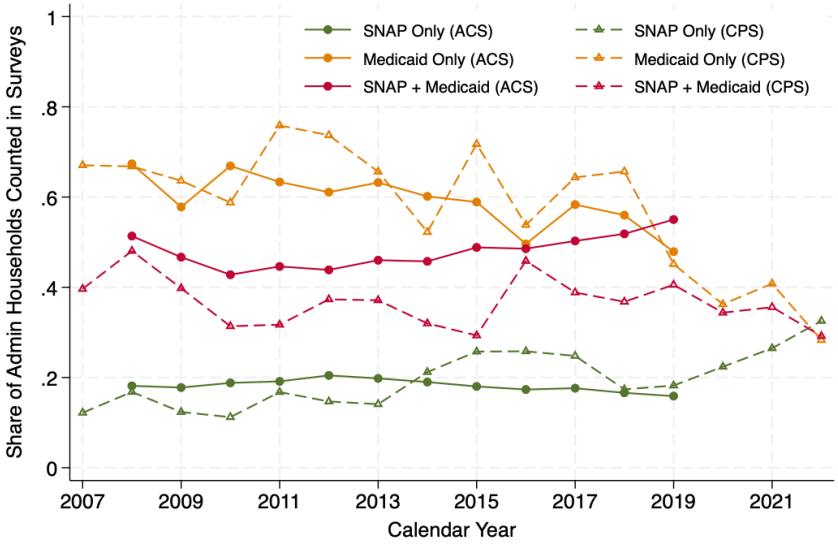
Notes: This figure shows the monthly counts of individuals receiving SNAP between January 2007 and December 2022 based on two administrative sources: aggregates calculated from the VDSS microdata (blue circles) and publicly available benchmarks from USDA FNS (orange triangles). The ratio of the VDSS microdata aggregates to publicly available benchmarks is shown in the dotted gray line.

Figure A.3: Comparison of ACS and CPS Totals to Microdata Aggregates

(a) Individual Programs

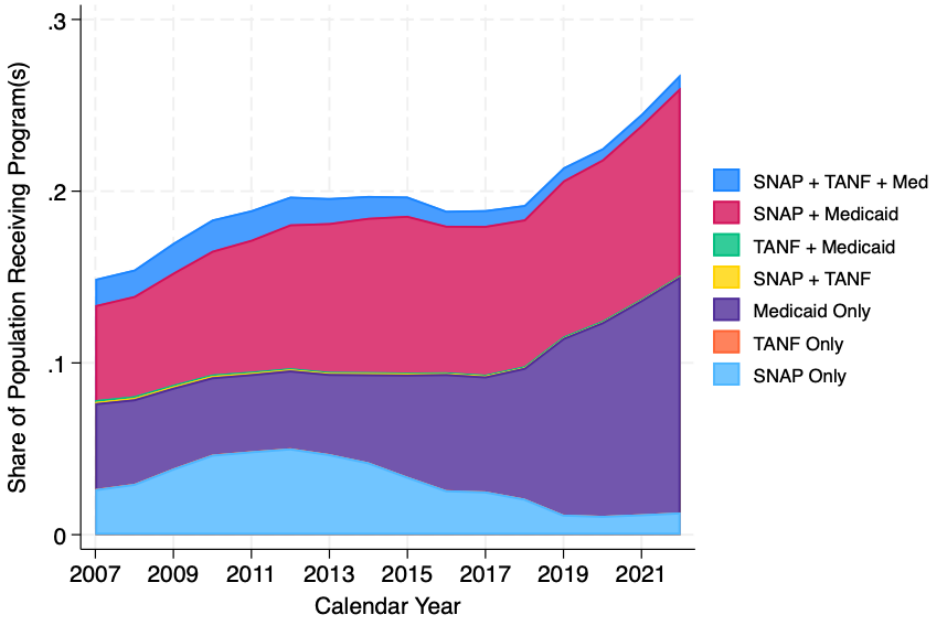


(b) Program Bundles



Notes: These figures show the ratios of survey-reported recipient households from the ACS and CPS ASEC to administrative cases in the VDSS microdata. Panel (a) shows comparisons for SNAP and Medicaid separately, while Panel (b) shows comparisons of program bundles between SNAP and Medicaid (we omit TANF because the ACS does not ask about it separately). In both surveys, SNAP receipt is recorded at the household level while Medicaid receipt is recorded at the individual level. The reference period is the calendar year in the CPS and in the administrative records, while it is the last 12 months in the ACS. Note that administrative cases can span multiple households.

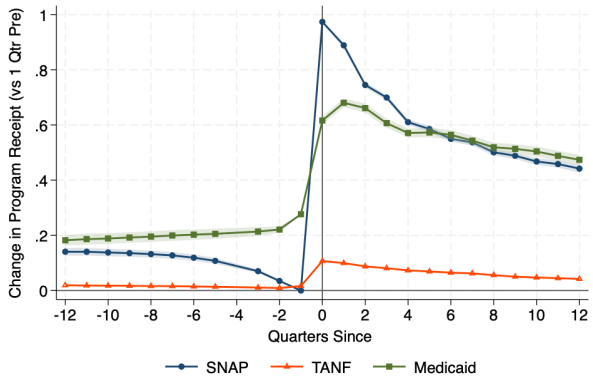
Figure A.4: Trends in Participation in Bundles of Programs



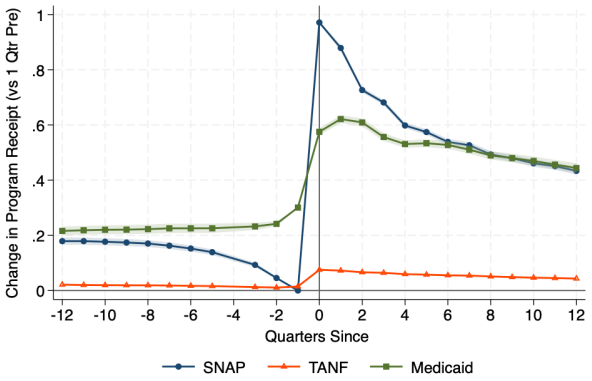
Notes: This figure illustrates shares of the population in Virginia receiving one of seven possible program bundles (between SNAP, TANF, and Medicaid) over time. Annual recipient counts are calculated from the VDSS microdata, while population estimates are drawn from the American Community Survey.

Figure A.5: Interactions Around Program Entry Conditional on Earnings Decline

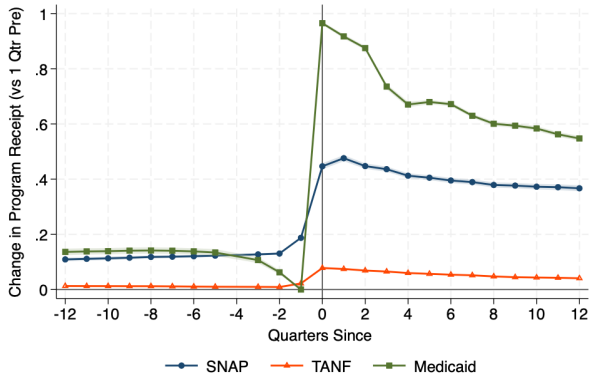
(a) SNAP Entry: No Earnings Loss $\geq$ 25%



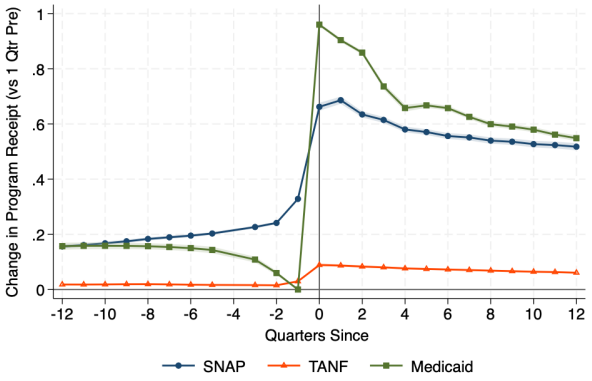
(b) SNAP Entry: Earnings Loss $\geq$ 25%



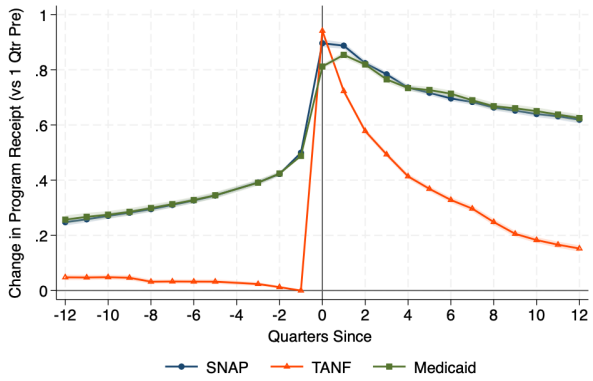
(c) Medicaid Entry: No Earnings Loss $\geq$ 25%



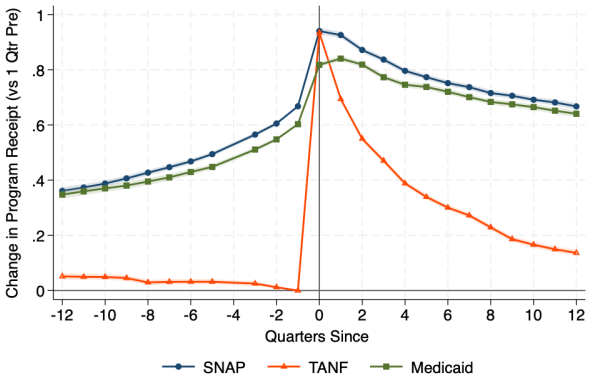
(d) Medicaid Entry: Earnings Loss $\geq$ 25%



(e) TANF Entry: No Earnings Loss $\geq$ 25%



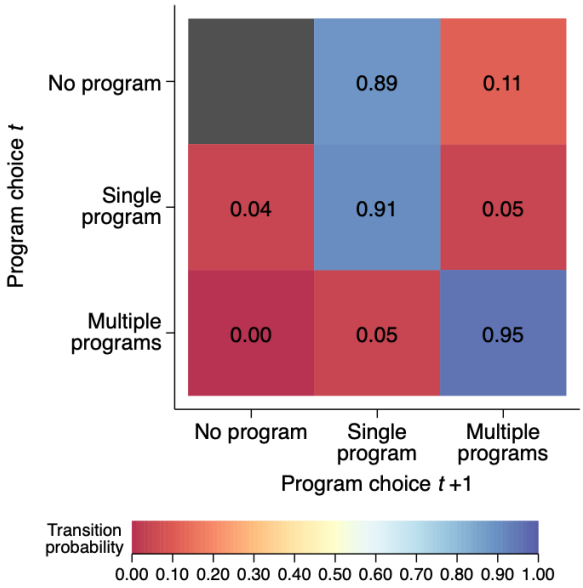
(f) TANF Entry: Earnings Loss $\geq$ 25%



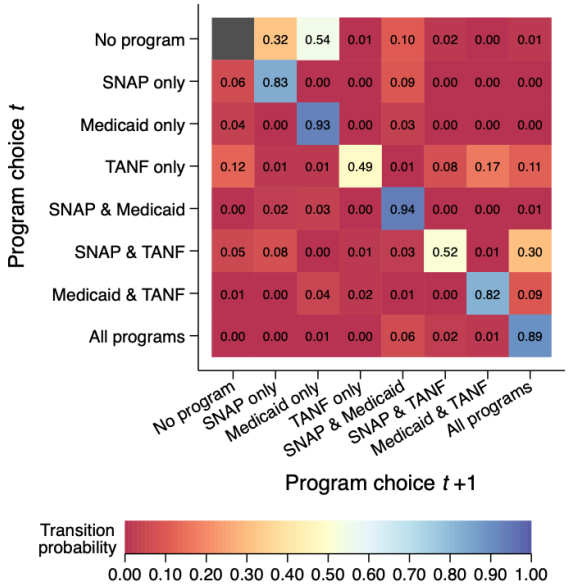
Notes: These figures plot event-study estimates of each program receipt on program entrants, focusing on individuals in households with children who enter a given program and did not receive the program in the immediate prior quarter. We use a 10% random sample of entrants. Left (right) panels condition on individuals whose households did not (did) suffer an earnings loss of 25%+ between quarters in the 12 months leading up to initial receipt. Our sample consists of those entering between January 2010-December 2019. Shaded regions are 95% confidence intervals.

Figure A.6: Month-to-Month Transition Matrices between Program Choices

(a) Single and Multiple Programs



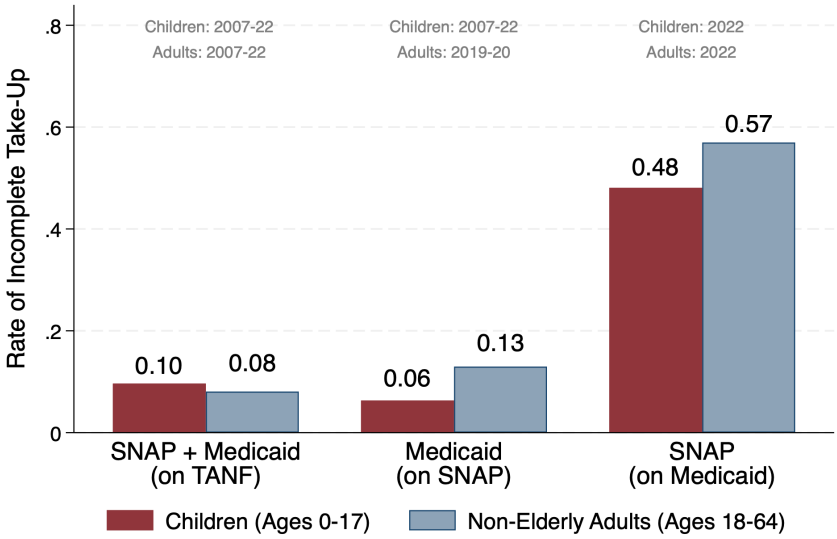
(b) Bundles of Programs



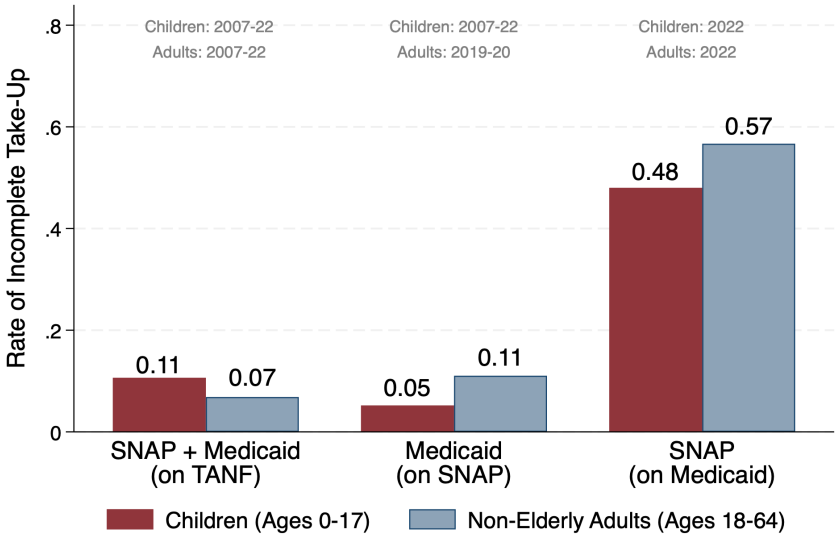
Notes: This figure illustrates transition matrices using monthly data between January 2007 and December 2022 among individuals living in households with children. Matrices exclude no program-to-no program transitions since they are based on recipient data.

Figure A.7: Evidence of Incomplete Take-Up of Multiple Programs (Varying Minimum Receipt Length of Conditioning Program)

(a) Minimum of 3 Months

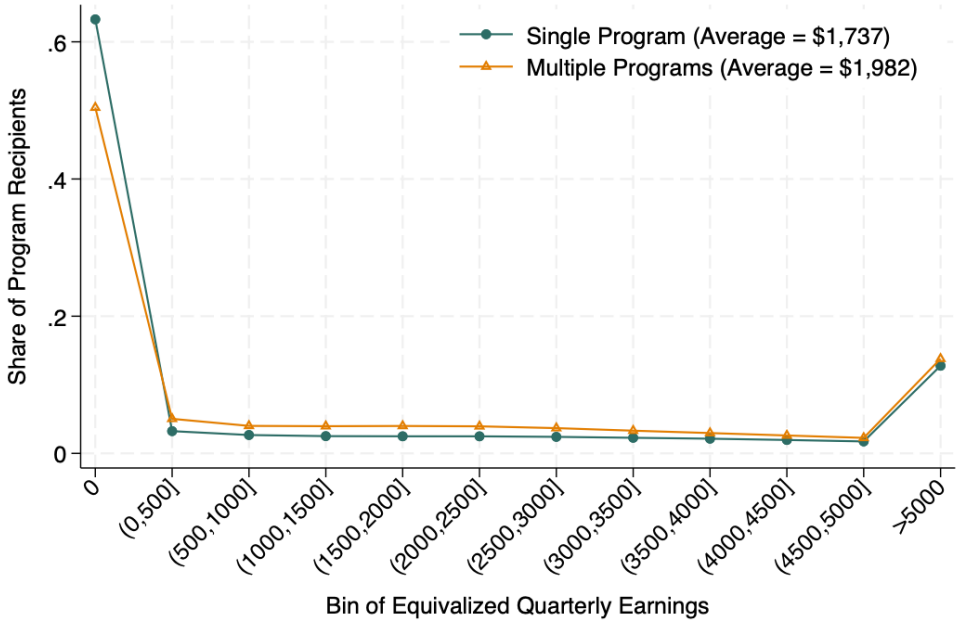


(b) Minimum of 6 Months



Notes: These figures plot conditional participation rates of different programs among children and adults who are inferred to be eligible based on their participation in other programs. Panel (a) calculates rates for those who are enrolled in the conditioning program for at least 3 months in a calendar year. Panel (b) calculates rates for those who are enrolled in the conditioning program for at least 6 months in a calendar year. See notes to Figure 3 for more details.

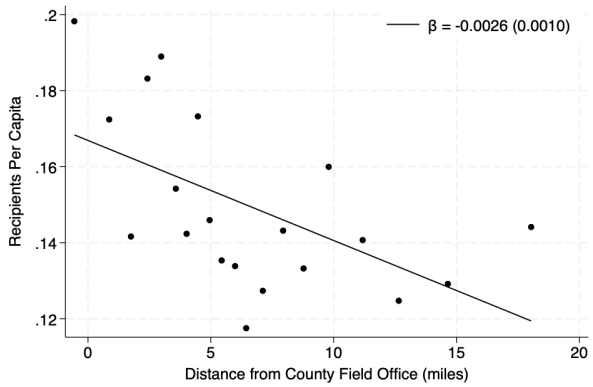
Figure A.8: Wages Among Single vs. Multiple Program Recipients



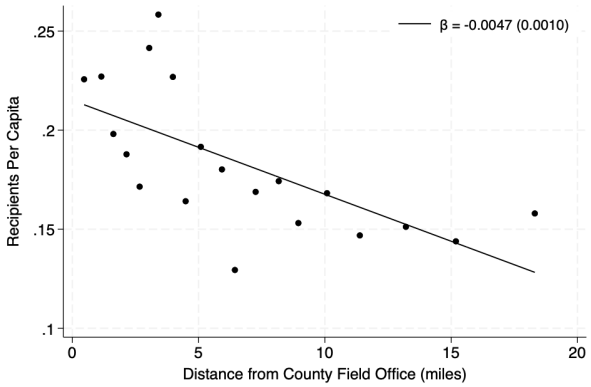
Notes: This figure compares the distribution of quarterly household wages for single and multiple program recipients among households with children between 2007-22. Household earnings are equivalized using the National Academy of Sciences equivalence scale.

Figure A.9: Receipt Rates Decrease in Distance to County Field Office

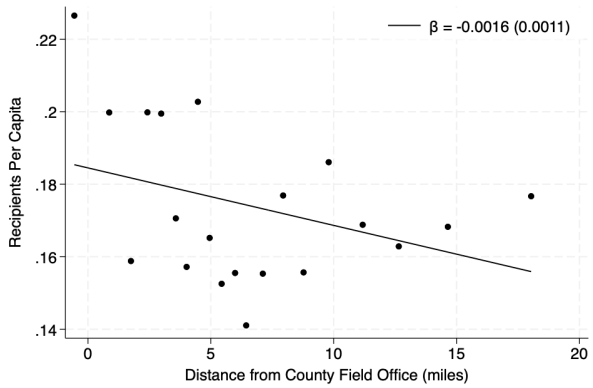
(a) Multiple Program Receipt (with controls)



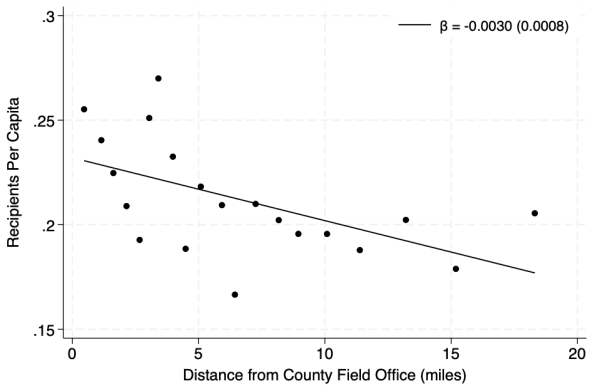
(b) Multiple Program Receipt (no controls)



(c) Single Program Receipt (with controls)



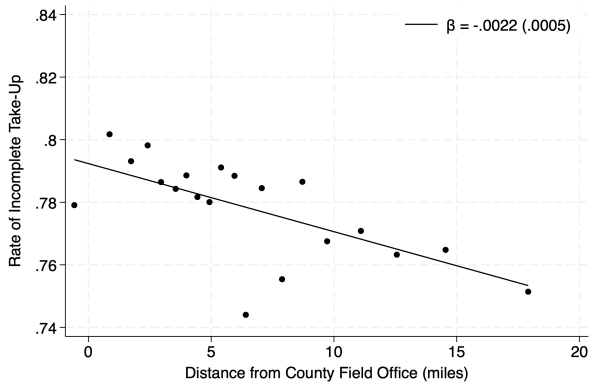
(d) Single Program Receipt (no controls)



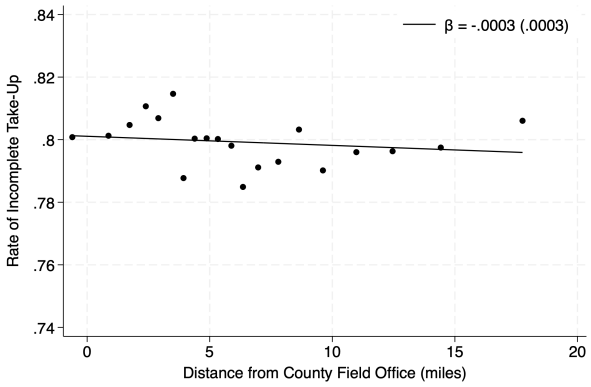
Notes: These figures illustrates the association between zip-level multiple and single program participation rates and zip code distance to their assigned county field office using monthly data between January 2007 and December 2022. Participation rates are expressed in per capita terms using population estimates of zip code tabulation areas (ZCTAs) from the 5-year ACS, scaled to the zip-county level using the HUD-USPS crosswalk. Binscatters are weighted by any program receipt. Panels (a) and (c) control for ZCTA-level poverty rate, zip-level population share of Blacks and Hispanics, and county fixed effects (drawn from 5-year ACS estimates), while Panels (b) and (d) do not include any controls.

Figure A.10: Conditional Take-Up Rates Decrease in Distance to County Field Office

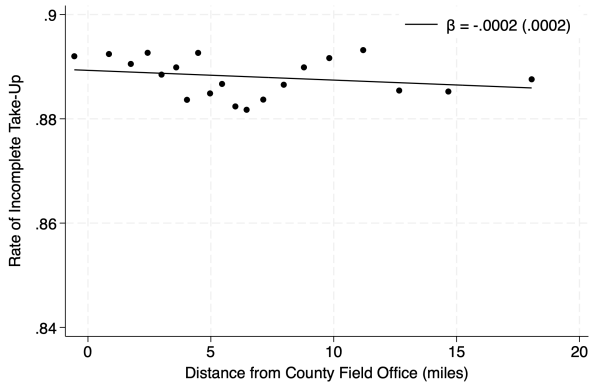
(a) Children: SNAP + Medicaid (on TANF)



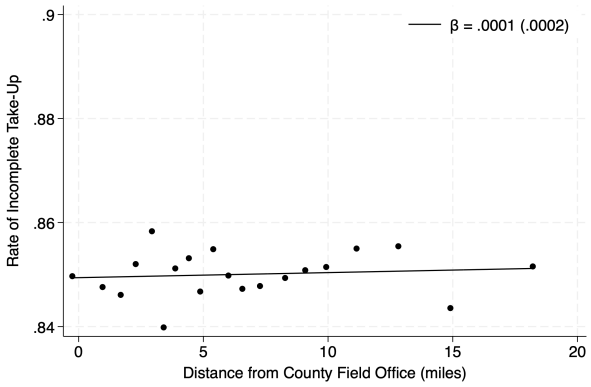
(b) Adults: SNAP + Medicaid (on TANF)



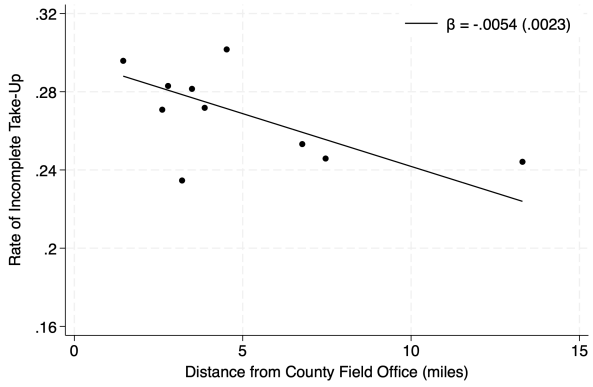
(c) Children: Medicaid (on SNAP)



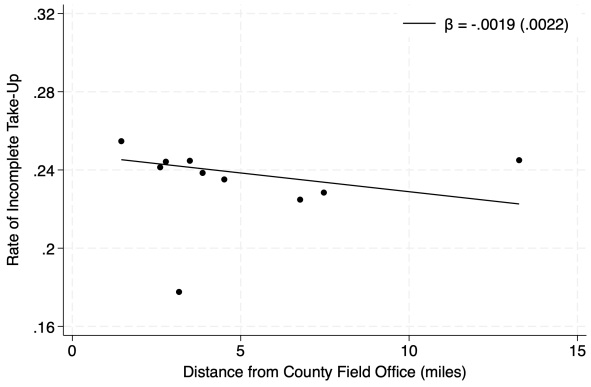
(d) Adults: Medicaid (on SNAP)



(e) Children: SNAP (on Medicaid)



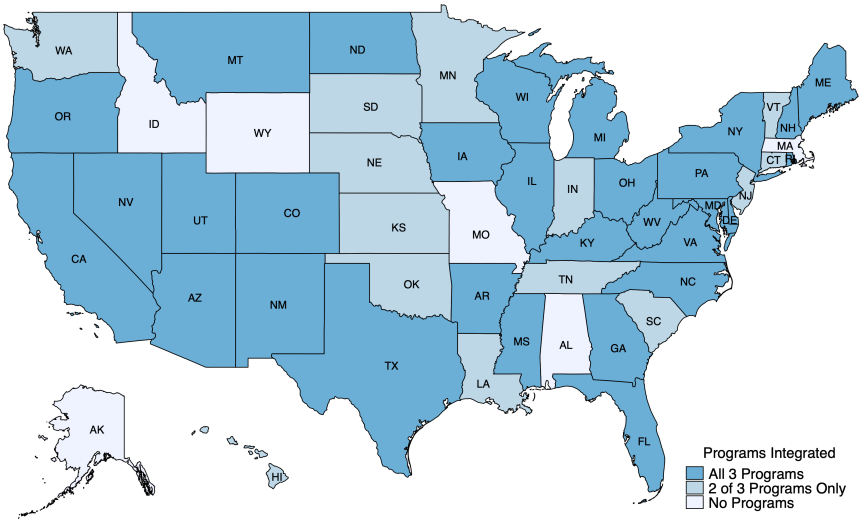
(f) Adults: SNAP (on Medicaid)



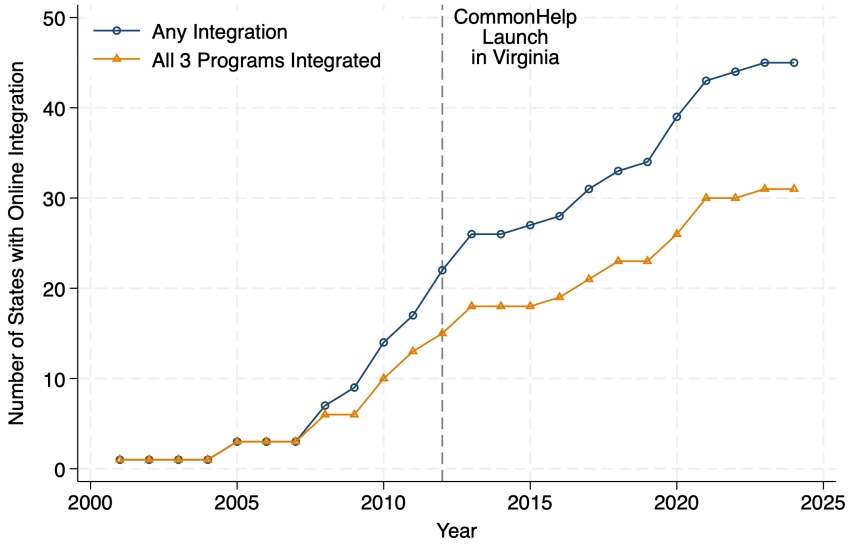
Notes: These figure illustrates the association between zip-level take-up rates of multiple programs (conditional on being inferred to be eligible based on participation in other programs) and zip code distance to their assigned county field office. Following Figure 3, Panels (a) and (b) are calculated over 2007-22, Panel (c) is calculated over 2007-22, Panel (d) is calculated over 2019-20, and Panels (e) and (f) are calculated over 2022 only. Binscatters are weighted by any program receipt and control for ZCTA-level poverty rate, zip-level population share of Blacks and Hispanics, and county fixed effects (drawn from 5-year ACS estimates).

Figure A.11: Adoption of Integrated Online Applications Across States

(a) Adopting States (as of 2024)



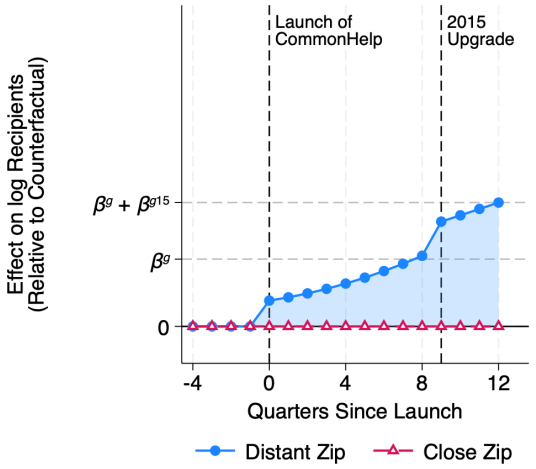
(b) Number of States Adopting Over Time



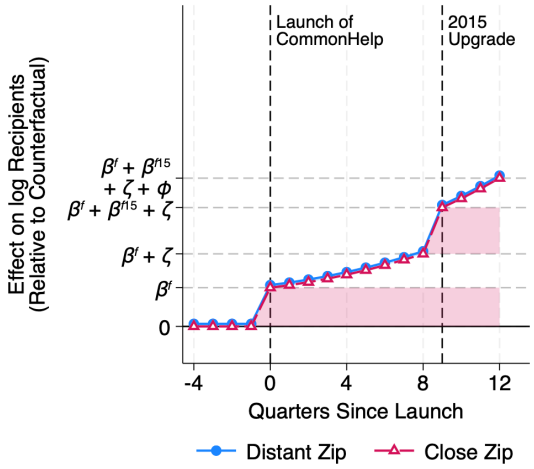
Notes: Panel (a) shows which states have adopted integrated online applications across multiple programs (SNAP, TANF, and Medicaid) as of 2024, based on Code for America's Benefits Enrollment Field Guide (https://codeforamerica.org/explore/benefits-enrollment-field-guide/). The differential shading distinguishes states with no integration across programs, states with integration over only two out of three programs (generally SNAP and TANF), and states with all three programs integrated into a single online application. Panel (b) shows the trend over time in the number of states adopting integrated online applications. The count of states in 2024 is based on Code for America data, while counts for earlier years are based on manual searches for the start date of integration. We show separate trends for the number of states with any integration across two or three programs (blue circles) and with integration across all three programs (orange triangles).

Figure A.12: Expected Impacts on Program Receipt under Different Assumptions, with Dynamic Impacts

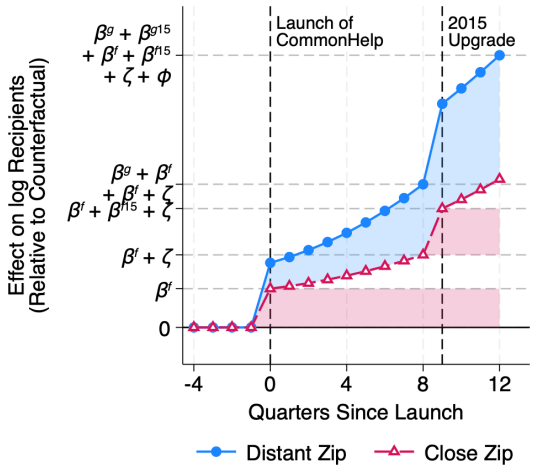
(a) Reduced Geographic Burdens Only



(b) Reduced Fixed Burdens Only



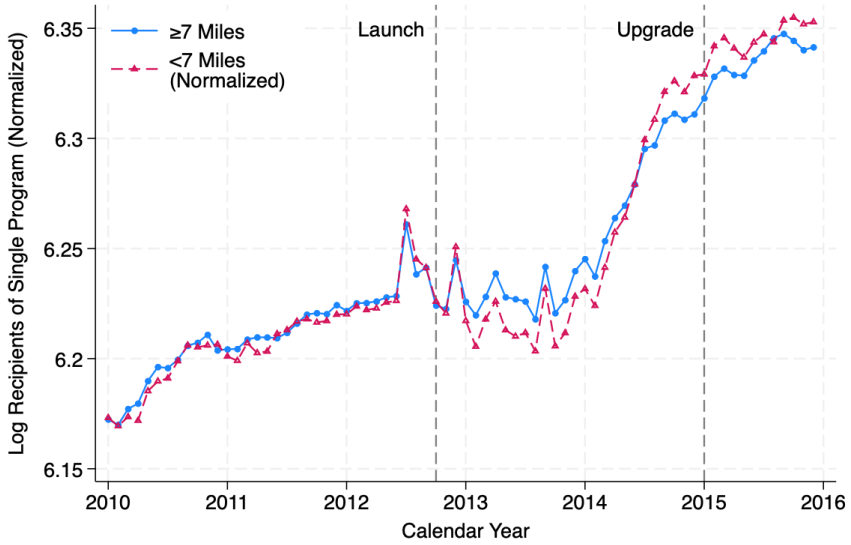
(c) Reduced Geographic + Fixed Burdens



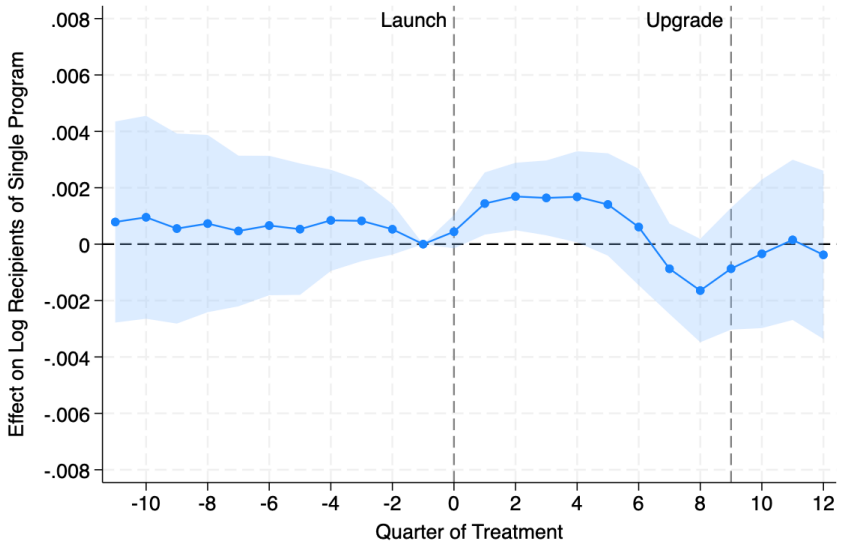
Notes: See notes to Figure 5.

Figure A.13: Raw Trends and Event-Study Estimates for Single Program

(a) Single Program: Raw Trends



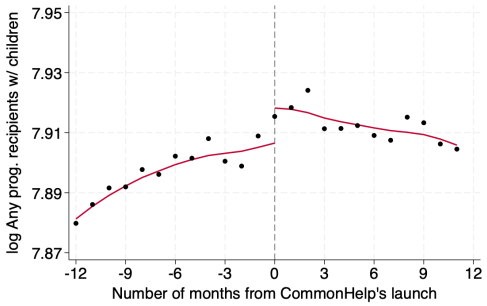
(b) Single Program: Event Study



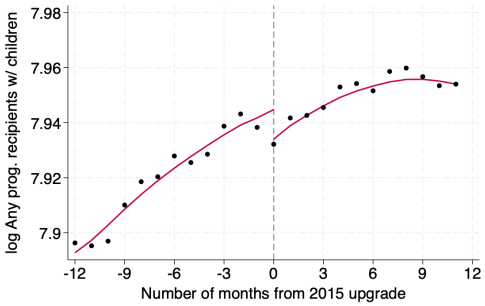
Notes: Panel (a) plots raw monthly trends in log program enrollment (among those in households with children) for the average zip code based on distance to assigned county field office. The level of zip codes less than 7 miles away is normalized to the level of zip codes more than 7 miles away in September 2012. Panel (b) illustrates event-study estimates of $\{\beta_k^g\}$ from Equation (2). Shaded regions are 95% confidence intervals based on robust standard errors clustered at the county level. In all panels, zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

Figure A.14: Visual Evidence of Discontinuities for Single and Program Receipt

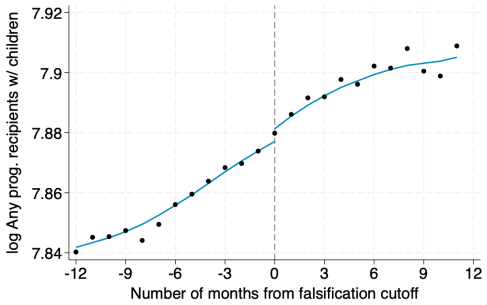
(a) Any Program: Launch Treatment Window, October 2012



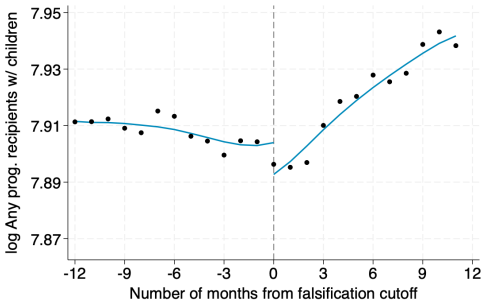
(b) Any Program: Upgrade Treatment Window, January 2015



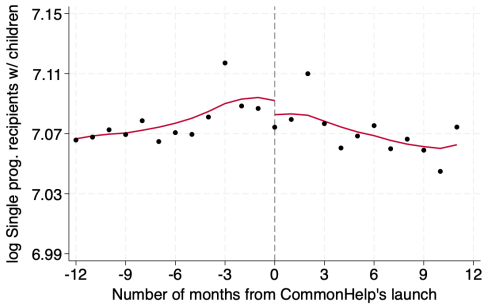
(c) Any Program: Launch Control Window, October 2011



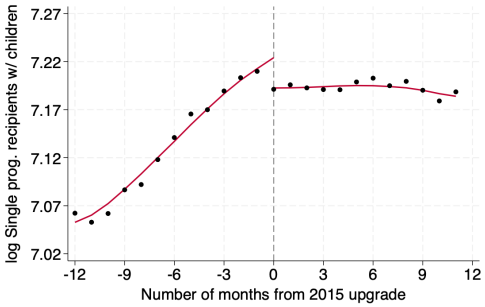
(d) Any Program: Upgrade Control Window, January 2014



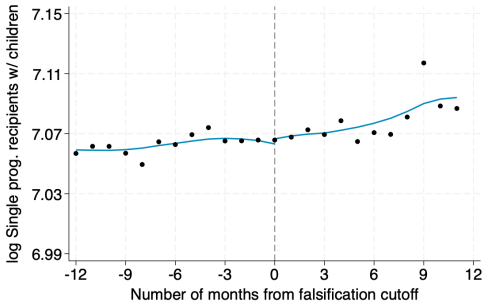
(e) Single Program: Launch Treatment Window, October 2012



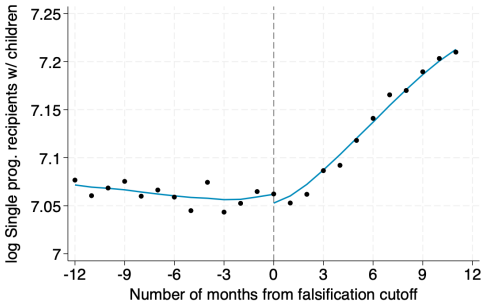
(f) Single Program: Upgrade Treatment Window, January 2015



(g) Single Program: Launch Control Window, October 2011

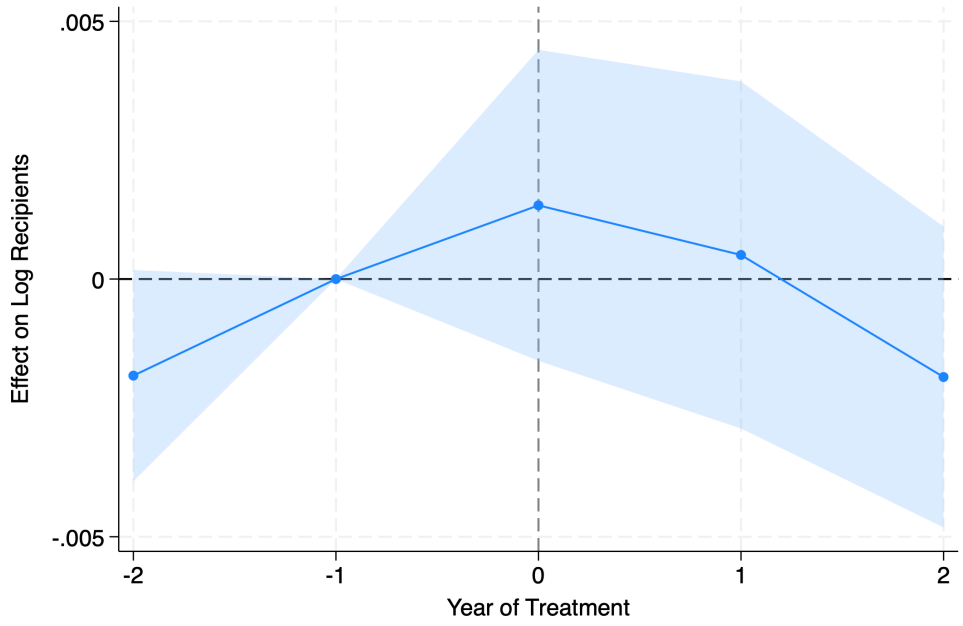


(h) Single Program: Upgrade Control Window, January 2014



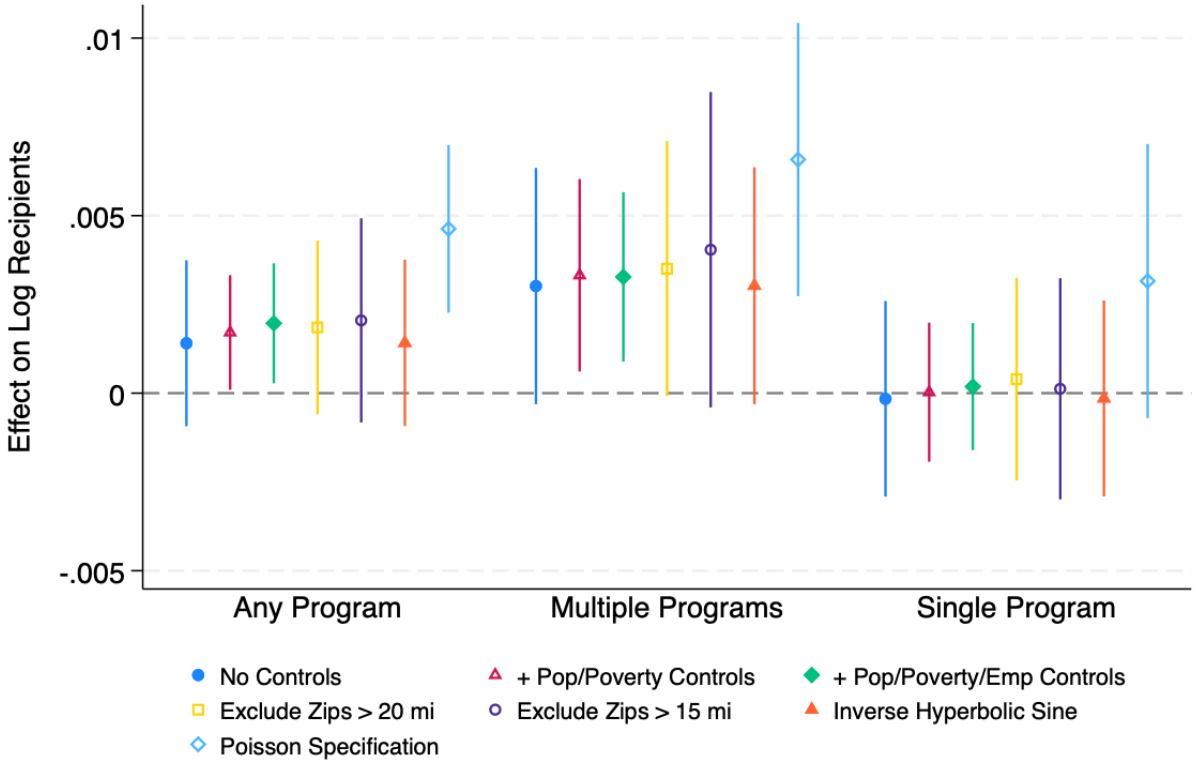
Notes: See notes of Figure 7.

Figure A.15: Event-Study Estimates on Participation in UI (Falsification Check)



Notes: This figure illustrates event-study estimates of CommonHelp's per-mile effect on Unemployment Insurance (UI) receipt. This should act as a falsification check because UI is not a program administered by CommonHelp. Because we only have calendar-year information on UI receipt, we adopt a modified version of Equation (2), where the event-time dummies are in terms of calendar years rather than quarters. Specifically, 2013 is the initial year of treatment (event-time 0), and our time period of interest spans 2011 (event time -2) through 2015 (event time 2). Shaded regions are 95% confidence intervals based on robust standard errors clustered at the county level. Zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

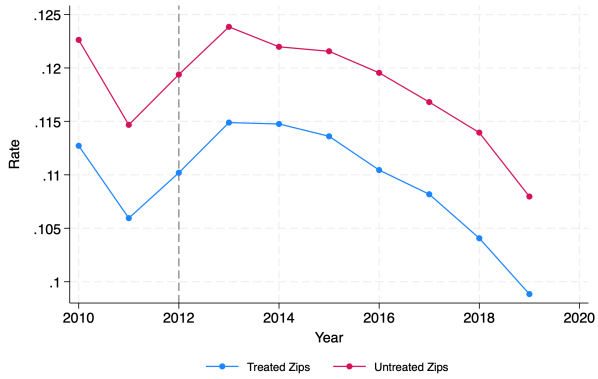
Figure A.16: Sensitivity of DD Estimates to Controls and Alternative Specifications



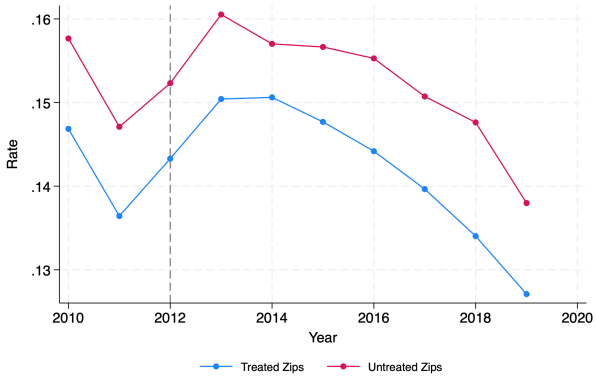
Notes: This figure shows the sensitivity of our pre-post difference-in-differences (DD) estimates to the inclusion and choice of covariates (taken from ACS 5-year estimates) and to alternative functional form assumptions. We show estimates separately for three outcomes, reflecting the counts of any program recipients, multiple program recipients, and single program recipients (logged unless specified otherwise). The blue circles show estimates with no zip-level covariates, the red triangles show estimates after incorporating zip- and time-varying population and poverty controls (log total population, share Black non-Hispanic, share age 65+, share female, and poverty rate), and the green diamonds show estimates after additionally incorporating zip- and time-varying employment controls (unemployment and out-of-labor-force rates). Keeping the full set of controls, the yellow squares and purple circles show estimates that exclude zips that are farthest from their assigned county office (more than 20 miles and 15 miles away, respectively). The final two sets of estimates continue to keep the full set of controls and zips but rely on alternative transformations of the outcome (other than log) that can still allow coefficients to be approximately interpreted in percent terms. Specifically, the orange triangles apply an inverse hyperbolic sine transformation to the outcome, while the blue diamonds employ a Poisson specification. Confidence bands at the 95% confidence level are based on robust standard errors clustered at the county level. Zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

Figure A.17: Raw Covariate Trends Between Zips with Distance Below/Above 7 Miles

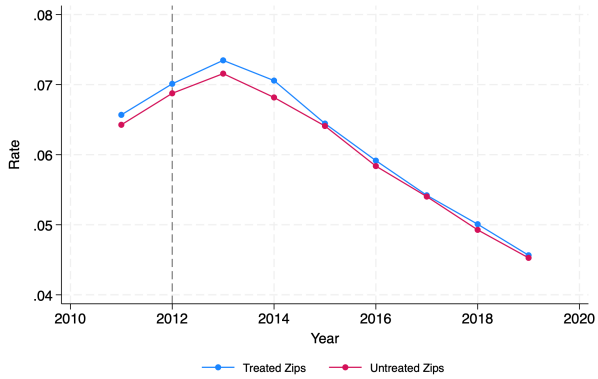
(a) Poverty Rates



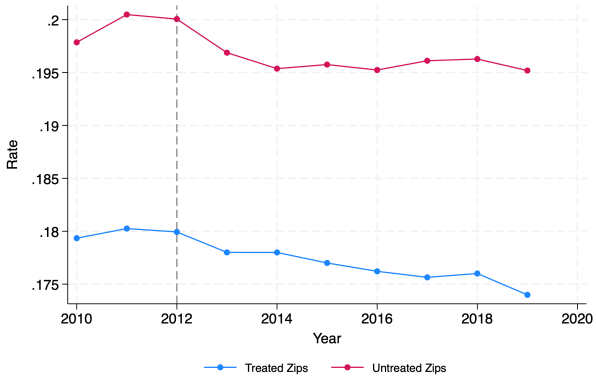
(b) Child Poverty Rates



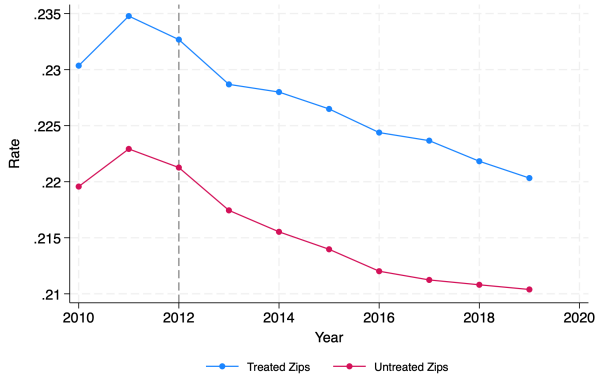
(c) Unemployment Rate



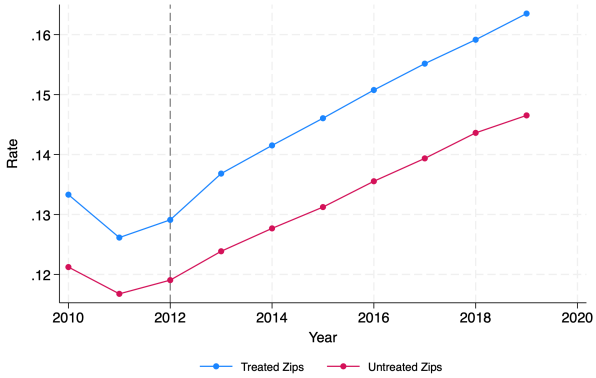
(d) Share of Black non-Hispanic Population



(e) Share of Population Age 0–17



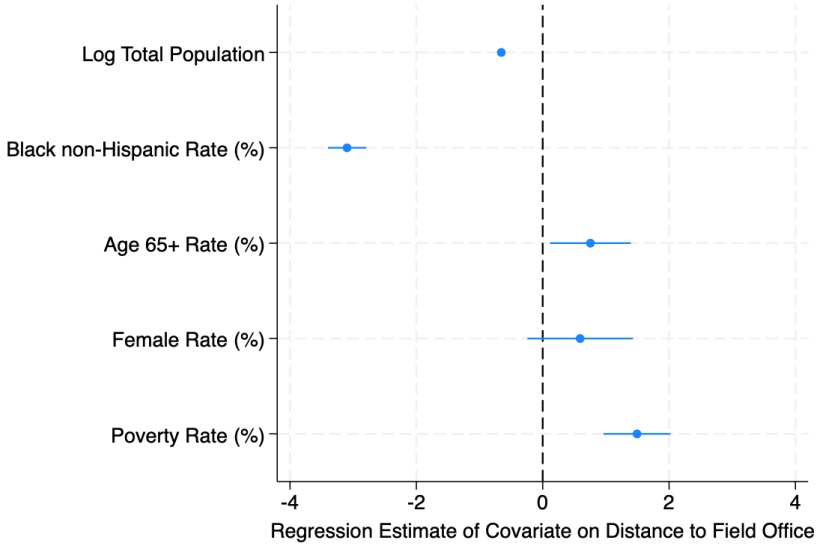
(f) Share of Population Age 65+



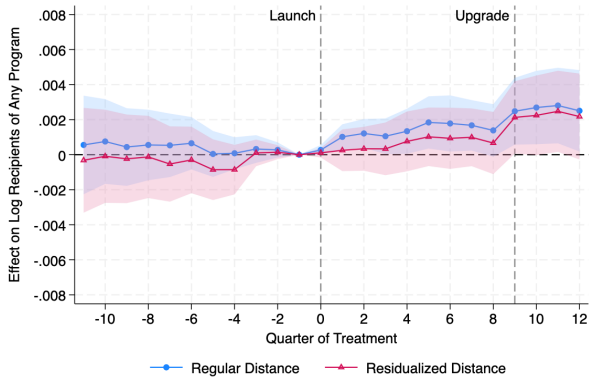
Notes: These figures show raw trends in ZCTA-level covariates (taken from ACS 5-year estimates) for zip codes that are more than 7 miles from their assigned field office (treated) and less than 7 miles from their assigned field office (untreated). Estimates are weighted by annual zip-level population counts.

Figure A.18: Robustness of Event-Study Estimates to Differences in Baseline Composition

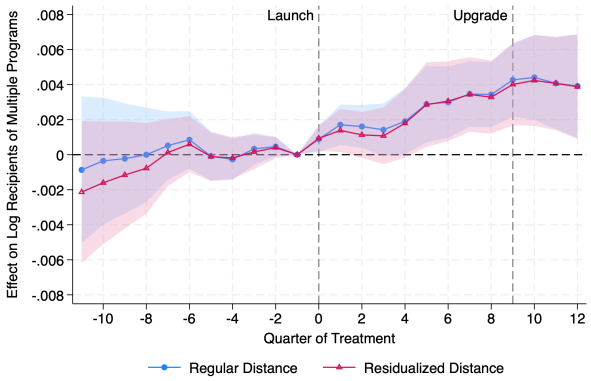
(a) Correlates of Distance to Field Office



(b) Event-Study: Any Program

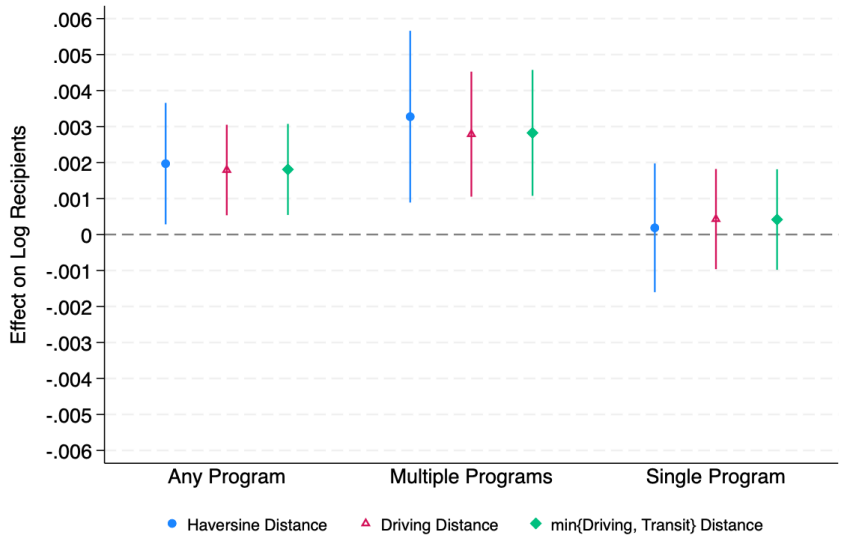


(c) Event-Study: Multiple Programs



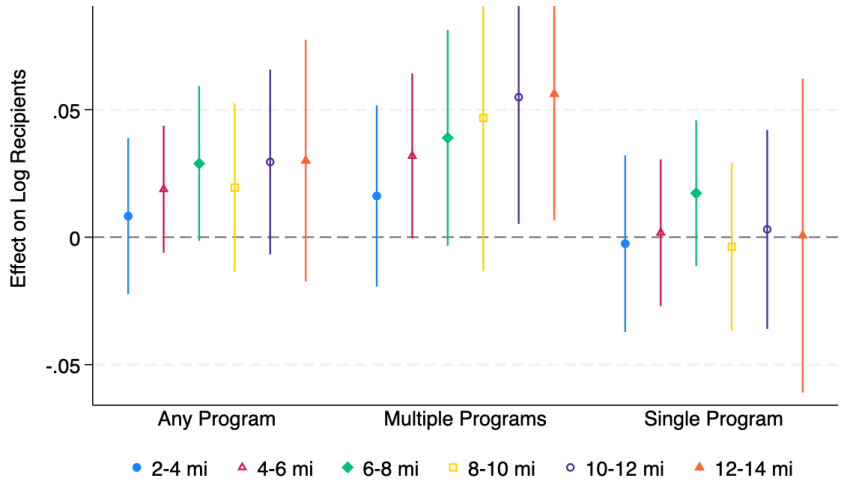
Notes: Panel (a) reports coefficients on various zip-level covariates in a regression of distance to field office on those covariates. The sample period consists of October 2009-September 2012 (prior to CommonHelp). Confidence bands at the 95% confidence level are based on robust standard errors. Panels (b) and (c) illustrate estimates of $\{\beta_k^g\}$ from Equation (2), where the outcomes are log participants in any program (Panel a) and multiple programs (Panel b) and the regressors reflect residualized distance based on the covariates in Appendix Figure A.18a. Shaded regions are 95% confidence intervals based on robust standard errors clustered at the county level. Zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

Figure A.19: Sensitivity of DD Estimates to Alternative Distance Metrics



Notes: This figure shows the sensitivity of our pre-post difference-in-differences (DD) estimates to different definitions of calculations of distance from field office. Our main estimates (blue circles) use Haversine distance, but we also show robustness checks using driving distance and the minimum of driving and transit distance (both calculated using Google API). Confidence bands at the 95% confidence level are based on robust standard errors clustered at the county level. Zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

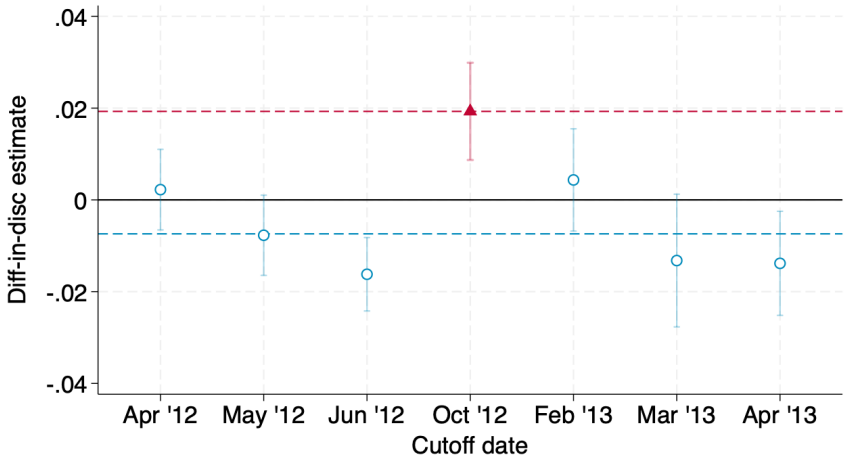
Figure A.20: CommonHelp's Impact on Enrollment Increases in Distance to Field Office



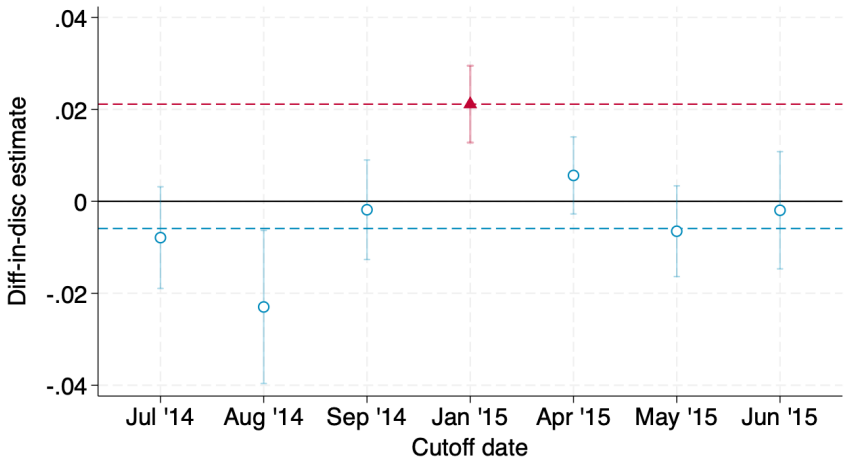
Notes: This figure reports estimates of β^g based on a pre-post version of Equation (2) in which D_z is replaced with $\mathbb{1}\{D_z \in [d_0, d_1)\}$ that is estimated using zip codes with $D_z < 2$ or $D_z \in [d_0, d_1)$ for different choices of d_0, d_1. Vertical lines are 95% confidence intervals based on robust standard errors clustered at the county level. Zip codes are weighted by the average count of recipients of any program between October 2009–September 2012.

Figure A.21: Placebo DRD Estimates on Multiple Program Receipt

(a) Launch Estimates



(b) Upgrade Estimates

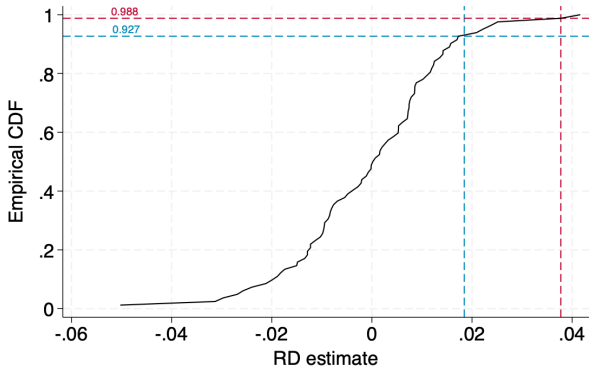


Notes: This figure plots DRD placebo estimates, where placebo cutoff months are 4–6 months away from the treatment month (to avoid having the effect at the true cutoff to contaminate the placebo effect in its estimation window). Besides changing the cutoff month, the DRD regression specifications are identical as before. Vertical capped bars are 95% confidence intervals based on standard errors clustered at the county level. The left (right) panel excludes the RD estimate around the upgrade (launch) date. Dashed red lines indicate the DRD estimate; dashed blue lines indicate the mean DRD estimates from the placebo estimates.

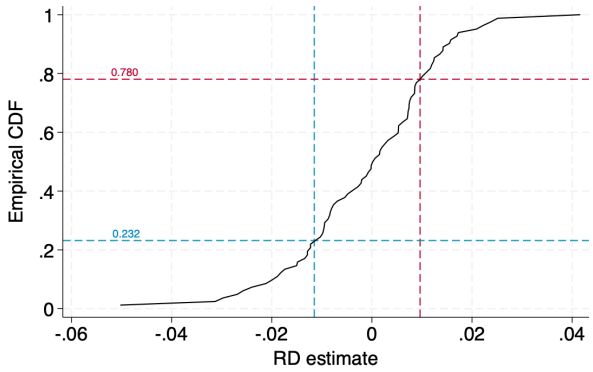
Figure A.22: Placebo RD Effect Estimates on Multiple Program Receipt

(a) Empirical CDFs of Placebo RD Estimates

(i) Launch Estimates

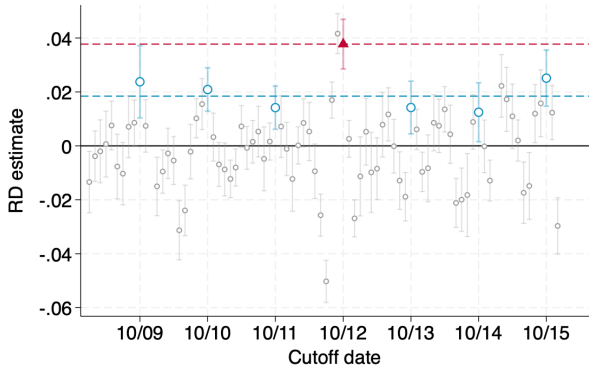


(ii) Upgrade Estimates

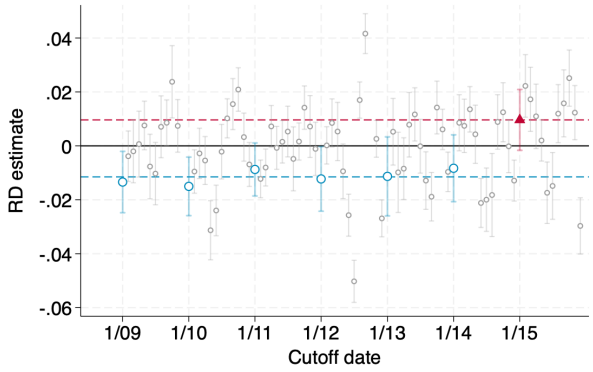


(b) All Placebo RD Estimates

(i) Launch Estimates



(ii) Upgrade Estimates



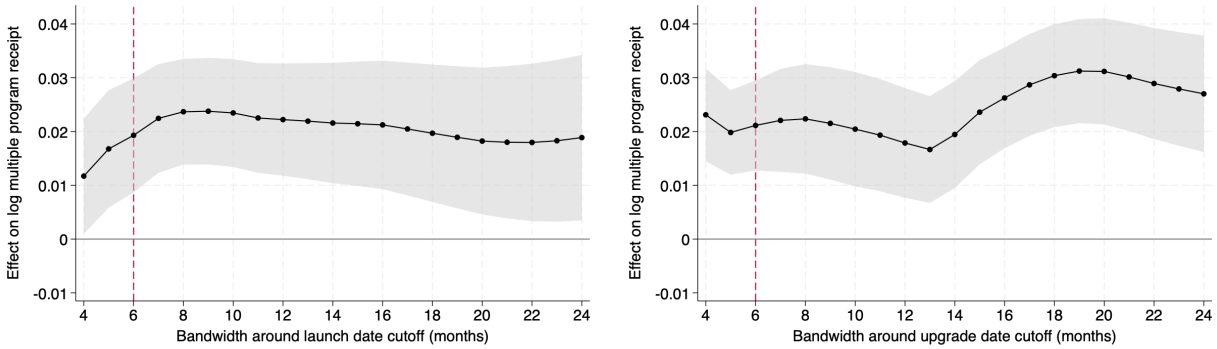
Notes: The top panels plot empirical cumulative distribution functions of RD placebo estimates, where placebo cutoff dates are all months between January 2009–December 2015. The bottom panels plot RD placebo estimates, where placebo cutoff dates are all months between January 2009–December 2015. Vertical capped bars are 95% confidence intervals based on standard errors clustered at the county level. The left (right) panels exclude the RD estimate around the upgrade (launch) date. Dashed red lines indicate the RD estimate corresponding to the main cutoff date; dashed blue lines indicate the mean RD estimate from the control windows.

Figure A.23: Robustness of DRD Effect Estimates on Multiple Program Receipt to Alternative Bandwidths and Distance Thresholds

(a) Varying Bandwidth

(i) Launch Estimates

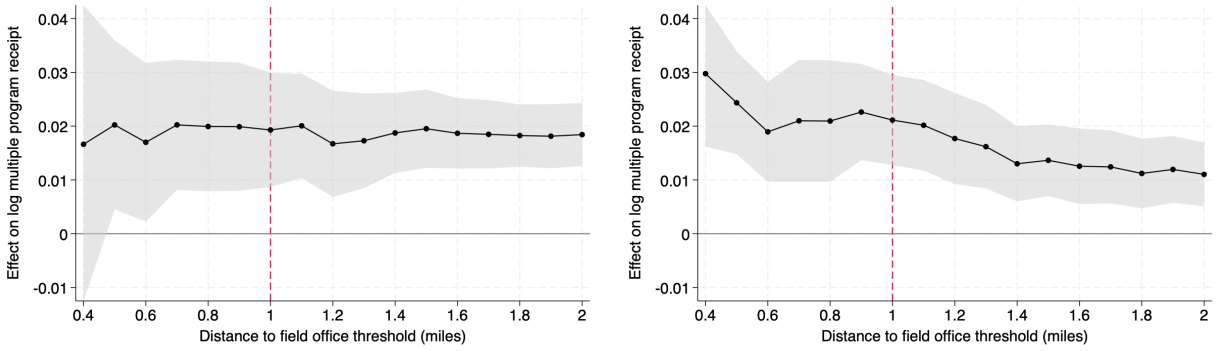
(ii) Upgrade Estimates



(b) Varying Distance Threshold

(i) Launch Estimates

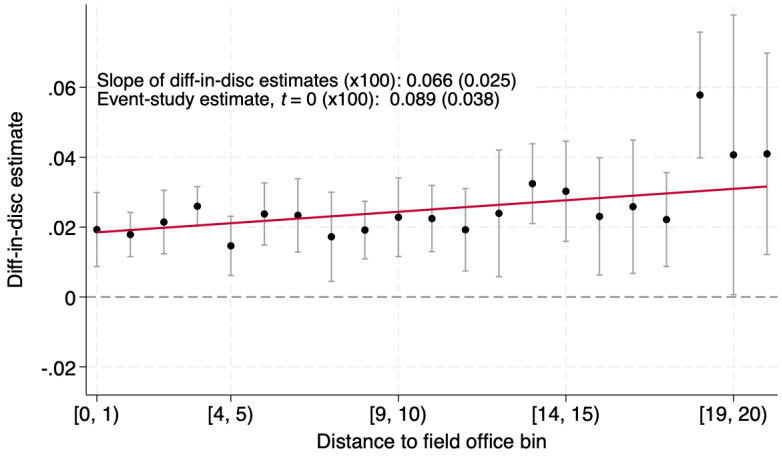
(ii) Upgrade Estimates



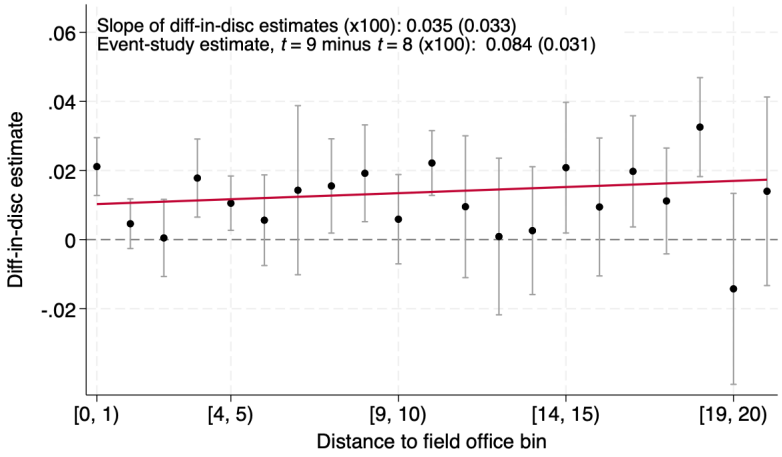
Notes: The top panels plot DRD estimates for different choices of bandwidths from the date cutoff. The bottom panels plot DRD estimates for different distance thresholds $\bar{d}$ that selects the analytical sample. Dashed red lines indicate the selected bandwidth and distance threshold using our preferred specification. Shaded areas are 95% confidence intervals based on standard errors clustered at the county level.

Figure A.24: DRD Estimates on Multiple Program Receipt by Distance Bin (Total Effects)

(a) Launch

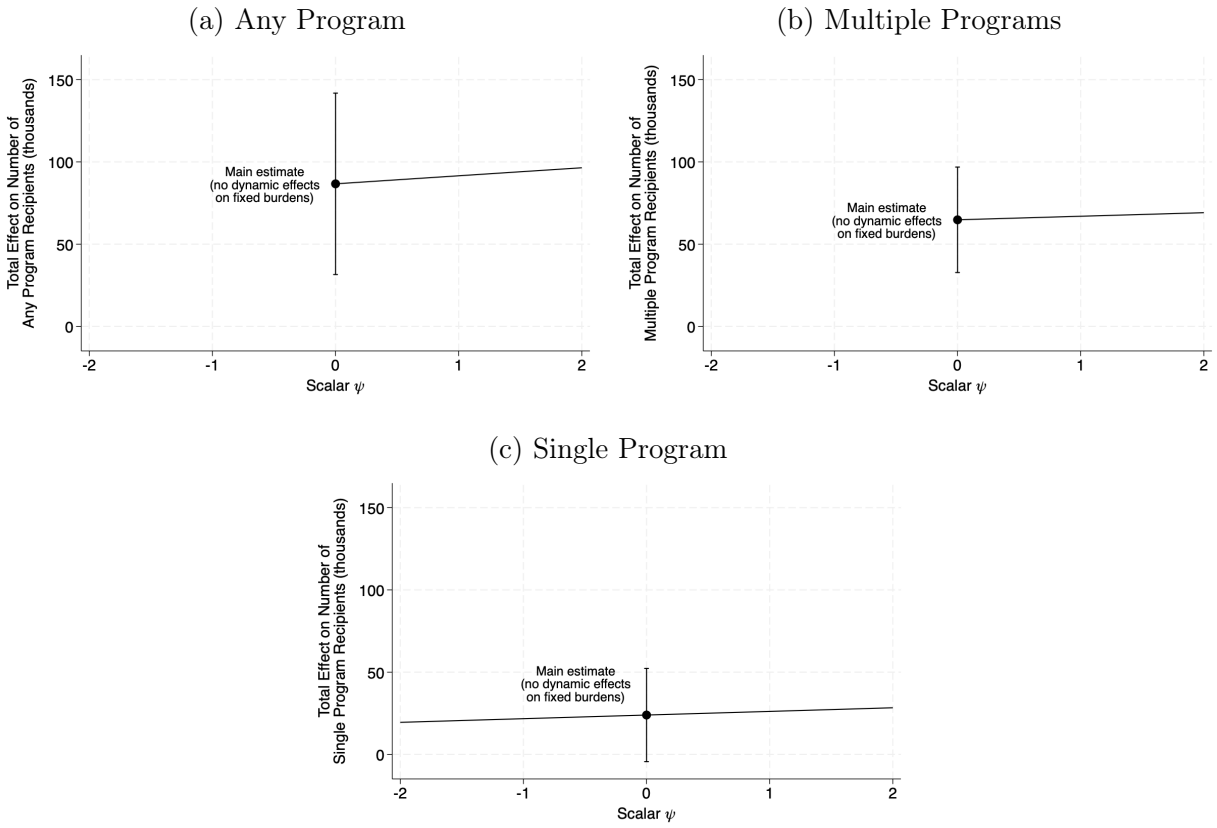


(b) Upgrade



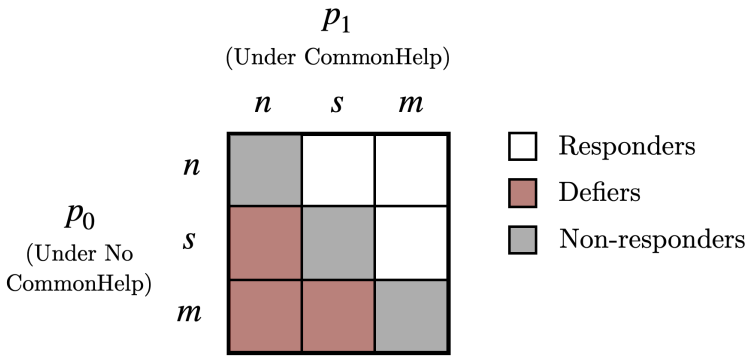
Notes: This figure plots DRD estimates among zip codes lying in different distance bins. All regressions employ triangular kernels. The red line is the best linear fit of the estimates, weighting estimates by their inverse squared standard errors. Its slope estimate is reported in the text box along with the estimate of β_0 (left panel) and $\beta_9 - \beta_8$ (right panel) from Equation (2) for comparison; standard errors are in parentheses. Gray vertical capped lines are 95% confidence intervals based on robust standard errors clustered at the county level.

Figure A.25: Robustness to Allowing for Dynamic Effects on Reduced Fixed Burdens



Notes: This figure plots estimates of ΔY based on different assumptions of the growth rate of impacts on fixed burdens, which is parameterized as $\psi \cdot \hat{g}_{\beta g}$ where ϕ is some constant and $\hat{g}_{\beta g}$ is the estimated growth rate of impacts on dynamic burdens based on the event-study DD estimates. The bottom panel allows $\psi < 0$ since single program receipt may have non-monotonic response patterns. Capped vertical lines are 95% confidence intervals based on 250 county-level block bootstrap replications.

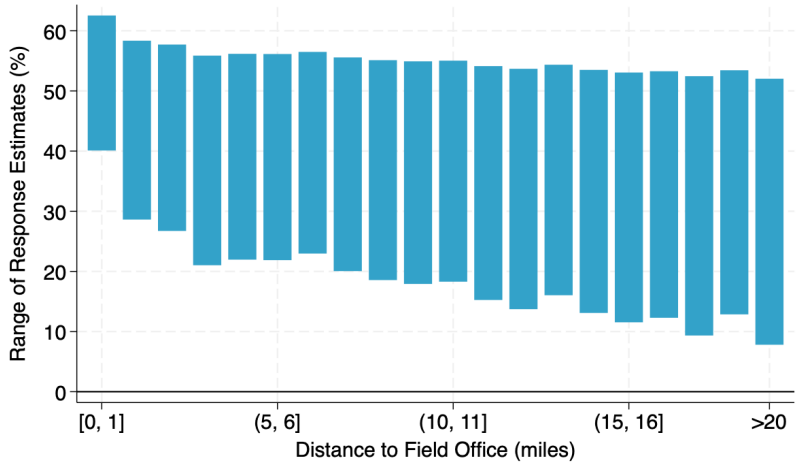
Figure A.26: Permissible Responder Types Under the Monotonicity Assumption



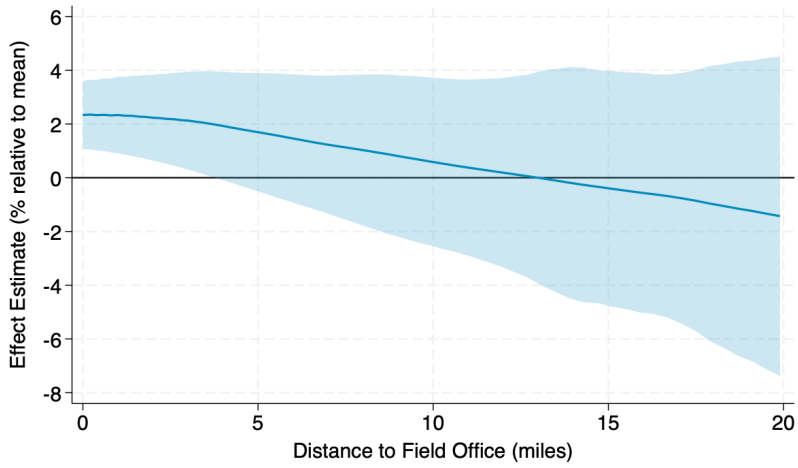
Notes: This figure visualizes the content of the monotonicity assumption for the responder type analysis. The white and gray squares indicate the permissible latent types of responders. See text for more details.

Figure A.27: Multiple-Program Targeting Effects by Distance to Field Office

(a) Range of Shares of Intensive-Margin Responses



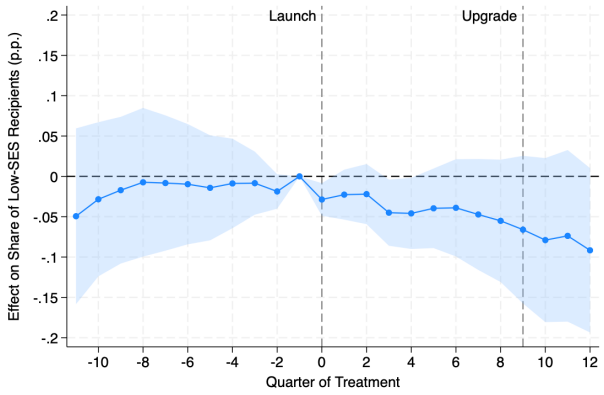
(b) Impact on Share of Low-Socioeconomic Status Recipients



Notes: The top panel reports estimated ranges of intensive-margin multiple program responders as a share of total responders among zip codes grouped by 1-mile distance bins; see main text for details. The bottom panel plots estimated effects on the share of multiple program recipients with low-socioeconomic status (SES) (defined in the main text). Estimates of the total effect by distance (Equation B.3) are divided by the share of low-SES multiple program recipients at each distance, estimated via a local linear regression of low-SES shares measured quarter prior to CommonHelp's launch on distance. Shaded region is 95% confidence intervals based on county-clustered standard errors.

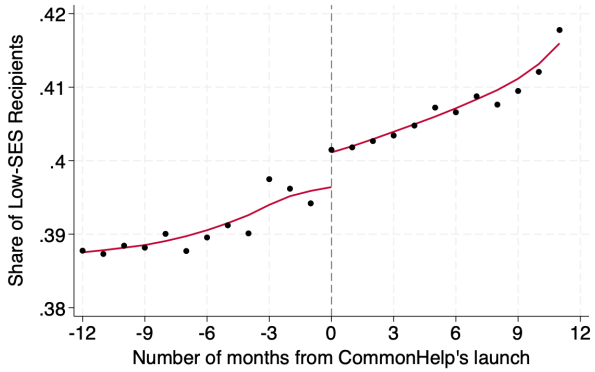
Figure A.28: Visual Evidence of Multiple-Program Targeting on Observables Analysis: Share of Low-Socioeconomic Status (SES) Recipients

(a) Effects from Geographic Burdens: Event-Study Estimates

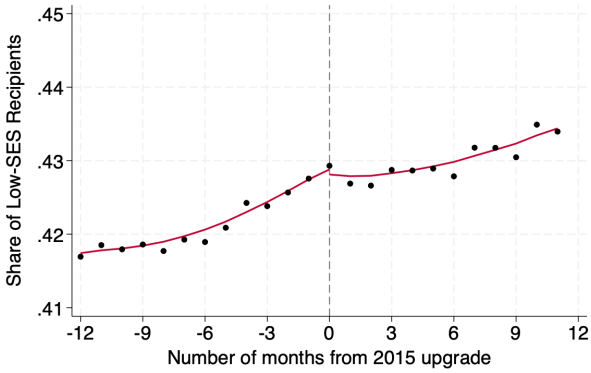


(b) Effects from Fixed Burdens: Regression Discontinuity Evidence

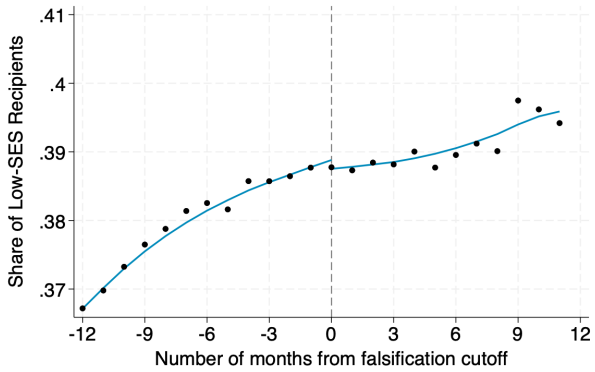
(i) Launch Treatment Window: October 2012



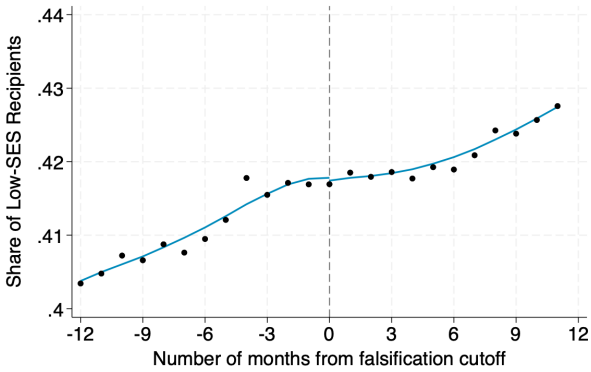
(ii) Upgrade Treatment Window: January 2015



(iii) Launch Control Window: October 2011

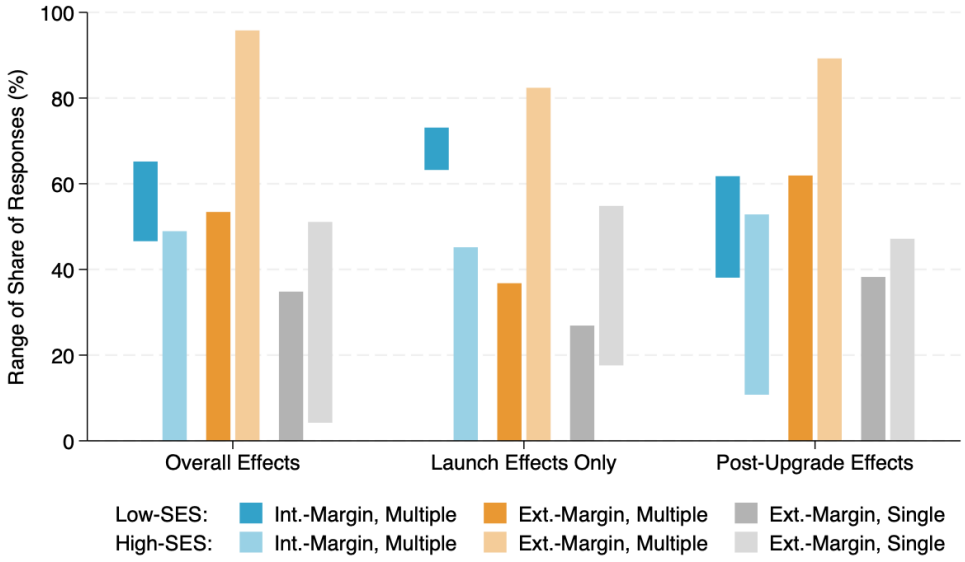


(iv) Upgrade Control Window: January 2014



Notes: The top panel plots estimates from an event-study analog of Equation (B.1). Shaded area is 95% confidence intervals based on robust standard errors clustered at the county level. The bottom panels plot visual discontinuity evidence based on Equation (B.2). See notes of Figure 7 for details.

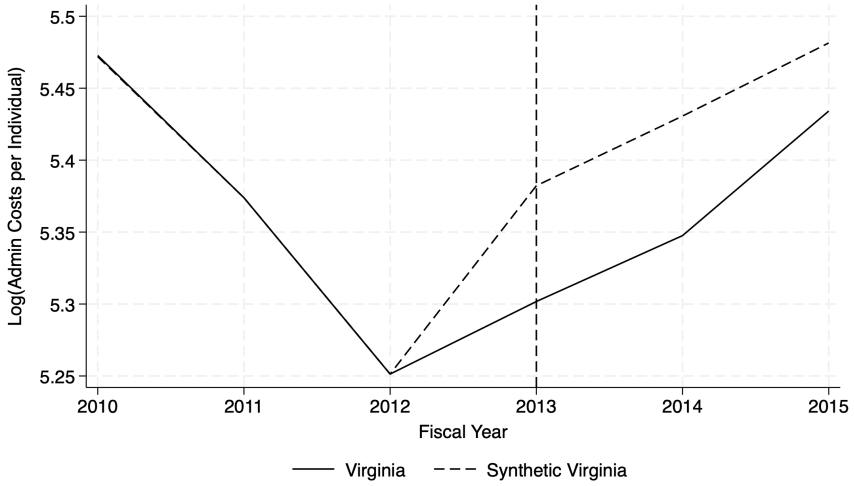
Figure A.29: Share of Intensive- and Extensive-Margin Responses, by SES group



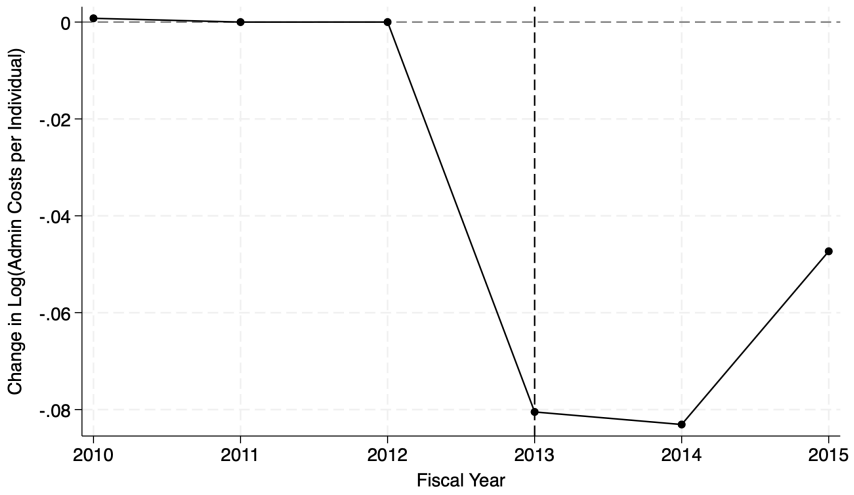
Notes: This figure plots the estimated percent range of responses (among all responses to CommonHelp), separately by SES group, based on Equation (8). See main text for details.

Figure A.30: Synthetic Control Estimates for Log SNAP Admin Costs Per Recipient

(a) Virginia vs. Synthetic Virginia Trends

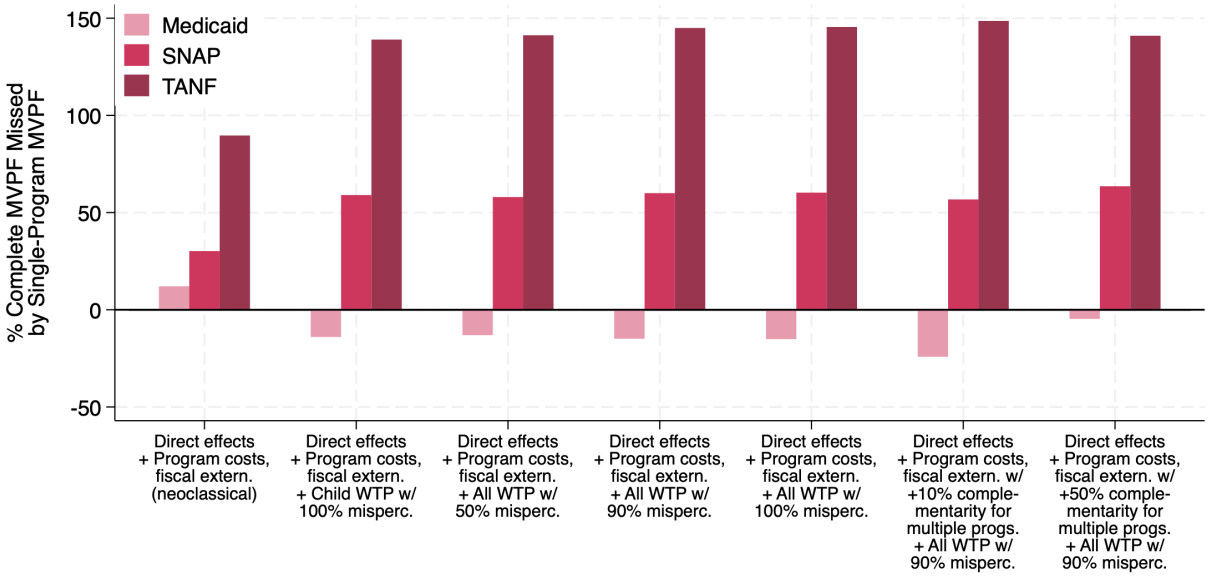


(b) Difference in Trends



Notes: These figures use the synthetic control method to compare changes in log SNAP administrative costs per individual recipient in Virginia to changes in other states. Panel (a) compares the fiscal year trends in log SNAP admin costs per recipient between Virginia and synthetic Virginia, with the synthetic control group constructed based on matching to Virginia on the following set of pre-treatment covariates: the outcome, population, and unemployment rate for each of the fiscal quarters preceding treatment; population cuts by gender, race, and age; median income; and labor force participation rate. Panel (b) plots the fiscal year differences in log SNAP admin costs per recipient between Virginia and synthetic Virginia.

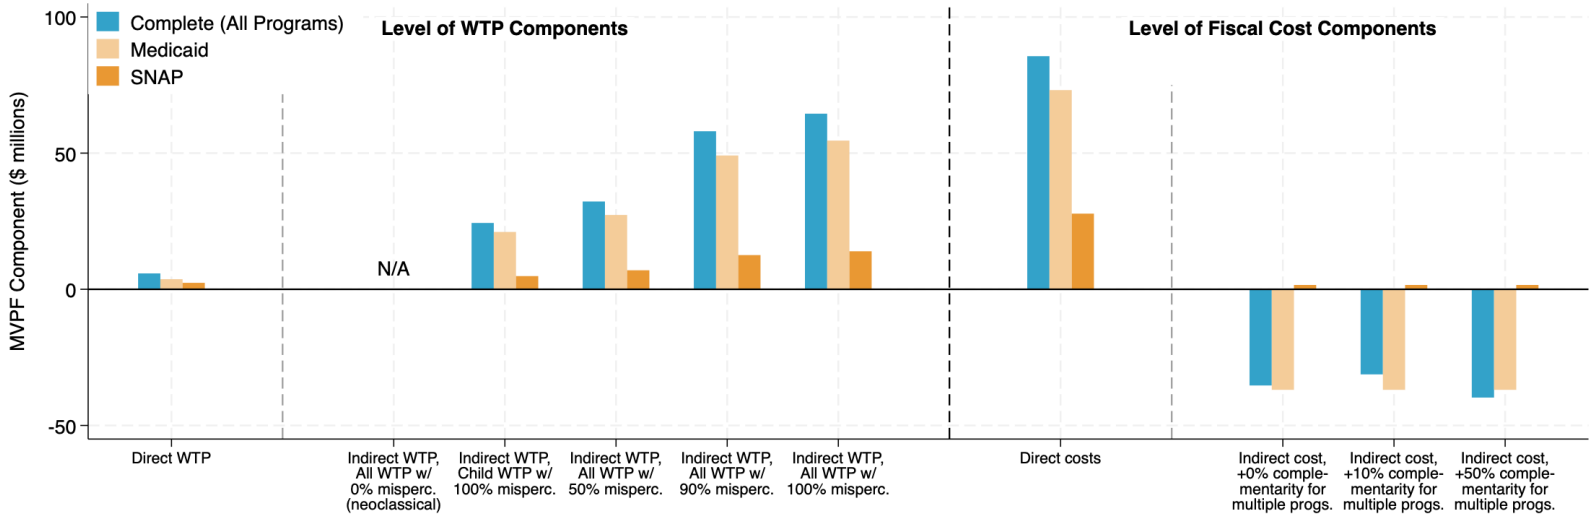
Figure A.31: Differences in MVPF between Single vs. All Programs, by Program and Modeling Assumption



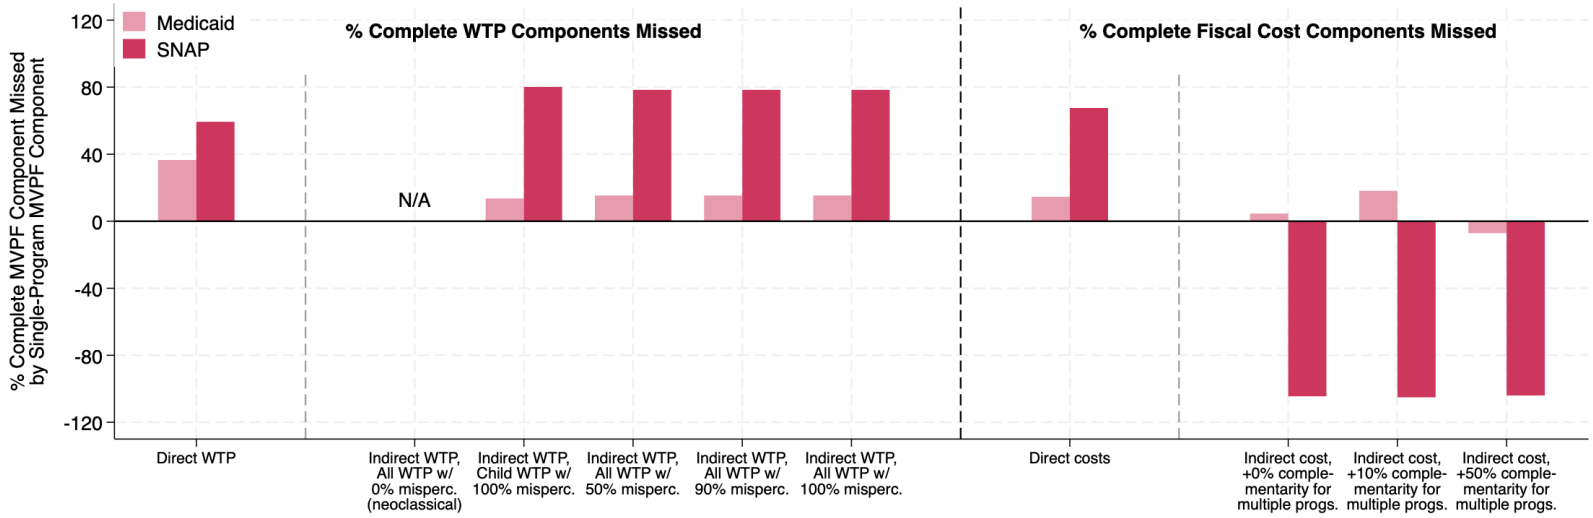
Notes: This figure plots the percent of the multiple-program MVFP that is missing from each single-program MVPF.

Figure A.32: Components of the MVPF of CommonHelp: Multiple vs. Single Programs

(a) Levels of Components



(b) Differences between Single vs. Multiple Programs



Notes: Panel (a) plots estimates of the components underlying the MVPF, for multiple programs and by Medicaid and SNAP, across a range of modeling assumptions. Components are based on direct and indirect effects on willingness-to-pay (WTP) and fiscal costs; see main text for details. Panel (b) plots the percent of the multiple-program MVPF component that is missing from each respective single-program MVPF component.

Table A.1: Overview of SNAP, Medicaid, and TANF Eligibility Criteria in Virginia (2012)

	SNAP	Medicaid	TANF
Demographic elig.	All families	Children ages 0–18; parents, caregivers, and pregnant women; elderly; disabled*	Families with cohabiting children ages 0–17 or with children age 18 enrolled in school
Income limits	Gross income $\leq$ 130% FPL (no limit for elderly or disabled); net income $\leq$ 100% FPL†	Children and pregnant women $\leq$ 200% FPL; elderly and disabled $\leq$ 80% FPL; parents $\leq$ 33% FPL‡	Gross income $\leq$ 35–41% FPL**
Asset limits	Non-elderly/non-disabled $\leq$ \$2,000; Elderly or disabled $\leq$ \$3,250	\$2,000 excluding home, personal belongings	None
Work requirements	None§	None	30 hours per week¶
Time limits	None	None	5 years

Notes: Based on 2012 program rules.
*ABAWDs became eligible for Medicaid after January 2019.
†Gross income limits for SNAP increased to 200% and asset limits were eliminated after July 2021.
‡Income limits for Medicaid increased to 138% for elderly/disabled, parents, and ABAWDs after Jan 2019.
**Population-weighted average across different regions in Virginia. Range varying based on case size.
§ABAWDs ages 18–49 faced work requirements for SNAP after October 2013.
¶Exemptions granted to minors, 18-year olds in secondary school, the disabled, elderly, and primary caretakers of a child under 12 months or a disabled family member.

Table A.2: Summary Effects on Enrollment Growth, by Individual Program

	(1)	(2)	(3)	(4)	(5)
	Program type			Number of observations	
	SNAP	Medicaid	TANF	Zip codes	Zip-county-months
A. Effects from Reducing Geographic Burdens (DD)					
Overall effect (%/miles)	0.187	0.294***	0.508**	982	85,176
	(0.123)	(0.100)	(0.229)		
Launch effect (%/miles)	0.136	0.193**	0.435*	982	85,176
	(0.117)	(0.085)	(0.230)		
Upgrade effect (%/miles)	0.167**	0.330***	0.237*	982	85,176
	(0.082)	(0.116)	(0.121)		
Outcome mean (count/zip)	437.5	480.7	56.2	982	85,176
B. Effects from Reducing Fixed Burdens (DRD)					
Overall effect (%)	2.647***	0.538	−7.832***	47	4,200
	(0.455)	(0.367)	(2.773)		
Launch effect (%)	2.760***	−0.045	−7.119***	47	4,200
	(0.434)	(0.291)	(2.498)		
Upgrade effect (%)	−0.366	1.893***	−2.318	47	4,200
	(0.385)	(0.452)	(1.697)		
Outcome mean (count/zip)	794.6	939.7	90.6	47	4,200
C. Total Average Effects ($D_z = \bar{D}$)					
Overall effect (%)	3.866***	2.456***	−4.518	982	85,176
	(0.822)	(0.804)	(3.037)		
Launch effect (%)	3.644***	1.211*	−4.280	982	85,176
	(0.765)	(0.653)	(2.791)		
Upgrade effect (%)	0.723	4.049***	−0.773	982	85,176
	(0.607)	(0.961)	(1.799)		
Outcome mean (count/zip)	437.5	480.7	56.2	982	85,176

Notes: See notes of Table 2 for details.

Table A.3: Summary Effects on Enrollment Growth, by Program Bundle

	(1)	(2)	(3)	(4)	(5)	(6)
	Program type				Number of observations	
	SNAP only	Medicaid only	SNAP & Medicaid	All programs	Zip codes	Zip-county -months
A. Effects from Reducing Geographic Burdens (DD)						
Overall effect (%/miles)	−0.172	0.162	0.202*	0.521**	982	85,176
	(0.213)	(0.150)	(0.108)	(0.263)		
Launch effect (%/miles)	−0.113	0.185	0.128	0.429*	982	85,176
	(0.181)	(0.115)	(0.102)	(0.243)		
Upgrade effect (%/miles)	−0.192	−0.075	0.243**	0.299*	982	85,176
	(0.170)	(0.199)	(0.098)	(0.155)		
Outcome mean (count/zip)	114.7	157.0	272.8	45.8	982	85,176
B. Effects from Reducing Fixed Burdens (DRD)						
Overall effect (%)	3.226***	−5.533***	4.021***	−7.236**	47	4,200
	(0.947)	(1.232)	(0.675)	(3.124)		
Launch effect (%)	5.135***	−6.005***	3.272***	−6.581**	47	4,200
	(0.922)	(1.338)	(0.591)	(2.663)		
Upgrade effect (%)	−6.203***	1.533**	2.436***	−2.127	47	4,200
	(1.492)	(0.593)	(0.497)	(2.283)		
Outcome mean (count/zip)	197.5	343.8	514.8	74.1	47	4,200
C. Total Average Effects ($D_z = \bar{D}$)						
Overall effect (%)	2.104	−4.477***	5.340***	−3.835	982	85,176
	(1.546)	(1.369)	(1.022)	(3.150)		
Launch effect (%)	4.397***	−4.799***	4.103***	−3.780	982	85,176
	(1.426)	(1.258)	(0.915)	(2.726)		
Upgrade effect (%)	−7.454***	1.047	4.021***	−0.176	982	85,176
	(1.830)	(1.356)	(0.881)	(2.365)		
Outcome mean (count/zip)	114.7	157.0	272.8	45.8	982	85,176

Notes: See notes of Table 2 for details.

Table A.4: Total Effects of CommonHelp on Program Enrollment, by Reform Period

	(1)	(2)	(3)	(4)	(5)	(6)
	Any		Multiple		Single	
	ΔY	%	ΔY	%	ΔY	%
A. Overall Effects						
1. Total: $g_z = \bar{\beta}^g \cdot D_z + \bar{\beta}^f$	58,424***	2.206	63,885***	4.642	−13,015	−1.048
	(20,953)		(15,503)	[2.360]	(10,115)	[−0.481]
2. Geographic only: $g_z = \bar{\beta}^g \cdot D_z$	34,278**	1.283	29,101***	2.062	1,535	0.125
	(17,876)		(12,090)	[1.075]	(8,144)	[0.057]
3. Fixed only: $g_z = \bar{\beta}^f$	24,764***	0.923	36,208***	2.579	−14,587***	−1.173
	(8,084)		(9,202)	[1.338]	(6,128)	[−0.539]
B. Launch Effects Only						
1. Total: $g_z = \beta^g \cdot D_z + \beta^f$	44,539***	1.673	47,771***	3.431	−6,431	−0.520
	(18,759)		(12,988)	[1.765]	(9,096)	[−0.238]
2. Geographic only: $g_z = \beta^g \cdot D_z$	21,717*	0.809	21,308**	1.502	1,252	0.102
	(15,925)		(11,001)	[0.787]	(6,763)	[0.046]
3. Fixed only: $g_z = \beta^f$	23,191***	0.864	27,253***	1.929	−7,698*	−0.622
	(7,353)		(7,877)	[1.007]	(5,495)	[−0.284]
C. Post-Upgrade Effects						
1. Total: $g_z = (\beta^g + \beta^{g15}) \cdot D_z + (\beta^f + \beta^{f15})$	69,913***	2.609	76,498***	5.895	−31,956**	−2.336
	(21,794)		(17,584)	[2.783]	(13,976)	[−1.162]
2. Geographic only: $g_z = (\beta^g + \beta^{g15}) \cdot D_z$	42,038**	1.553	24,989***	1.852	1,005	0.075
	(18,120)		(9,068)	[0.909]	(10,526)	[0.037]
3. Fixed only: $g_z = \beta^f + \beta^{f15}$	28,739***	1.056	53,388***	4.042	−33,011***	−2.411
	(11,488)		(12,277)	[1.942]	(10,468)	[−1.201]
Outcome mean (count/zip)	600.6		318.6		271.6	
N: zip codes	982		982		982	
N: zip-county-months	85,176		85,176		85,176	
Share of recipients (%)	100.0		53.2		45.4	

Notes: See notes to Table 3 for details. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: DRD Effect Estimates are Robust to Alternative Specifications

	(1)	(2)	(3)	(4)	(5)
	Main	Covariates	Distance kernel	Parametric linear	Parametric quadratic
A. Any program					
Overall effect (%)	0.923***	0.900***	0.917***	0.923***	1.359***
	(0.279)	(0.278)	(0.270)	(0.279)	(0.342)
Launch effect (%)	0.864***	0.849***	0.799***	0.961***	1.229***
	(0.248)	(0.249)	(0.249)	(0.270)	(0.327)
Upgrade effect (%)	0.192	0.166	0.384	0.196	0.422
	(0.290)	(0.289)	(0.424)	(0.282)	(0.305)
Outcome mean (count/zip)	1,146.7	1,146.7	1,146.7	1,146.7	1,152.1
B. Multiple programs					
Overall effect (%)	2.579***	2.830***	2.427***	2.579***	3.236***
	(0.601)	(0.548)	(0.746)	(0.601)	(0.635)
Launch effect (%)	1.929***	2.183***	1.758**	2.193***	2.371***
	(0.541)	(0.482)	(0.701)	(0.541)	(0.553)
Upgrade effect (%)	2.113***	2.105***	2.175***	2.298***	2.812***
	(0.427)	(0.430)	(0.465)	(0.525)	(0.611)
Outcome mean (count/zip)	588.9	588.9	588.9	588.9	595.8
C. Single program					
Overall effect (%)	−1.173**	−1.228**	−1.264***	−1.173**	−1.180**
	(0.477)	(0.476)	(0.450)	(0.477)	(0.577)
Launch effect (%)	−0.622	−0.686	−0.796	−0.902*	−0.508
	(0.479)	(0.478)	(0.483)	(0.458)	(0.571)
Upgrade effect (%)	−1.789***	−1.764***	−1.519**	−2.210***	−2.184***
	(0.511)	(0.515)	(0.625)	(0.673)	(0.771)
Outcome mean (count/zip)	541.3	541.3	541.3	541.3	540.4
Number of zip codes	47	43	47	47	47
Number of counties	40	37	40	40	40

Notes: This table examines DRD estimates from alternative specifications. Column (1) reports our main results. Column (2) includes time-averaged zip-level covariates from our DD specification (log total population; population shares of women, blacks, and elderly adults; and poverty, unemployment, and labor force participation rates). Column (3) additionally weighs zip codes by their proximity to their field office through a triangular kernel. Column (4) applies parametric linear fits instead of local linear fits. Columns (5) and (6) apply local quadratic and cubic fits, respectively, using a 12-month bandwidth. Standard errors in parentheses are clustered at the county level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Estimated Ranges of Shares of Intensive vs. Extensive Margin Responses, by Reform Period

	(1)	(2)	(3)
	Intensive-Margin, Multiple: $\omega_{(s,m)}$	Extensive-Margin, Multiple: $\omega_{(n,m)}$	Extensive-Margin, Single: $\omega_{(n,s)}$
A. Overall Effects			
Response share (p.p.)	[0.481, 2.360]	[0, 1.879]	[0, 1.879]
Assuming $\omega_{(n,m)} = 0$ (p.p.)	2.360	0	1.879
Assuming $\omega_{(n,s)} = 0$ (p.p.)	0.481	1.879	0
% Total responses	[20.372, 55.671]	[0, 79.628]	[0, 44.329]
B. Launch Effects Only			
Response share (p.p.)	[0.238, 1.765]	[0, 1.527]	[0, 1.527]
Assuming $\omega_{(n,m)} = 0$ (p.p.)	1.765	0	1.527
Assuming $\omega_{(n,s)} = 0$ (p.p.)	0.238	1.527	0
% Total responses	[13.461, 53.608]	[0, 86.539]	[0, 46.392]
C. Post-Upgrade Effects			
Response share (p.p.)	[1.162, 2.783]	[0, 1.620]	[0, 1.620]
Assuming $\omega_{(n,m)} = 0$ (p.p.)	2.783	0	1.620
Assuming $\omega_{(n,s)} = 0$ (p.p.)	1.162	1.620	0
% Total responses	[41.774, 63.201]	[0, 58.226]	[0, 36.799]

Notes: See notes to Table 4 for details.

Table A.7: CommonHelp's Multiple-Program Targeting Effects on Observables

	(1) Geographic per 10 miles (DD: far vs. close zips)	(2) Fixed (DRD: close zips only)	(3) Total (DD+DRD)	(4) Mean	(5) N
Low-Socioeconomic Status (p.p.)	−0.716	0.930***	0.481	35.165	371,416
	(0.495)	(0.256)	(0.428)		
A. Household-level Earnings History, 2009-2011					
Any employment (> \$0, p.p.)	0.501	−0.388*	−0.073	67.966	1,052,492
	(0.474)	(0.198)	(0.330)		
Fraction of quarters employed	0.009***	−0.005***	0.000	0.422	1,052,492
	(0.003)	(0.002)	(0.002)		
Mean quarterly wage earnings	62.348***	−30.681*	8.461	1,641.421	1,052,492
	(13.110)	(16.177)	(17.465)		
S.D. quarterly wage earnings	9.961	−23.136***	−16.883***	1,024.208	1,052,492
	(7.491)	(6.372)	(6.296)		
S.D. earnings among quarters with > \$0	10.065	−16.891**	−10.572	1,437.500	714,508
	(9.204)	(7.145)	(9.331)		
B. Individual-level Criminal Offense History, 1998–2011					
Any criminal offense (p.p.)	−0.236	0.324*	0.176	1.879	273,402
	(0.185)	(0.195)	(0.233)		
Number of offenses	−0.004	0.005**	0.003	0.024	273,402
	(0.003)	(0.003)	(0.003)		
Any property crime offense (p.p.)	0.025	0.195	0.211	1.127	273,402
	(0.115)	(0.172)	(0.191)		
Any violent offense (p.p.)	−0.099*	0.131*	0.069	0.245	273,402
	(0.058)	(0.076)	(0.085)		
Any drug offense (p.p.)	−0.162	0.117	0.015	0.527	273,402
	(0.123)	(0.172)	(0.193)		
Any juvenile offense (p.p.)	−0.308	0.134	−0.059	2.517	98,034
	(0.260)	(0.278)	(0.339)		
C. Household Characteristics					
Single-parent household (p.p.)	1.080***	0.027	0.705*	60.129	1,052,492
	(0.393)	(0.273)	(0.390)		
Two-parent household (p.p.)	−0.831**	−0.183	−0.705	27.809	1,052,492
	(0.404)	(0.277)	(0.426)		
Any children, ages 0–4 (p.p.)	−0.092	0.879***	0.822***	54.080	1,052,492
	(0.240)	(0.203)	(0.254)		
Any children, ages 5–9 (p.p.)	−0.112	−1.010***	−1.080***	26.901	1,052,492
	(0.340)	(0.189)	(0.308)		
Any children, ages 10–17 (p.p.)	0.210	0.130	0.262	18.985	1,052,492
	(0.183)	(0.225)	(0.254)		

Notes: This table reports estimates of γ^g, $\gamma^f = \gamma_1 - \gamma_0$, and γ from Equations (B.1)–(B.3) for various characteristics. Means are measured in the quarter prior to CommonHelp. Estimates for low-socioeconomic status and Panel B are restricted to the subsample of multiple program recipients in households with children in which either earnings is observable in 2009–2011 between ages 22–64 or juvenile crime is observable in 2000–2011 between ages 10–18. Panels A and C include all multiple program recipients in households with children. Standard errors are clustered at the county level and reported in parentheses. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Comparing Any vs. Multiple-Program Targeting Effects on Observables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Any program			Multiple programs			Diff.
	Total	Mean	N	Total	Mean	N	
Low-Socioeconomic Status (p.p.)	−0.265	46.764	682,411	0.481	35.165	371,416	0.746
	(0.341)			(0.428)			(0.547)
A. Household-level Earnings History, 2009-2011							
Any employment (>$0, p.p.)	0.562**	56.277	1,653,355	−0.073	67.966	1,052,492	−0.635
	(0.276)			(0.330)			(0.430)
Fraction of quarters employed	0.006**	0.358	1,653,355	0.000	0.422	1,052,492	−0.006**
	(0.002)			(0.002)			(0.003)
Mean quarterly wage earnings	40.685***	1,455.174	1,653,355	8.461	1,641.421	1,052,492	−32.224
	(12.880)			(17.465)			(21.701)
S.D. quarterly wage earnings	10.515	875.237	1,653,355	−16.883***	1,024.208	1,052,492	−27.398***
	(6.825)			(6.296)			(9.285)
S.D. earnings among quarters with >$0	10.397	1,483.141	1,037,627	−10.572	1,437.500	714,508	−20.969
	(10.497)			(9.331)			(14.045)
B. Individual-level Criminal Offense History, 1998–2011							
Any criminal offense (p.p.)	−0.013	2.121	500,002	0.176	1.879	273,402	0.189
	(0.140)			(0.233)			(0.272)
Number of offenses	−0.002	0.027	500,002	0.003	0.024	273,402	0.005
	(0.002)			(0.003)			(0.004)
Any property crime offense (p.p.)	0.010	1.182	500,002	0.211	1.127	273,402	0.201
	(0.134)			(0.191)			(0.233)
Any violent offense (p.p.)	0.049	0.300	500,002	0.069	0.245	273,402	0.020
	(0.053)			(0.085)			(0.100)
Any drug offense (p.p.)	−0.178**	0.663	500,002	0.015	0.527	273,402	0.193
	(0.083)			(0.193)			(0.210)
Any juvenile offense (p.p.)	0.463**	2.392	182,467	−0.059	2.517	98,034	−0.522
	(0.179)			(0.339)			(0.383)
C. Household Characteristics							
Single-parent household (p.p.)	0.689***	46.730	1,653,355	0.705*	60.129	1,052,492	0.016
	(0.241)			(0.390)			(0.458)
Two-parent household (p.p.)	−0.240	24.128	1,653,355	−0.705	27.809	1,052,492	−0.465
	(0.342)			(0.426)			(0.546)
Any children, ages 0–4 (p.p.)	−0.076	51.106	1,653,355	0.822***	54.080	1,052,492	0.898**
	(0.245)			(0.254)			(0.353)
Any children, ages 5–9 (p.p.)	−0.453*	26.691	1,653,355	−1.080***	26.901	1,052,492	−0.627
	(0.252)			(0.308)			(0.398)
Any children, ages 10–17 (p.p.)	0.547***	22.160	1,653,355	0.262	18.985	1,052,492	−0.285
	(0.138)			(0.254)			(0.289)

Notes: Columns (1) and (4) report estimates of γ from Equation (B.3) for any and multiple programs, respectively. Column (7) reports differences between Columns (4) and (1). See notes of Appendix Table A.7 for more details.

B Technical Appendix

B.1 Creating Program Participation Data at Zip-County Level

Since zip codes sometimes lie in multiple counties and are served by multiple field offices, our geographic unit of analysis is a zip-county pair. To create a frame of all zip-county pairs in Virginia, we rely on the HUD USPS Zip County Crosswalk Files, which track zip-county pairings on a quarterly basis. We take all unique zip-county pairings that appear at some point between 2010–2024 and assign them to every month of our analysis period. In total, this gives us approximately 1,170 unique ZIP codes and 1,700 unique ZIP-county combinations. Note that we keep both residential and non-residential (e.g., business) zip codes in our sample to account for all possible locations that applicants may list. We also rely on data collected by Bailey and Helmuth to incorporate historical zip code changes, based on zip code removals reported in the USPS bi-weekly bulletin board. Overall, these sources enable us to create a thorough and consistent dataset of zip-county pairs over time.

Given that the VDSS records report the zip code and county of residence of each program recipient, we calculate participation totals for each zip-county pair and month in the frame described above. If a zip code appears in the VDSS records but not in our frame (an occurrence that is very rare), then we assume the VDSS zip code is listed in error (or located out of state) and do not include counts corresponding to it. If the zip code of residence is not served by the LDSS county on record (which occurs in less than 2% of panels), we replace the county on record with the LDSS county serving that zip code that is closest to the county on record. We prioritize zip codes and counties based on program popularity (first SNAP, then Medicaid, then TANF). In a small minority of panels, individuals may also have conflicting geographic identifiers across programs in the same month. To resolve conflicts, we prioritize geographies for the program with the most recent start date and break remaining ties based on geographies for SNAP, Medicaid, and TANF (in that order). Finally, there are a few cases where a zip code lies in more than one county, but both counties happen to be served by the same LDSS office. In these cases, zip-county pairs are aggregated back up to zip codes.

B.2 Additional Details on Identification Strategies

This section provides additional details on the difference-in-differences (DD) and difference-in-regression discontinuity (DRD) designs described in Section 4.

B.2.1 Details on Difference-in-Differences Specification

Choice of Outcome Variable: Logarithmic Specification. Our primary DD specification uses log recipient counts as the dependent variable. By employing a log-linear specification, we require that zip codes grow at the same *proportional* rate. Thus, our DD estimator expresses CommonHelp's effects as percent changes. Given potential concerns with this approach (Chen and Roth 2024), we alternatively estimate a Poisson specification as a robustness check.

Covariates Used in Pre-Post Model. In the pre-post version of our DD model, we control for time-varying zip-level characteristics to improve precision. These include log population, poverty rate, unemployment rate, labor force participation rate, and population shares by race/ethnicity and gender. All controls are available starting in 2010, except for unemployment and labor force participation rates, which begin in 2011. Characteristics come from the Decennial Census and American Community Survey (ACS) at the Census-generated zip code tabulation areas (ZCTAs). We merge ACS data using zip-to-ZCTA crosswalks. All results are robust to excluding these controls.

B.2.2 Details on Difference-in-Regression Discontinuity Specification

Seasonality. We use a DRD—rather than a traditional RD—design given the possibility of seasonality in program enrollment. For example, seasonality in employment could drive seasonality in program receipt, and states may align application processing with the federal fiscal calendar. State administrative systems may process eligibility updates or backlog clearances in quarterly cycles. In addition, Medicaid enrollment records are reported at the quarterly level in the Medicaid Statistical Information System (MSIS) files submitted to the federal government and exhibit some seasonality. Another source of seasonality is the annual SNAP Cost-of-Living Adjustment (COLA) in October. SNAP benefits were frozen during the Great Recession, with no COLA increases in our October 2012 treatment window or October 2009-2011 control windows, but COLA increases resumed in October 2013. If COLA increases boost enrollment, then some of our control window estimates (starting in 2013) may be inflated and our DRD estimates would likely be understated.

Functional Form and Specifications. The ideal data for identifying these parameters would be daily address-level receipt data to enable finer identification around the cutoffs and restrict the analysis to households located immediately adjacent to field offices. Daily data would provide an (approximately) continuous running variable, permitting us to estimate discontinuities precisely at the launch and upgrade cutoff dates. Address-level data would allow us to focus on locations that are immediately neighboring field offices so that $\bar{d} \approx 0$, ensuring we identify impacts from reducing *only* fixed burdens. Instead, because our data are monthly and at the zip level, our DRD estimates may rely on observations farther from the cutoffs and may inadvertently capture impacts related to geographic burdens. We take a number of measures to minimize these potential biases. First, we employ a nonparametric local linear regression model by applying a triangular kernel that gives greater weight to months closer to the cutoff dates. Second, we restrict the sample to zip-counties within $\bar{d} = 1$ mile of their field office. This distance threshold is roughly 10% of the median distance to field office and includes only 47 zip codes (4.8% of our sample). We assess how sensitive our results are to this choice of bandwidth and alternative specifications in robustness checks.

B.3 Converting Regression Parameters to Percentage Point Effects

First, recall that β^g, β^{g15} are expressed as *per mile growth rate* in program receipt at the *zip code* level. In other words, for a zip code z with distance $D_z = d_z$,

$$\begin{aligned} g_z &\equiv (\beta^g + \beta^{g15}) \cdot d_z + \beta^f \\ &= \mathbb{E}\left[\log(Y_{z,post}(1)) - \log(Y_{z,post}(0)) \,|\, D_z = d_z\right] \\ &\approx \mathbb{E}\left[\frac{Y_{z,post}(1) - Y_{z,post}(0)}{Y_{z,post}(0)} \,|\, D_z = d_z\right], \end{aligned}$$

where $Y_{z,post}(0), Y_{z,post}(1)$ are the number of recipients in the post-CommonHelp period under no CommonHelp and under CommonHelp, respectively. Since $Y_{z,post} = Y_{z,post}(1) = (1 + g_z) \cdot Y_{z,post}(0)$, we can identify the counterfactual level of recipients in the post-CommonHelp period in each zip code by rescaling the factual level:

$$Y_{z,post}(0) \;=\; \frac{1}{1+g_z} \cdot Y_{z,post}.$$

Thus, the total change in the number of recipients of p in the population caused by CommonHelp is thus

$$\Delta Y \;\equiv\; \sum_{z \in \mathcal{Z}} \left[Y_{z,post}(1) - Y_{z,post}(0)\right] \;=\; \sum_{z \in \mathcal{Z}} \left(1 - \frac{1}{1+g_z}\right) \cdot Y_{z,post}.$$

B.4 Partial Identification Analysis

B.4.1 Mapping Percentage-Point Effects to Latent Responder Shares

Under appropriate parallel trends and smoothness assumptions, our DD and RD designs identify the causal percentage-point effect of CommonHelp on the share of recipients of program $p \in \mathcal{P}$ defined in Equation (6). Assuming expectations are taken over the population of recipients, this parameter corresponds to

$$\tilde{\beta}_p \;=\; \mathbb{E}[P_{p,it}(1) - P_{p,it}(0) \,|\, Post_t = 1].$$

The monotonicity assumption restricts behavioral responses to CommonHelp to be as

follows:[1]

$$(P_{it}(0), P_{it}(1)) = \begin{cases} (n,n) & \text{(none, non-responders)} \\ (n,s) & \text{(single, extensive-margin responders)} \\ (n,m) & \text{(single, intensive-margin responders)} \\ (s,s) & \text{(single, non-responders)} \\ (s,m) & \text{(multiple, intensive-margin responders)} \\ (m,m) & \text{(multiple, non-responders)} \end{cases}$$

From the law of total probability and given $n \subset s \subset m$,

$$\begin{aligned} \tilde{\beta}_p &= \Pr(P_{p,it}(1)=1, P_{p,it}(0)=0 \,|\, Post_t=1) - \Pr(P_{p,it}(1)=0, P_{p,it}(0)=1 \,|\, Post_t=1) \\ &= \Pr(P_{it}(0) \subset P_{it}(1) = p \,|\, Post_t=1) \;-\; \Pr(P_{it}(0) = p \subset P_{it}(1) \,|\, Post_t=1) \\ &= \sum_{\substack{p_0 \in \mathcal{P} \\ \text{s.t. } p_0 \subset p}} \omega_{(p_0,p)} - \sum_{\substack{p_1 \in \mathcal{P} \\ \text{s.t. } p \subset p_1}} \omega_{(p,p_1)}, \end{aligned}$$

where the penultimate equality follows from the monotonicity assumption and final equality follows from the definition of $\omega_{(p_0,p_1)}$. In particular,

$$\begin{aligned} \tilde{\beta}_s &= \omega_{(n,s)} - \omega_{(s,m)} \\ \tilde{\beta}_m &= \omega_{(n,m)} + \omega_{(s,m)}. \end{aligned}$$

Notice that the parameters corresponding to effects on "any program receipt" ($\tilde{\beta}_{s\cup m}$) and "no program receipt" ($\tilde{\beta}_n$) are linearly dependent on $\tilde{\beta}_s$, $\tilde{\beta}_m$: $\tilde{\beta}_{s\cup m} = \omega_{(n,s)} + \omega_{(n,m)} = \tilde{\beta}_s + \tilde{\beta}_m$ and $\tilde{\beta}_n = -\omega_{n,s} - \omega_{n,m} = -\tilde{\beta}_s - \tilde{\beta}_m$. In fact, the following lemma states that any other enrollment effect parameter based on $\mathcal{P}$ provides no additional identifying content:

Lemma 1. *If the identification assumptions (parallel trends, smoothness conditions) hold, all enrollment effect parameters are linearly spanned by $\tilde{\beta}_s$, $\tilde{\beta}_m$.*

Proof. Let $P_{it}^{\dagger} : \mathcal{P} \longrightarrow \mathbb{R}$ be some arbitrary program receipt variable with associated enrollment effect parameter $\tilde{\beta}^{\dagger}$. Note that $P_{it}^{\dagger}$ (and, analogously, its potential outcomes $P_{it}^{\dagger}(k),\ k \in \{0,1\}$) can always be represented as $P_{it}^{\dagger} = \sum_{p\in\mathcal{P}} \pi_p \cdot P_{p,it}$, where $\pi_p \in \mathbb{R}$. This implies that

$$\tilde{\beta}^{\dagger} = \mathbb{E}[P_{it}^{\dagger}(1) - P_{it}^{\dagger}(0)] = \sum_{p\in\mathcal{P}} \pi_p \cdot \mathbb{E}[P_{p,it}(1) - P_{p,it}(0)] = \sum_{p\in\mathcal{P}} \pi_p \cdot \tilde{\beta}_p,$$

where the first and last equalities follow from the identification assumptions holding. Given $n \in \mathcal{P}$, the fact that $\tilde{\beta}_n = -\tilde{\beta}_s - \tilde{\beta}_m$ completes the proof. □

1. These behavioral response groups are analogous to the "never takers," "compliers," and "always takers" commonly used in the literature employing instrumental variables (IV) (Angrist et al. 1996). We do not use this taxonomy to emphasize we are not employing an IV framework, but rather a simple "treatment" (CommonHelp) and multinomial "outcome" (program receipt types) framework.

Thus, $\tilde{\beta}_s$, $\tilde{\beta}_m$ exhausts all information about the underlying response shares. Given that these two identified parameters depend on three unknown response shares, we must pursue a partial identification analysis.

B.4.2 Derivation of Sharp Bounds

Define the caseload (i.e., the sum of single and multiple program recipients) as $\sum_{z\in\mathcal{Z}} \bar{Y}_{z,post}$ where $\bar{Y}_{z,post} = \bar{Y}_{z,post,s} + \bar{Y}_{z,post,m}$. We know that $\omega_{(p_0,p_1)} \in [0,1]$. We can further tighten the bounds by drawing on the shares of non-responders and the adding up constraint $\sum_{(p_0,p_1)} \omega_{(p_0,p_1)} = 1$.

The share of single non-responders corresponds to the population share of recipients of s under no CommonHelp in the post-CommonHelp period, less the share of intensive-margin compliers $\omega_{(s,m)}$. This is precisely

$$\omega_{(s,s)} \;=\; \frac{\sum_{z\in\mathcal{Z}} Y_{z,post,s}(0)}{\sum_{z\in\mathcal{Z}} \bar{Y}_{z,post}} - \omega_{(s,m)} \;=\; \underbrace{\frac{\sum_{z\in\mathcal{Z}} \frac{1}{1+g_{z,s}} \cdot Y_{z,post,s}}{\sum_{z\in\mathcal{Z}} \bar{Y}_{z,post}}}_{\theta_{(s,s)}} - \omega_{(s,m)}.$$

Similarly, the share of multiple non-responders is

$$\omega_{(m,m)} \;=\; \frac{\sum_{z\in\mathcal{Z}} Y_{z,post,m}(0)}{\sum_{z\in\mathcal{Z}} \bar{Y}_{z,post}} \;=\; \underbrace{\frac{\sum_{z\in\mathcal{Z}} \frac{1}{1+g_{z,m}} \cdot Y_{z,post,m}}{\sum_{z\in\mathcal{Z}} \bar{Y}_{z,post}}}_{\theta_{(m,m)}}.$$

Meanwhile, since our baseline population are recipients in the post-CommonHelp period, the share of no-program non-responders does not exist by construction: $\omega_{(n,n)} = 0$.

Based on the aforementioned linear equations, we possess the following linear restrictions:

$$\begin{aligned}
\omega_{(n,s)} - \omega_{(s,m)} &= \tilde{\beta}_s \\
\omega_{(n,m)} + \omega_{(s,m)} &= \tilde{\beta}_m \\
\omega_{(n,s)} + \omega_{(n,m)} &= \tilde{\beta}_s + \tilde{\beta}_m \\
\omega_{(s,m)} + \omega_{(s,s)} &= \theta_{(s,s)}
\end{aligned}$$

These deliver the following sharp bounds on each responder share:

$$\begin{aligned}
\max\{0,\ \tilde{\beta}_s\} \;\le\; \omega_{(n,s)} \;&\le\; \min\{1+\tilde{\beta}_s,\ \tilde{\beta}_s+\tilde{\beta}_m\} \\
0 \;\le\; \omega_{(n,m)} \;&\le\; \min\{\tilde{\beta}_m,\ \tilde{\beta}_s+\tilde{\beta}_m\} \\
\max\{0,\ -\tilde{\beta}_s\} \;\le\; \omega_{(s,m)} \;&\le\; \min\{1-\tilde{\beta}_s,\ \tilde{\beta}_m,\ \theta_{(s,s)}\}
\end{aligned}$$

B.4.3 Characterizing Responders

Let $X_i \in \{0,1\}$ indicate whether recipient i is low-socioeconomic status (SES). Our aim is to estimate the probability that a given responder type $(p_0,p_1) \in \{(s,m),(n,m),(n,s)\}$ is

low-SES:

$$\Pr\left(X_i = 1 \;\middle|\; (P_{it}(0), P_{it}(1)) = (p_0, p_1)\right).$$

First, we estimate (β^g, β^f) separately for each subgroup $X_i = 0$ and $X_i = 1$, we repeat the previous analyses to recover bounds on each responder-type share, by subgroup. Define the responder share of subgroup $X_i = x$ as

$$\omega^x_{(p_0,p_1)} \;\equiv\; \Pr(P_{it}(0) = p_0, P_{it}(1) = p_1 \,|\, X_i = x)$$

and denote its sharp lower and upper bounds as $\underline{\omega}^x_{(p_0,p_1)}$ and $\bar{\omega}^x_{(p_0,p_1)}$, respectively.

Note that we can measure SES only for a subsample of our main data. Therefore, the overall responder shares of the subsample may differ from the overall responder shares of the full sample. We abuse notation and denote $\omega_{(p_0,p_1)}$ as the overall responder share for this subsample. We re-estimate its sharp bounds, denoted $\underline{\omega}_{(p_0,p_1)}$ and $\bar{\omega}_{(p_0,p_1)}$.

From Bayes' rule, we can recover the target parameter:

$$\Pr(X_i = 1 \;|\; (P_{it}(0), P_{it}(1)) = (p_0, p_1)) \;=\; \frac{\omega^1_{(p_0,p_1)}}{\omega_{(p_0,p_1)}} \cdot \Pr(X_i = 1)$$

The challenge, however, is that $\omega^1_{(p_0,p_1)}$ and $\omega_{(p_0,p_1)}$ can take a range of values. To identify bounds on $\Pr(X_i = 1 \;|\; (P_{it}(0), P_{it}(1)) = (p_0, p_1))$, we pursue the following algorithm:

1. Draw a value $\omega_{(p_0,p_1)}$ from a fine grid of $[\underline{\omega}_{(p_0,p_1)},\ \bar{\omega}_{(p_0,p_1)}]$.
2. Find the permissible set of $\omega^1_{(p_0,p_1)}$, denoted $\Omega^1_{(p_0,p_1)}$, given $\omega_{(p_0,p_1)}$. To do this:
   (a) Draw a value of $\omega^0_{(p_0,p_1)}$ from a fine grid of $[\underline{\omega}^0_{(p_0,p_1)},\ \bar{\omega}^0_{(p_0,p_1)}]$.
   (b) Find the implied value of $\omega^1_{(p_0,p_1)}$ satisfying the law of total probability:

$$\omega_{(p_0,p_1)} \;=\; \omega_{(p_0,p_1)\,|\,0} \cdot [1 - \Pr(X_i = 1)] + \omega_{(p_0,p_1)\,|\,1} \cdot \Pr(X_i = 1).$$

   (c) If the implied value $\omega^1_{(p_0,p_1)} \in [\underline{\omega}^1_{(p_0,p_1)},\ \bar{\omega}^1_{(p_0,p_1)}]$, then $\omega^1_{(p_0,p_1)} \in \Omega^1_{(p_0,p_1)}$.
   (d) Repeat this process for all values $\omega^0_{(p_0,p_1)}$ over the grid $[\underline{\omega}^0_{(p_0,p_1)},\ \bar{\omega}^0_{(p_0,p_1)}]$.
3. For each permissible value $\omega^1_{(p_0,p_1)} \in \Omega^1_{(p_0,p_1)}$ and given $\omega_{(p_0,p_1)}$, compute $\Pr(X_i = 1 \,|\, (P_{it}(0), P_{it}(1)) = (p_0, p_1))$ via Bayes' rule. Denote this set of values as $\mathcal{P}(\omega_{(p_0,p_1)})$. Define $\underline{p}(\omega_{(p_0,p_1)}) \equiv \min\{\mathcal{P}(\omega_{(p_0,p_1)})\}$ and $\bar{p}(\omega_{(p_0,p_1)}) \equiv \max\{\mathcal{P}(\omega_{(p_0,p_1)})\}$.
4. Repeat this process for all values $\omega_{(p_0,p_1)}$ over the full grid $[\underline{\omega}_{(p_0,p_1)},\ \bar{\omega}_{(p_0,p_1)}]$. Estimate the bounds of $\Pr(X_i = 1 \,|\, (P_{it}(0), P_{it}(1)) = (p_0, p_1))$ as

$$\left[\min_{\omega_{(p_0,p_1)} \in [\underline{\omega}_{(p_0,p_1)},\ \bar{\omega}_{(p_0,p_1)}]} \left\{\underline{p}(\omega_{(p_0,p_1)})\right\},\ \max_{\omega_{(p_0,p_1)} \in [\underline{\omega}_{(p_0,p_1)},\ \bar{\omega}_{(p_0,p_1)}]} \left\{\bar{p}(\omega_{(p_0,p_1)})\right\}\right].$$

B.5 Conventional Targeting Analysis

Empirical Analysis. To assess changes in recipient composition, we adapt our DD and DRD models to the individual level using repeated monthly cross-sections of program recipients. For a given predetermined characteristic C_{it}, we estimate:

$$C_{it} = \gamma^g(D_{z(i)} \times PostCH_t) + \delta_{z(i)} + \lambda_t + \theta_1 Age_i + \theta_2 Age_i^2 + \epsilon_{it}, \tag{B.1}$$

$$\begin{aligned} C_{it} = {} & (\gamma_1 \cdot \mathbb{1}\{t \geq 0\} + f_{1,pre}(t) \cdot \mathbb{1}\{t < 0\} + f_{1,post}(t) \cdot \mathbb{1}\{t \geq 0\}) \times Treat_t \\ & + (\gamma_0 \cdot \mathbb{1}\{t \geq 0\} + f_{0,pre}(t) \cdot \mathbb{1}\{t < 0\} + f_{0,post}(t) \cdot \mathbb{1}\{t \geq 0\}) \times Control_t \\ & + \theta_1 Age_i + \theta_2 Age_i^2 + \epsilon_{zt}, \end{aligned} \tag{B.2}$$

where we control for a quadratic of i's age at CommonHelp's launch to account for potential life-cycle dynamics that influence the measurement of C_{it}.

Here, γ^g and $\gamma^f \equiv \gamma_1 - \gamma_0$ capture how reductions in geographic and fixed burdens affected the observable composition of multiple program recipients.[2] We combine these parameters to assess CommonHelp's overall effect on recipient characteristics in a given zip code using:

$$\Delta C_z \equiv \gamma^g \cdot D_z + \gamma^f. \tag{B.3}$$

This quantity reflects the change in socioeconomic status among multiple program enrollees for a zip code with distance D_z from its field office.

Results. Table A.7, Columns (1) and (2) report estimates from the individual-level DD and DRD regressions, revealing starkly opposite patterns in how reductions in geographic versus fixed burdens affected the composition of multiple program recipients. For our summary low-SES indicator, CommonHelp led to a 0.72 percentage-point *decrease* in low-SES multiple program recipients for every 10 miles of distance, but a 0.93 percentage-point *increase* in low-SES recipients in zip codes closest to field offices. Appendix Figure A.28 provides visual evidence supporting the validity of these estimates. Although the summary low-SES estimate from reduced geographic burdens is not statistically significant, Appendix Table A.7 reports significant and consistent patterns among its components, including precise increases in recipients from households with longer employment histories and higher mean wages (Panel A) and a significant decline in recipients with violent criminal offense records (Panel B).

Table A.7, Column (3) reports *average* targeting effects among multiple program recipients by setting $D_z = \bar{D} = 6.6$ miles in Equation (B.3). On average, CommonHelp's impact on the composition of multiple program recipients was small and statistically insignificant (0.48 percentage-point increase in the share of low-SES recipients), with little to no significant changes in employment, earnings, or criminal history. Given the targeting effects from reducing geographic and fixed burdens move in opposite directions, it is unsurprising that average effects across all zip codes mask this important heterogeneity.

2. Thus, if C_{it} is binary, then γ^g and γ^f reflect *percentage-point* changes (rather than *percent* changes).

B.6 Details on Marginal Value of Public Funds Analysis

We closely follow Hendren and Sprung-Keyser (2020) (hereafter, HSK) to assess the Marginal Value of Public Funds (MVPF) of CommonHelp. The MVPF is defined as the ratio of a policy's willingness-to-pay (WTP) to its net fiscal cost, incorporating fiscal externalities from behavioral responses. Here, we describe in detail each set of MVPF components that we calculate: (1) direct WTP, (2) direct fiscal costs (in terms of benefit amounts and administrative costs), and (3) behavioral responses and indirect effects. We incorporate roughly 125 parameters drawn from our estimates, external data sources, and the broader literature. Where relevant, we distinguish between inputs that are relevant for calculating multiple-program MVPFs versus single-program MVPFs.

B.6.1 Details on Direct Willingness-to-Pay (WTP)

CommonHelp directly reduced administrative burdens for applicants, generating time savings that we convert into WTP. We consider two forms of time savings associated with CommonHelp. First, the transition from a paper-based form to an online system reduced the application completion time from approximately two hours to just 30 minutes, saving 90 minutes per application (Virginia Department of Social Services 2012; Code for America 2024). Second, using roundtrip distances from residential ZIP codes to local field offices (estimated using Google Maps API), we calculate travel time savings assuming two field office visits would have been required to meet with caseworkers and submit the application. This assumption is consistent with the evidence in Bartlett et al. (2004), who find that SNAP applicants made approximately 2.4 trips to the field office during the application process.

We value all time at \$14.50/hour (twice the minimum wage rate, a customary practice in the literature) and scale these WTP estimates by the number of new household units who received programs *through CommonHelp* between October 2012 and December 2015. While we cannot directly observe which recipients applied via CommonHelp versus the traditional application process, we scale the total number of entrants by 15% based on internal discussions with VDSS staff (who suggested 15% was the approximate share of enrollees entering via CommonHelp). For multiple-program MVPFs, we include all new recipients across SNAP, Medicaid, and TANF; for single-program MVPFs, we restrict to the respective program's new recipients. We also incorporate time savings from recertifications, assuming these occur twice for roughly 60% of recipients (based on empirical survival curve estimates) and require half the time of an initial application. By estimating time savings for *all* enrollees and recertifiers, we account for the fact that CommonHelp yielded time savings among both marginal and inframarginal recipients.

Note that these WTP estimates are likely to be conservative for a number of reasons. First, they do not capture the change in WTP for applicants who may have been rejected (as we only have information on applicants who ultimately enrolled). Second, our estimates of travel savings reflect only the opportunity cost of time and exclude out-of-pocket travel costs (e.g., gasoline). Finally, we do not account for the wait time savings from avoiding in-person office congestion. Rowe et al. (2015) suggest that wait times could be several hours depending on the time of day or the month, even though clients were generally interviewed within 60 minutes of arrival. Bartlett et al. (2004) also estimate that applicants spent approximately

3.9 hours (inclusive of time spent waiting for assistance) in their field office completing the application process.

B.6.2 Details on Direct Fiscal Costs

We next discuss CommonHelp's impacts on direct fiscal costs, starting with two types of administrative costs (both the up-front development cost of the platform and the changes in per-recipient administrative costs). We then consider the increase in benefit disbursements associated with enrollment growth, where we use imputed benefit amounts to re-estimate the policy's impact on benefit outlays via the DD and DRD regressions.

Change in Fixed Administrative Costs Due to CommonHelp. According to an article from GovTech, the up-front development cost of the CommonHelp platform totaled $10,779,739 (Pittman 2012). We assume that this cost was fully borne over the three years after CommonHelp (which constitutes our post-treatment period of analysis). Note that we include the full up-front cost of CommonHelp, even though our analysis focuses only on individuals in households with children.

Change in Variable Administrative Costs Due to CommonHelp. In addition to increases in fixed administrative costs, we allow for CommonHelp to have affected variable (per-recipient) administrative costs. Specifically, CommonHelp may have generated efficiency gains by reducing agency workloads via streamlining the processing of applications. To assess this, we conduct a synthetic control analysis that uses annual SNAP administrative cost data to compare changes in per-recipient costs in Virginia against those in a "synthetic" control state between FY 2013–2015. Comparable analyses are infeasible for TANF or Medicaid, as TANF data are subject to significant measurement error and Medicaid administrative cost trends are confounded by Medicaid expansion and other policy changes across states. we assume that the effects for SNAP extend proportionally to these programs, as all are administered by the same state and field offices under VDSS.

Benefit Amounts for SNAP. As our VDSS microdata contain only data on benefit receipt and not amounts, we impute amounts based on binned averages from the public-use SNAP Quality Control (QC) microdata files. Specifically, using the QC data, we collapse benefit amounts using cells based on case size (ranging from 1 to 6+ individuals), monthly case-level wages ($0, $1-$500, $501-$1,000, $1,001-$1,500, and $1,501+), and fiscal year limiting to SNAP recipients in Virginia. Individual cases are weighted using survey weights. We ensure that all cells have at least 5 unweighted observations in the QC data, and we aggregate across wage categories if needed to meet the minimum cell size threshold. In the VDSS microdata, we calculate case size using SNAP case identifiers. We also calculate case-level wages by linking to the VEC quarterly wage records, converting quarterly wages to monthly wages by dividing by three, and summing across individual wages in a given case.

Benefit Amounts for TANF. We also use binned averages from the QC data for TANF benefit amounts, which we have only for TANF recipients who also receive SNAP. However,

this is a relatively innocuous restriction, as approximately 90% of all TANF recipients receive SNAP. Specifically, using the QC data, we collapse benefit amounts using cells based on case size (fewer than 3, exactly three, and 4+ individuals), whether or not a case has monthly wages, and fiscal year limiting to SNAP recipients receiving TANF in Virginia. Individual cases are weighted using survey weights. We ensure that all cells have at least 5 unweighted observations in the QC data, and we aggregate across wage categories if needed to meet the minimum cell size threshold. In the VDSS microdata, we calculate case size using TANF case identifiers. We also calculate indicators for case-level wages based on whether any individual has earnings based on the VEC quarterly wage records.

Benefit Amounts for Medicaid. For Medicaid, we assign individual-level amounts that are equal to spending per-enrollee by age group buckets (0-18, 19-44, 45-64, 65-84, and 85+). We obtain overall spending averages by state (focusing on the amounts for Virginia) and year from the Centers for Medicare and Medicaid Services (CMS) Office of the Actuary. We then multiply these annual state-level amounts by fractional adjustments for each age group, which are obtained from national CMS data.

B.6.3 Details on Behavioral Responses and Indirect Effects

Finally, CommonHelp's MVPF may be shaped by behavioral responses and other indirect effects. These include fiscal externalities and shifts in WTP from informational frictions. It is beyond the scope of this paper to estimate these directly with our data, so we draw on studies used in HSK and the broader literature to assign relevant values.

Some studies calculate the MVPF by aggregating the WTP and fiscal externalities from as many outcomes or behavioral responses as possible. While this is useful for obtaining the true MVPF of a particular program, it complicates our primary objective: assessing the sensitivity of the MVPF of CommonHelp when accounting for single vs. multiple program participation. This is because different program MVPFs possess different ingredients based on the available evidence on a program's causal impacts on outcomes and behavioral responses, which can mechanically lead to different MVPF calculations based on the programs included in the calculation. For example, Bailey et al. (2024) recently estimated that the MVPF of SNAP's historical introduction is 62.3—much of which is driven by their estimates of SNAP's impacts on life expectancy for young children. In contrast, there appears to be no existing evidence of the causal impacts of TANF on life expectancy. Comparing the MVPF of CommonHelp based on only TANF versus TANF *and* SNAP would be an uninterpretable exercise if the WTP and fiscal externalities of SNAP included life expectancy. In order to facilitate more interpretable comparisons, we focus on outcomes and behavioral responses that are arguably considered "first-order" from a policymaker's perspective and in the academic literature.

For fiscal externalities, we restrict to estimates based on labor supply, long-run earnings, and cost savings (in Medicaid's case). For example, we use Hoynes and Schanzenbach (2012) for labor supply effects of SNAP and Bailey et al. (2024) for long-run impacts on children's income. In all cases, we obtain separate estimates for program impacts across four age groups: young children (0–5), older children (6–17), adults (18–64), and seniors (65+). When computing multiple-program MVPFs, we must make additional modeling assumptions about how fiscal externalities interact across programs. We consider three scenarios:

(1) externalities are additive and independent across programs; (2) externality rates are amplified by 10%; and (3) externality rates are amplified by 50%. The latter two might arise from sharper behavioral responses to benefit cliffs (Moffitt 1979) or complementarities across programs (Rossin-Slater and Wüst 2020).

We also allow for the possibility that increased program receipt raises the WTP for CommonHelp. Under a neoclassical model with full information, induced recipients are marginal and thus possess a WTP of \$0. However, credible evidence suggests that informational frictions limit take-up. In our context, substantial rates of incomplete multiple program take-up (Fact 3) and findings from Finkelstein and Notowidigdo (2019)—who document 90% misperception of SNAP eligibility among elderly Medicaid recipients—suggest that induced recipients would derive higher WTP. We therefore incorporate WTP estimates from the literature, disaggregated by age group and scaled based on different assumptions about the degree of misperception: 0% (neoclassical), 50%, 90%, and 100% (full misperception). We also try incorporating the WTP of only recipients below age 18, which assumes that their parents who apply for programs do not internalize the value of benefits reaped by their children (i.e., are not "altruistic"). For multiple-program MVPFs, we sum the valuations across programs; for single-program MVPFs, we include only the respective program's WTP.

Below, we discuss in detail the sources of the fiscal externality and WTP estimates for each of the three programs we analyze: SNAP, Medicaid, and TANF.

SNAP. We focus on the fiscal externalities and WTP of SNAP from the program's historical introduction (then known as Food Stamps). With respect to behavioral responses and outcomes, we focus on SNAP's impacts on parent and child earnings.

We begin with the fiscal externalities. Drawing on estimates from Hoynes and Schanzenbach (2012), HSK calculate that each \$1 of SNAP benefits decreases labor earnings by \$0.16 among non-elderly household heads with 12 or fewer years of education. This number is obtained by scaling the reduced-form impact of SNAP's introduction on tax revenue from earnings by a 6% participation rate, and thus represents impacts *per household head*. We assume this fiscal externality applies to all SNAP recipients between ages 18–64. Meanwhile, for children, Bailey et al. (2024) calculate that the impacts of receiving SNAP among children *between ages 0–5* on long-run earnings generates a fiscal externality of −\$0.07. This number is obtained by scaling the reduced-form impact of SNAP's *availability* between ages 0–5 by a 16% participation rate between ages 0–5. Therefore, the \$0.07 figure represents the impact of receiving SNAP *per child of ages 0–5*, irrespective of duration in SNAP. Like HSK, we assume no fiscal externalities among children between ages 6–17 and elderly adults age 65 or older. We ignore impacts on other outcomes such as crime and mortality.

For WTP, we draw on HSK, Bailey et al. (2024), and Finkelstein and Notowidigdo (2019). HSK employ a revealed preference argument that accounts for the SNAP benefit–earnings schedule and negative earnings responses for the marginal complier to estimate a \$0.62 WTP for female household heads (based on Hoynes and Schanzenbach 2012). We assume this holds for each non-elderly adult.[3] HSK also use estimates from Bailey et al. (2024) to estimate a

3. Note that HSK approximate SNAP benefit amounts based on average household size. For female heads of households, it is reasonable that they value all SNAP benefits provided to household. Extrapolating this to two-parent households is an admittedly strong assumption.

$0.45 WTP per child between age 0–5, which is driven by their later-life after-tax earnings. Similar to before, we assume $0 WTP per child between ages 6–17. Meanwhile, Finkelstein and Notowidigdo (2019) cite studies on the fungibility of SNAP benefits to estimate that elderly adults have a $0.80 WTP.

Medicaid. We focus on fiscal externalities and WTP of Medicaid using different settings and natural experiments depending on the age group examined. The estimates for children (age 0-17) leverage the discontinuity in eligibility for children born around September 30, 1983. The estimates for non-elderly adults (age 18-64) are based on the Oregon Health Insurance Experiment (OHIE), under the assumption that the estimates for low-income adults without dependents (the population of interest in the OHIE) can be extrapolated to the low-income adults in our setting (many of whom are likely to be parents, for whom there are very few estimates in the literature). Finally, there are very few studies examining the impacts of Medicaid for elderly adults (age 65+), but our estimates are based on a structural model in Hackmann (2019). With respect to behavioral responses and outcomes, we focus on changes in moral hazard and hospitalization costs/uncompensated care for children and non-elderly adults, and on quality of care via access to nursing homes for the elderly.[4]

We again start by discussing fiscal externalities. For children, we assume that every $1 in Medicaid expenditure translates to a fiscal externality of −$1. Based on estimates from Wherry et al. (2018), HSK suggest that the cost savings to government (in terms of hospitalization and emergency room costs) pays for Medicaid itself by the time children are age 54. It could very well be that the fiscal externality is larger than −$1, but using a lower bound of −$1 is sensible given that we do not account for impacts on mortality when calculating WTP. For non-elderly adults, we assume a fiscal externality of 0 given the lack of fiscal impacts detected in terms of labor market changes or future healthcare cost reductions. For elderly adults, we again assume a fiscal externality of 0. Theoretically, in there are offsetting forces in terms of decreased costs on home and community-based services upon moving to nursing homes and increased costs from the expansion of the nursing home market.

Next, we discuss estimates for WTP. For children, HSK lay out an estimate that is based on the WTP per dollar of Medicaid expansion, but we recalculate this estimate so that it is in terms of Medicaid receipt. Starting with an initial cost of providing Medicaid to children of $1,484 per Wherry et al. (2018), we spread out the average cost over ages 8-14 and discount back to age 8 to obtain a discounted costs of eligibility of $1,181. Lo Sasso and Seamster (2007) estimate that the change in uncompensated costs per additional eligible child is $17, which after scaling up by the take-up rate (44.6%) and discounting over seven years becomes $31. Multiplying the cost of Medicaid provision by the moral hazard rate of 0.6 from Card and Shore-Sheppard (2004) and subtracting uncompensated care costs yields $678, which we divide by $1,181 to get a WTP of 0.57 per dollar of Medicaid expenditure.

For adults, we rely on Finkelstein et al. (2019) to obtain a WTP of $1.16 per recipient dollar. Note that this is the WTP under a "complete information approach," and Finkelstein

4. We do not explicitly account for mortality to keep our Medicaid estimates more consistent with those for SNAP and TANF (which do not account for mortality). However, note that the WTP estimates for OHIE in principle reflect long-run impacts on health, but in practice these effects are negligible.

et al. (2019) estimate the WTP under alternative assumptions including a "consumption-based optimization approach" (0.99) and a "health-based optimization approach" (0.55). We follow HSK in taking the complete information approach as our baseline. Finally, for elderly adults, we employ estimates from Table 4 of Hackmann (2019), who estimates that a universal 10 percent increase in Medicaid reimbursement rates would lead to an increase in consumer surplus of $212 million and an increase in spending of $331 million. These are based purely on changes in the nursing home industry, which primarily affects the elderly. Under the assumption that these intensive margin effects translate to extensive margin changes, we take the ratio of these magnitudes to obtain a WTP of 0.64 for the elderly.

TANF. Unfortunately, scarce studies examine the impacts of TANF participation. Thus, we instead focus on the fiscal externalities and WTP of TANF's predecessor, AFDC, based on the historical program's expansion in benefit generosity.[5] With respect to behavioral responses and outcomes, we focus on TANF's impacts on parent and child earnings. We modify HSK's methodology to meet our purposes.

HSK estimate that $1,000 additional AFDC benefits to beneficiaries generates a fiscal externality of $147.77 among single mothers based on estimates from Meyer and Rosenbaum (2001). We assume this effect estimate, which captures intensive-margin changes in AFDC benefit generosity, translates to the extensive-margin TANF receipt for all non-elderly adults. Thus, we assume that every $1 in TANF expenditure arising from *TANF receipt* translates to a fiscal externality of $0.15 *per non-elderly adult*. For children, HSK use estimates from Currie and Cole (1993), which focuses on impacts of AFDC receipt during pregnancy on child birthweight, to approximate a fiscal externality of −$137.77 *per birth* and *per $1,000 additional AFDC benefits*. Like the externality estimate for non-elderly adults, we assume this can be translated to public expenditures arising from extensive-margin TANF receipt. This yields a fiscal externality of −$0.14 *per birth* based on a 20% tax rate on future earnings. In their MVPF calculations, HSK further scale this by the fertility rate among AFDC participants. Instead, we assign this fiscal externality to children between ages 0–5, allowing us to use the same age groups used for SNAP. This assumes that increased financial resources received during early childhood may yield similar effects as those received during pregnancy. We also assume $0 fiscal externalities among children between ages 6–17 and elderly adults.[6]

With respect to WTP, HSK project a $1-to-$1 WTP for single mothers. We assume this holds for all non-elderly and elderly adults. Meanwhile, based on after-tax earnings from AFDC births, we assume children between ages 0–5 have a WTP of $0.55 per benefit dollar while children between ages 6–17 do not have any WTP for TANF.

5. HSK evaluate the MVPF of time limits on AFDC and consider various other welfare-to-work experiments. These policy experiments are closer to the spirit of the TANF program, although the MVPF of *reforms* is not informative of the MVPF of TANF itself.

6. Admittedly, assuming no fiscal externalities for older children is a strong assumption given evidence on the intergenerational transmission of welfare receipt.

Appendix References

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91 (434): 444–455.

Bailey, Martha J., Hilary Hoynes, Maya Rossin-Slater, and Reed Walker. 2024. "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program." *Review of Economic Studies* 91 (3): 1291–1330.

Bartlett, Susan, Nancy Burstein, and William Hamilton. 2004. *Food Stamp Program Access Study: Final Report.* Technical report. November: U.S. Department of Agriculture, Economic Research Service.

Card, David, and Lara D. Shore-Sheppard. 2004. "Using Discontinuous Eligibility Rules to Identify the Effects of the Federal Medicaid Expansions on Low-Income Children." *Review of Economics and Statistics* 86 (3): 752–766.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with Zeros? Some Problems and Solutions." *The Quarterly Journal of Economics* 139 (2): 891–936.

Code for America. 2024. *Benefits Enrollment Field Guide.* Accessed: December 2, 2024. https://codeforamerica.org/explore/benefits-enrollment-field-guide/?att=online&j=VA&app=102&dev=d#explore.

Currie, Janet, and Nancy Cole. 1993. "Welfare and Child Health: The Link Between AFDC Participation and Birth Weight." *The American Economic Review* 83 (4): 971–985.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127 (6): 2836–2874.

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

Hackmann, Martin B. 2019. "Incentivizing Better Quality of Care: The Role of Medicaid and Competition in the Nursing Home Industry." *American Economic Review* 109 (5): 1684–1716.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics* 135 (3): 1209–1318.

Hoynes, Hilary Williamson, and Diane Whitmore Schanzenbach. 2012. "Work Incentives and the Food Stamp Program." *Journal of Public Economics* 96 (1-2): 151–162.

Lo Sasso, Anthony T., and Dorian G. Seamster. 2007. "How Federal and State Policies Affected Hospital Uncompensated Care Provision in the 1990s." *Medical Care Research and Review* 64 (6): 731–744.

Meyer, Bruce D., and Dan T. Rosenbaum. 2001. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." *The Quarterly Journal of Economics* 116 (3): 1063–1114.

Moffitt, Robert. 1979. "Cumulative Effective Tax Rates and Guarantees in Low-Income Transfer Programs." *The Journal of Human Resources* 14 (1): 122–129.

Pittman, Elaine. 2012. "Virginia Launches Online Public Assistance System." Published by Government Technology magazine. Accessed: 2025-05-07. https://www.govtech.com/health/Virginia-Launches-Online-Public-Assistance-System.html.

Rossin-Slater, Maya, and Miriam Wüst. 2020. "What Is the Added Value of Preschool for Poor Children? Long-Term and Intergenerational Impacts and Interactions with an Infant Health Intervention." *American Economic Journal: Applied Economics* 12 (3): 255–286.

Rowe, Gretchen, Andrew Gothro, Elizabeth Brown, Lisa Dragoset, and Megan Eguchi. 2015. *Assessment of the Contributions of an Interview to SNAP Eligibility and Benefit Determinations: Final Report.* Technical report. May: U.S. Department of Agriculture, Food and Nutrition Service, Office of Policy Support.

Virginia Department of Social Services. 2012. *A Quick Guide to Using CommonHelp.* https://www.dss.virginia.gov/files/division/pa/commonhelp_community/guidance_procedures/qrg.pdf. Accessed: 2025-05-07.

Wherry, Laura R., Sarah Miller, Robert Kaestner, and Bruce D. Meyer. 2018. "Childhood Medicaid Coverage and Later-Life Health Care Utilization." *Review of Economics and Statistics* 100 (2): 287–302.