NETWORK LISER

Discussion Paper Series

IZA DP No. 18873

August 2026

Who Benefits from Digitizing Public Programs?

Neil A. Cholli
Cornell University

Derek Wu
University of Virginia
and IZA@LISER

IZA@LISER NETWORK
ISSN: 2365-9793

Who Benefits from Digitizing Public Programs?*

Abstract

Governments increasingly digitize access to public programs, but it remains unclear who benefits. Using administrative microdata from Virginia, we study the impact of an online portal that streamlined SNAP, Medicaid, and TANF applications. Among children and non-elderly adults, the portal raised enrollment by 1.7–2.0% for every ten miles between their residence and local program office. Yet it had no detectable effect on the elderly, despite their lower participation and comparable enrollment barriers. Predicted digital access accounts for 84% of the variation in effects across these groups, showing that digitization can expand access while excluding vulnerable groups like the elderly.

JEL classification

I38, H53, O33

Keywords

digital government, welfare programs, take-up, administrative burdens, elderly

Corresponding author

Derek Wu

derek.wu@virginia.edu

* We thank Manasi Deshpande, Matt Notowidigdo, and seminar and conference participants at the University of Virginia and Russell Sage Foundation Pipeline Conference for useful comments, 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 Heidi Clesner, Dylan Craig, and Jimmy Yung 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.

1 Introduction

Digital technology is remaking the way individuals access public programs. Interactions that once required a trip to a government office to submit paper forms are increasingly handled through online portals that screen for eligibility and accept applications electronically. These reforms are often justified as a way to reduce the time and effort of dealing with the government, which can otherwise lead eligible households to forgo benefits in settings ranging from college financial aid to tax filing (e.g., Bettinger et al. 2012; Herd and Moynihan 2018; Benzarti 2020). In principle, the households that would gain the most from digitization should be those who face the steepest barriers to applying in person. Yet that promise holds only for those able to use the online channel in the first place. If vulnerable groups are precisely the ones least able to get online, then digitization could change *who* the government reaches and not just *how many* it serves—leaving these groups behind even if enrollment rises overall.

This tension is especially acute in the U.S. social safety net, where barriers to enrollment are substantial and digitization has been both rapid and expansive.[1] In 2007, only two states offered online application portals across multiple means-tested programs, but by 2024 that number had reached 45 (Appendix Figure A.1). The next wave of digitization is already taking shape, as agencies experiment with AI tools to guide applicants, assist caseworkers, and automate eligibility verification (Conlan 2024; U.S. Department of Agriculture, Food and Nutrition Service 2026; Gosciak et al. 2026).

In this paper, we ask who benefits from digitizing safety net access, and who is left behind. We take up this question in Virginia, which introduced an online portal that offered 24/7 eligibility screening and streamlined applications and recertifications for some of the most prominent means-tested programs, including the Supplemental Nutrition Assistance Program (SNAP), Medicaid, and Temporary Assistance for Needy Families (TANF).

To identify the impacts of the online portal, we link the universe of individual-level records on participation in these three programs to the historical locations of field offices administering these programs. Traditionally, households would visit field offices to reach caseworkers and apply in person, imposing greater time and travel costs to those living farther from their office (Rossin-Slater 2013; Deshpande and Li 2019; Hicks 2024; Bitler et al. 2025; Cholli and Wu 2026). By allowing households to apply remotely, the portal reduced enrollment costs more for those living farther away. We therefore estimate a difference-in-differences model that compares program participation before versus after the portal's launch across zip codes varying in distance from their assigned office. This research design identifies

1. Beyond the U.S., examples of digitized safety net programs include Australia's Centrelink/myGov system, Brazil's Bolsa Família program, India's National Rural Employment Guarantee Scheme, and the U.K.'s Universal Credit system.

the impact of digitization on take-up among those facing greater geographic burdens, with more distant zip codes receiving a higher "dose" of treatment.

Our central finding is that digitization increased enrollment across the board for the non-elderly, while it had no effects on the elderly. Safety net attachment—measured by the total number of programs received—significantly increased for children by 1.7% per ten miles and non-elderly adults by 2.0% per ten miles, with effects strikingly consistent across finer age bins and household composition. In contrast, among the elderly (those aged 65 or older), digitization had no detectable effect, raising enrollment by a statistically insignificant 0.1% per ten miles. We can reject equality of the elderly and non-elderly adult effects ($p = 0.03$), while the child and non-elderly adult effects are statistically indistinguishable ($p = 0.53$). The null effect for the elderly persisted even following a 2015 portal upgrade that allowed electronic submission of supplemental documents.

The elderly's lack of response is striking for several reasons. Policymakers have long treated the elderly as a priority population, exempting them from SNAP's gross income test and granting those with low incomes categorical eligibility for Medicaid. Yet they have the lowest participation rates in SNAP of any age group, whether measured as a share of the overall or (approximately) eligible populations. And as we show, the elderly are no less exposed to geographic burdens, as all groups face similar distance gradients in per-capita program enrollment. Their low take-up is also not simply a matter of weak demand, since the elderly have been shown to respond to outreach and application assistance delivered through other means (e.g., Kauff et al. 2014; Finkelstein and Notowidigdo 2019). Nor did the reform displace existing enrollment options, as Virginia's portal operated alongside the paper application method.

Why did the elderly fail to respond to digitization? The evidence points most prominently to one mechanism: the elderly face greater barriers to digital access. Computer ownership and internet access fall steeply after age 65, especially among those most likely to qualify for these programs. If digital access were the binding constraint, then the reform's effects should be larger where such access is more widespread. Consistent with this prediction, we find that effects are concentrated in counties with greater broadband penetration, particularly among the non-elderly. As a sharper test, we construct a measure of predicted digital access for each group based on their average internet access and county's broadband availability. This predicted measure alone accounts for roughly 84% of the variation in digitization's effects across groups.

We consider and largely rule out other explanations for the heterogeneity in elderly versus non-elderly effects. It is not driven by differences in the programs the elderly enrolled in, as it persists for SNAP and Medicaid examined separately. Differences in benefit amounts also cannot explain the gap, as the elderly received considerably more Medicaid than younger enrollees. Nor do we find much evidence that the elderly simply value these programs less, as subgroups of the elderly with high revealed willingness-to-pay—including those previously connected to a program, with low pre-retirement wages, or with a criminal record—responded no more strongly than the rest. The gap also cannot be explained by the elderly being less likely to live with children, who tend to qualify households for more programs, or by the portal disproportionately easing earnings-documentation burdens for workers relative to retirees. Finally, we find little evidence that receipt of Medicare or Social Security substitutes for the programs we study.

Digitization thus operates differently from other reforms that expand access to public programs. Whereas simplifying forms or procedural requirements reaches all applicants alike, moving access online works through a channel that not everyone can use. Its impacts hinge on digital access, which can vary sharply across groups facing similar levels of enrollment burdens. Policymakers aiming to broadly increase take-up should therefore consider pairing online portals with avenues that do not require independent internet use, which will only grow more pressing as agencies integrate AI tools into enrollment.

Our results speak to several literatures. First, we complement a growing literature on the effects of digitizing government services and enrollment procedures. Prior work has found mixed evidence on participation, with some evidence of heterogeneity across subpopulations, but differences between elderly and non-elderly populations remain largely unexplored (e.g., Kopczuk and Pop-Eleches 2007; Ebenstein and Stange 2010; Schwabish 2012; Gray 2019; Briscese et al. 2022; Gallegos et al. 2023). Dahlstrand (Forthcoming) demonstrates that digital platforms alleviate geographic burdens in the context of Swedish medical care. The most cautionary evidence comes from reforms in which digital channels *replaced* in-person caseworker assistance (Eubanks 2018; Wu and Meyer, Forthcoming). Our results show that the limits of digitization are not confined to such cases, as even a *supplemental* portal systematically left older adults behind. In a companion paper, we show that the same portal promoted multiple program participation among households with children (Cholli and Wu 2026). Here, we identify a digital access barrier underlying the disparate effects we document across the broader population. In doing so, we connect to broader research on the digital divide, which has documented persistent age gaps in technology adoption (Hunsaker and Hargittai 2018; Pew Research Center 2021) but has seldom traced their consequences for participation in public programs.

Second, we contribute to a large body of work on administrative burdens and incomplete take-up (for reviews, see Currie 2006; Ko and Moffitt 2024). Much of this work has emphasized average effects or heterogeneity by income, race, and health (e.g., Deshpande and

Li 2019; Homonoff and Somerville 2021; Unrath 2021; Shepard and Wagner 2025; Ericson et al., Forthcoming). A separate strand of literature has documented the low take-up of programs—particularly SNAP—among the elderly (Haider et al. 2003; Gabor et al. 2002; Zuo and Heflin 2023) and shown that this population *does* respond to targeted interventions that reduce enrollment barriers (Cody and Ohls 2005; Kauff et al. 2014; Sama-Miller et al. 2014; Finkelstein and Notowidigdo 2019). Our results bridge these literatures by showing that digitization reduces burdens for many groups but not for the elderly, despite their responsiveness to other forms of enrollment assistance.

2 Data and Descriptive Facts

2.1 Programs and Data

Our primary data are individual-level administrative records covering the universe of SNAP, Medicaid, and TANF recipients in Virginia, obtained from the Virginia Department of Social Services (VDSS) and spanning 2007–2022. SNAP provides monthly benefits to low-income households that can be used to purchase groceries. Medicaid provides health insurance coverage primarily to low-income children, parents, pregnant women, and elderly or disabled individuals. TANF provides cash assistance and employment-related services to low-income families with children.

In addition to program-specific indicators for monthly receipt, these microdata contain some limited demographics such as age, gender, and zip code and county of residence. We approximate household characteristics such as presence of children and number of parents by linking individuals who share a case in any program during a given month.[2] We aggregate receipt data at the zip code-by-month level for various groups of interest (e.g., children, non-elderly adults, elderly).[3]

Our primary outcome is the total number of programs received by individuals in the relevant group within a given zip code and month. This measure reflects changes in both *whether* individuals receive any programs at all (extensive margin) and *how many* programs they receive conditional on participating (intensive margin). In contrast to take-up studies that focus on enrollment in a particular program, our count-based measure captures these multiple margins of participation that may be at play in our setting (Cholli and Wu 2026). Our primary source of within-county variation is the Haversine ("as the crow flies") distance

2. For example, this procedure recovers parent–child links even when the child appears on a child-only case, as long as the parent and child are listed jointly on a case in at least one other program.

3. Because zip codes can cross county lines and field office assignments are based on county of residence, we conduct our analysis at the zip-by-county level. We refer to zip-by-county pairs as "zip codes" for simplicity.

between each zip code's centroid and the address of its assigned county field office. We collect these addresses from archived versions of the VDSS website. We also assess robustness to alternative distance metrics like driving distance.

Finally, we bring in zip-level characteristics from the American Community Survey (ACS), which are reported at the Census-generated zip code tabulation area (ZCTA) and can be merged using zip-to-ZCTA crosswalks beginning in 2010. These data provide zip- and time-varying covariates for our descriptive and causal analyses. We additionally use the ACS to obtain counts of the full and low-income population by age and to characterize demographic patterns and digital access across the age distribution. Finally, we link the administrative VDSS records to Unemployment Insurance (UI) wage records and criminal conviction records to conduct additional heterogeneity analyses.

2.2 Descriptive Facts

Figure 1 illustrates three descriptive facts about safety net attachment that motivate our analysis. First, Panel (a) shows that attachment to the safety net declines steeply with age within Virginia's full population. This pattern partly reflects how these programs are designed. SNAP, Medicaid, and TANF are all more accessible to households with children, whether through more generous eligibility thresholds, larger benefit amounts, or fewer work requirements.[4] Among 1-year-olds in the full population, for example, 47% receive Medicaid, 31% receive SNAP, and 7% receive TANF—translating to more than 0.85 programs received per capita. By age 65, this drops to about 0.12 programs per capita. Program participation then rises slightly at the oldest ages, driven by Medicaid whose eligibility expands for the elderly just as the demand for medical care rises with age. SNAP and TANF receipt, by contrast, decline near-monotonically across the entire age distribution.

Does the low participation we observe at older ages simply reflect lower eligibility, or lower take-up among those who are eligible? To try and distinguish between these margins, we focus on SNAP, which has the most broadly accessible eligibility criteria across ages among the three programs we study. Panel (b) plots SNAP receipt by age both as a share of the full population and as a share of low-income individuals, defined as those with household income below 200% of the federal poverty line (FPL). While the latter is not a precise measure of SNAP take-up, it allows for an informative comparison.[5] Specifically, the age gradient

4. For Medicaid in particular, income eligibility thresholds during much of our sample period were 205% of the federal poverty line (FPL) for children but only 33% FPL for non-disabled parents and 80% FPL for the disabled and elderly. Virginia did not expand Medicaid to low-income non-elderly adults until 2019.

5. Virginia's SNAP gross income threshold was 130% FPL prior to 2021. We use 200% FPL for two reasons: (i) ACS income is reported annually while SNAP eligibility is determined monthly, and (ii) Census surveys are well known to under-report incomes (e.g., Meyer et al. 2015). Our 200% FPL benchmark misses

Figure 1: Program Participation by Age and Distance from Field Office, 2007–2019

(a) Number of Programs Received by Age

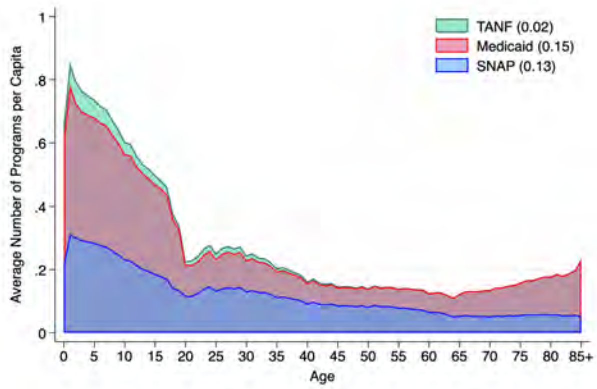

(b) SNAP Received by Age

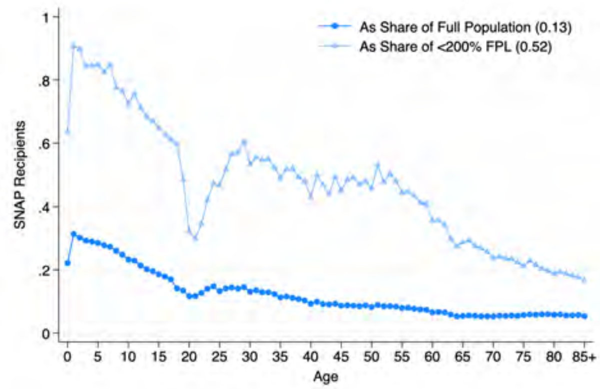

(c) Distance Gradient by Age Group

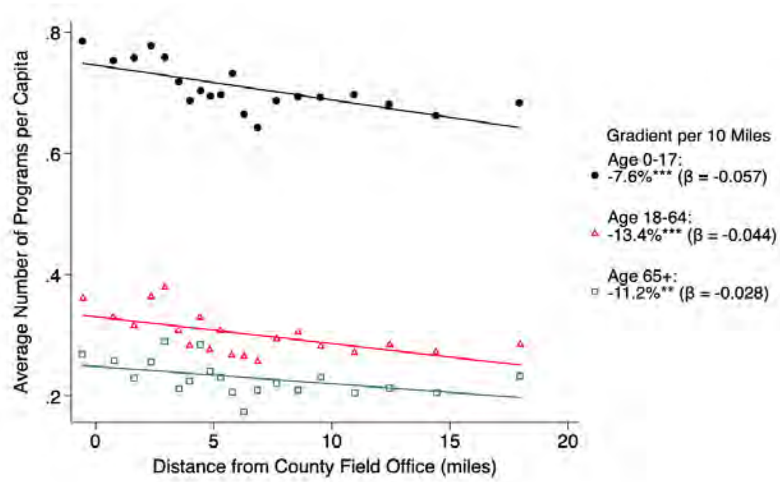

Notes: Panel (a) shows the average number of programs (out of SNAP, Medicaid, and TANF) received per capita (in the full population) by age in Virginia, pooled across 2007–2019. Overall receipt shares are reported in the legend. Panel (b) shows the share of SNAP recipients in the full population and in households with income below 200% of the federal poverty line (FPL) by age. Panel (c) illustrates zip-level binscatters of the number of programs received per capita by distance from the nearest county field office, separately for children (age 0–17), non-elderly adults (18–64), and the elderly (65+), controlling for zip code tabulation area (ZCTA)-level poverty rate and population shares of Blacks, Hispanics, children, elderly, those without a high school diploma, and college graduates. The legend reports the percent change per ten miles relative to the fitted intercept, with the underlying slope (per ten miles) in parentheses. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

in enrollment is steeper after conditioning on the low-income population—the opposite of what one would expect if eligibility differences across age were the primary driver of the

two features of the eligibility rules that cut in opposite directions. On one hand, individuals above age 60 are exempt from SNAP's gross income test, which may make the true eligible population larger than what our benchmark suggests. On the other hand, SNAP imposes an asset test that we do not incorporate, which may be more binding for the elderly if they are more likely to hold assets.

participation gradient. Less than 30% of low-income individuals above age 65 receive SNAP, compared with over 50% of low-income adults aged 25–55 and over 80% of low-income children below age 5. This holds even as the share of individuals below 200% FPL rises after age 50 (Appendix Figure A.2). The age gradient in SNAP participation thus appears to largely reflect differential take-up across ages.

Finally, Panel (c) establishes that all age groups—including the elderly—face geographic burdens to enrolling. The number of programs received per capita declines with distance to the assigned field office for every age group, even after controlling for zip-level characteristics that may also vary with distance. For the elderly, safety net attachment falls by a statistically significant 11.2% per ten miles, falling between the distance gradients for children (−7.6% per ten miles) and non-elderly adults (−13.4% per ten miles). These patterns hold within each individual program (Appendix Figure A.3) and when using driving or transit distance instead of Haversine distance (Appendix Figure A.4). Given their low participation rates and comparable exposure to geographic burdens, the elderly stand out as a vulnerable group that may especially benefit from reforms that lower enrollment burdens.

3 Empirical Strategy

3.1 The CommonHelp Reform

Previously, SNAP, Medicaid, and TANF applications in Virginia were administered exclusively through local VDSS field offices, with roughly one office serving each county. Households submitted a single paper application covering their desired programs to their assigned office (in person or by mail), and were required to return for documentation verification and periodic recertification. As in most states, the application was written in technical language, and applicants who needed help had to consult office caseworkers during business hours only.

In October 2012, Virginia launched CommonHelp, a centralized online application portal, alongside the existing paper-based system. A 24/7 online-screening tool recommended programs based on basic demographic and income inputs, reducing information frictions for households unaware of benefits they qualified for. The application process itself was guided, written at a 7th-grade reading level, and tailored to the applicant's selected programs (Code for America 2025). In January 2015, the portal was upgraded to allow applicants to upload electronic copies of supporting documents (e.g., proof of identity, residence, and earnings), after which the entire application and recertification process could be completed remotely.

A key feature of the reform for our identification strategy is that it disproportionately reduced burdens for households living farther from their assigned county field office. Prior to the reform, travel and time costs rose with distance. CommonHelp largely eliminated these costs for those who could complete the application online, thereby sparing households trips to the office. Households living in more distant zip codes thus stood to gain more from the online portal. As online applications have been launched in nearly every state, our findings are relevant and potentially generalizable across the country.

3.2 Research Design

To estimate the effect of CommonHelp's reduction in geographic burdens, we exploit this variation through a continuous "treatment dosage" difference-in-differences design. This compares participation before versus after CommonHelp's launch across zip codes located at varying distances from their field office. Our continuous specification offers advantages over a binarized "far-versus-close" comparison: it exploits all of the variation in distance and estimates an interpretable semielasticity parameter.

Using 2010–2015 data, we estimate a Poisson event-study specification of the form:

$$\mathbb{E}[Y_{zt} \mid \boldsymbol{W}_{zt}] = \exp\left(\sum_{\substack{-11 \leq k \leq 12 \\ k \neq -1}} \beta_k (D_z \times \mathbb{1}\{q(t) = k\}) + \alpha D_z + \theta_{c(z)} + \delta_t + \lambda_{c(z)} t \right), \tag{1}$$

where Y_{zt} is the number of programs received in zip code z in month t (summed over individuals in the relevant group), D_z is distance in miles to the county field office, $q(t)$ is the calendar quarter, $c(z)$ is the county of zip code z, $\theta_{c(z)}$ and δ_t are county- and month-fixed effects, and $\boldsymbol{W}_{zt}$ collects all covariates. Because our design compares zip codes served by the same county office, these estimates are not confounded by broad rural-urban differences across the state. The coefficients β_k capture the proportional change in zip-level program receipt per mile of distance from the field office in event-quarter k. Throughout, we report these effects as percent changes per ten miles, given by $\exp(10\beta_k) - 1$. We also include county-specific time trends, $\lambda_{c(z)}t$, to absorb secular trends in within-county demographic composition that could otherwise confound our estimates. We restrict the sample to zip-county pairs averaging at least 5 programs received per month over the pre-period (January 2010–September 2012). Because this threshold is based on total receipt rather than group-specific counts, every analysis uses the same set of zip codes. We cluster standard errors at the county level.

The main advantage of the Poisson specification over log-linear alternatives is that it accommodates zero-valued outcomes without ad hoc transformations (Chen and Roth 2024). This is particularly relevant for our setting, where some zip-month-group cells can contain zero program enrollees. We nonetheless confirm robustness to $\log(Y+1)$ and inverse hyperbolic sine specifications and find quantitatively similar results. The key identifying assumption is the "strong parallel trends" assumption, which requires that, absent CommonHelp, average program receipt would have followed the same (proportional) trend at every distance level within the same county. Callaway et al. (2024) show that, under this assumption, equation (1) identifies a positive weighted sum of average responses over the distribution of distances.

To summarize the overall average effects of CommonHelp on program enrollment, we also estimate a Poisson pre-post difference-in-differences specification:

$$\mathbb{E}[Y_{zt} \mid \boldsymbol{W}_{zt}, \boldsymbol{X}_{zt}] = \exp\bigl(\beta(D_z \times \mathbb{1}\{q(t) \geq 0\}) + \alpha D_z + \boldsymbol{X}_{zt}\boldsymbol{\gamma} + \theta_{c(z)} + \delta_t + \lambda_{c(z)}t\bigr)\,. \tag{2}$$

The coefficient β can be thought of as a weighted average of the post-launch event-study coefficients β_k's from equation (1). As with the event-study coefficients, we report β as a percent change per ten miles. To improve precision, we control for a vector $\boldsymbol{X}_{zt}$ of time-varying zip-level characteristics drawn from ACS 5-year estimates.[6] Finally, it is worth noting that equations (1) and (2) identify the *relative* impacts of CommonHelp on program participation in farther versus closer zip codes, reflecting reductions in geographic burdens. Any common reductions in burdens affecting all zip codes equally, such as those generated by simplified application forms, are absorbed by the time-fixed effects and not captured by β.[7] That said, geographic burdens are central to the case for digital reform, as they motivate the very design of online portals as substitutes for in-person field-office visits.

4 Main Results

Pooled across all individuals, CommonHelp generated a statistically significant 1.6% increase in the number of programs received per ten miles of distance from the field office ($p < 0.05$). Panel (a) of Figure 2 shows that the effect jumps to 1% in the quarter immediately after launch, rises gradually over the ensuing quarters, and jumps to over 3% following the January 2015 upgrade. This pattern is consistent with progressive learning about the portal and an additional discrete reduction in burdens associated with the ability to electronically upload documents. The pre-launch coefficients are flat and precisely estimated near zero, supporting

6. $\boldsymbol{X}_{zt}$ includes log total population; shares of Black, elderly, and female residents; and poverty rate.

7. Cholli and Wu (2026) identify CommonHelp's effects on distance-invariant burdens by detecting discontinuities in program receipt among zip codes closest to field offices. Due to its local nature and restricted sample, this method demands substantial statistical power, which is not feasible when estimating effects across the granular age groups we examine here.

Figure 2: CommonHelp's Effects on Program Receipt in the Population

(a) Event-Study Estimates, Full Population

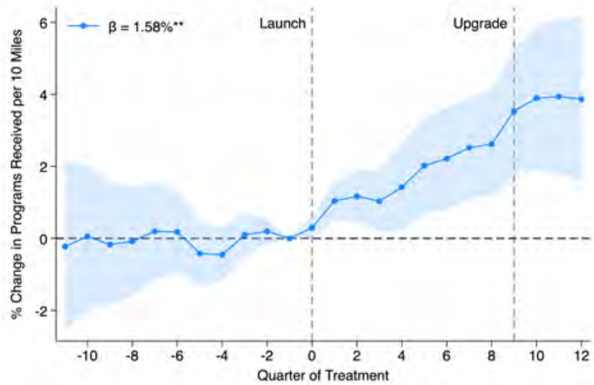

(b) Event-Study Estimates, by Age Group

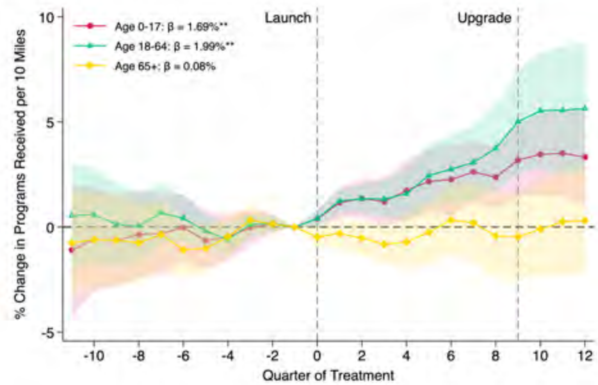

(c) Summary Effects by Age and Household Type Subgroups

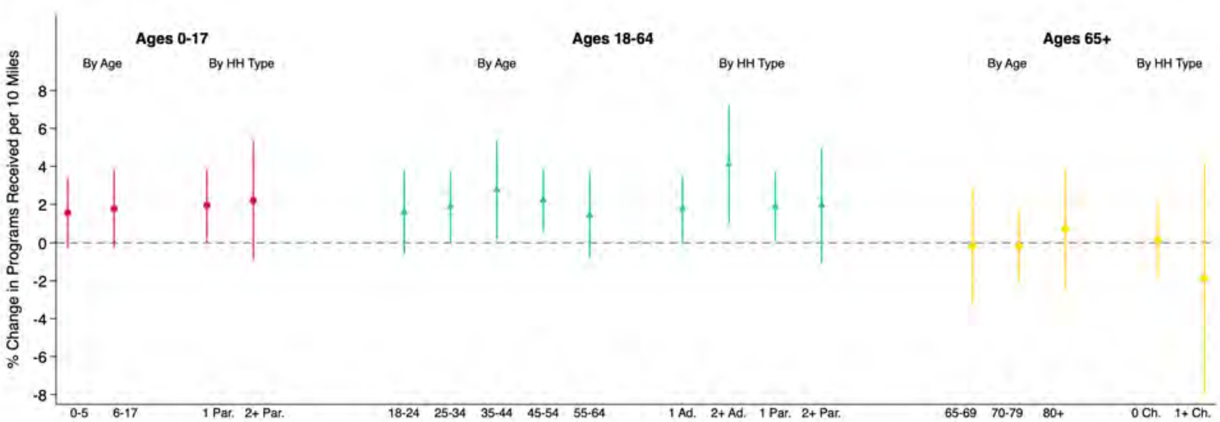

Notes: Panels (a) and (b) plot the percent change per ten miles implied by β_k from equation (1) for the full population and separately for children (ages 0–17), non-elderly adults (ages 18–64), and elderly (age 65 and above). These are given by $100 \times [\exp(10\beta_k) - 1]$, with standard errors obtained by the delta method. Pooled estimates of β from equation (2) are reported in the legend on the same scale. Panel (c) reports the corresponding estimates from equation (2) for more granular age bins and for bins by age × household type. Confidence bands/intervals are at the 95% level.

the parallel trends assumption; any alternative explanation would need to reproduce not only this flat pre-period but also the sharp increases at both the launch and the 2015 upgrade.[8] Aggregating over the distribution of zip code distances and caseloads, we estimate that the digital reform's reductions in geographic burdens increased overall program receipt by 1%.

However, the pooled effect masks a sharp divide by age. For children (age 0–17), CommonHelp increased safety net attachment by 1.7% per ten miles ($p < 0.05$), translating

8. Appendix Figure A.5 plots raw trends in the log number of programs received in distant versus close zip codes, providing further visual support for parallel trends prior to CommonHelp's launch.

into an aggregate effect of 1.1% across zip codes. Because children usually do not apply on their own, this effect likely reflects the responses of parents or guardians who enroll them, whether on a shared or child-only case. For non-elderly adults (age 18–64), the summary effect is slightly larger at 2.0% per ten miles ($p < 0.05$), or an aggregate effect of nearly 1.3%. Panel (b) shows that the event-study trajectories for these two groups evolve nearly identically in the first eight quarters after launch but diverge after the 2015 upgrade, with the effect for non-elderly adults rising to over 5% per ten miles compared to roughly 3% for children. For the elderly (age 65+), by contrast, we document essentially no effect of CommonHelp. The pooled point estimate is 0.1% per ten miles and statistically indistinguishable from zero, and none of the post-launch or post-upgrade event-study coefficients are significantly different from zero. We can reject equality of the elderly versus non-elderly adult effects ($p = 0.03$) and of the elderly versus children effects ($p = 0.09$), but cannot reject equality of the children versus non-elderly adult effects ($p = 0.53$).

The principal margin of heterogeneity thus appears to be between the elderly and the non-elderly. To reinforce this conclusion, we examine whether substantial heterogeneity exists *within* either group by estimating effects across more granular subgroups. Panel (c) reports pre-post estimates corresponding to β from equation (2) for these subgroups.[9] Across seven age bins covering the non-elderly population (age 0–64), the effect sizes are strikingly consistent. Prime-aged adults show somewhat larger magnitudes, but five of the seven bins fall within a tight range of 1.4–1.9% per ten miles, and a joint test of equality across the seven cannot reject that the effects are the same ($p = 0.89$). A second cut by household composition tells a similar story for the non-elderly. We examine effects separately for children in single- versus multi-parent households and for non-elderly adults disaggregated by single adults without children, multiple adults without children, single parents, and multiple parents. With the exception of non-elderly adults living with other adults but without children, the effects for the remaining five subgroups fall within a similarly tight range of 1.7–2.2% per ten miles, and we cannot reject equality across these six subgroups ($p = 0.63$).

The elderly, in contrast, show effects close to zero across all granular age subgroups (ages 65–69, 70–79, and 80+), indicating that the null result is not driven by the oldest members of the group. We also find no evidence of responses to CommonHelp for elderly living with or without children, which we discuss further below. In sum, all non-elderly subgroups responded to CommonHelp at a similar rate, while no elderly subgroup did. This reinforces that the heterogeneity runs between the elderly and non-elderly, not within either group.

9. Appendix Figure A.6 reports the corresponding event-study estimates and Appendix Tables A.1 and A.2 report the coefficient estimates (both overall and separately for the initial effect after the launch and the additional effect after the upgrade).

We confirm the robustness of our estimates to a range of specification choices, reported in Appendix Table A.3. They are stable to the inclusion of time-varying zip-level covariates, to alternative functional forms ($\log(Y+1)$ and inverse hyperbolic sine), and to the exclusion of the most distant zip codes. They are also nearly identical under alternative definitions of distance, including driving distance and the minimum of driving and transit distances calculated via the Google Maps API. Finally, we estimate separate effects for discrete distance bins rather than imposing a linear relationship (Appendix Figure A.7). Effects are near zero in the closest bins and grow at greater distances, tracing out the dose-response our continuous specification assumes and easing the weighting concerns that arise in continuous-treatment designs (Callaway et al. 2024).

5 Why Did the Elderly Not Respond?

Why did the elderly fail to respond to CommonHelp? The evidence points most consistently to digital access. The portal lowered enrollment costs only for those who could get online, and this was the barrier the elderly were unable to overcome. We first present evidence for this mechanism, before showing that alternative explanations cannot explain the difference in effects across ages.

5.1 Digital Access

Because CommonHelp was initially optimized for desktop computers, accessing it likely required a computer and home internet connection. If digital access falls steeply with age, then this mechanism would generate precisely the pattern we observe: meaningful effects among the digitally connected non-elderly and null effects among the less connected elderly.

Using estimates from the ACS, we start by showing sharp declines in computer ownership, internet access, and high-speed internet with age. Among the full population in Virginia, Panel (a) of Figure 3 shows high and stable rates of internet and computer access across non-elderly ages that sharply decline around or after age 65. Only 30–50% of the elderly had access to high-speed internet in 2013, compared to around 50–60% at younger ages. Among the elderly with incomes below 200% FPL, Panel (b) shows that this share drops further to 10–40%. These gaps in access may be further compounded by differences in functional limitations, as the share of elderly individuals with cognitive or visual difficulties also rises sharply with age (Appendix Figure A.8). Even an elderly applicant with a computer and home internet may have struggled to complete a digital application if they could not easily read the screen or work through the steps.

Figure 3: Digital Access and the Effects of CommonHelp

(a) Digital Access by Age, Full Population

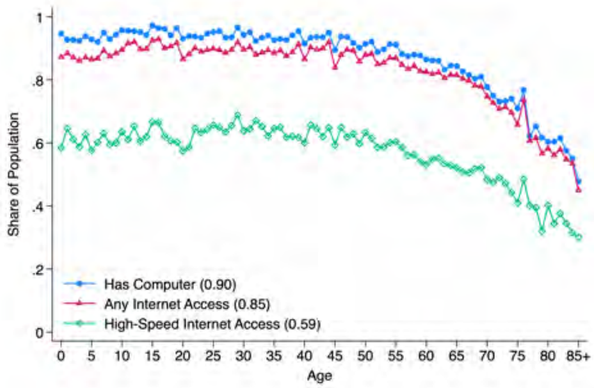

(b) Digital Access by Age, Income <200% FPL

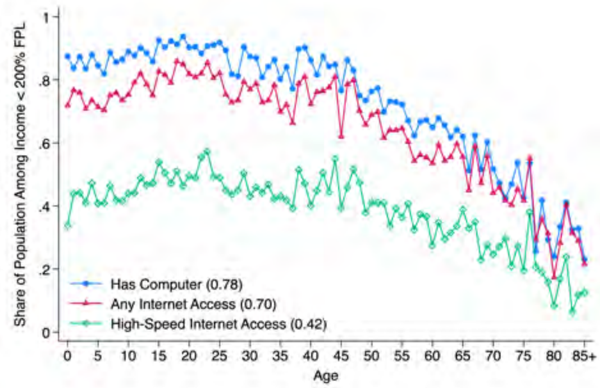

(c) Event-Study Estimates by Broadband Level

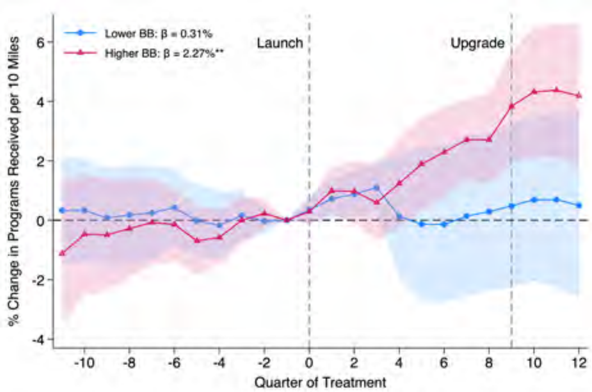

(d) Effects by Predicted Digital Access

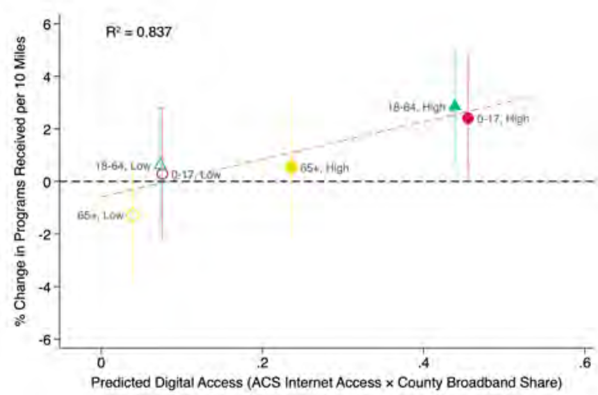

Notes: Panels (a) and (b) report the share of individuals with computer and internet access by age using the 2013 ACS for Virginia (with population-weighted means for each year in the legend). Panel (a) reports estimates for the full population. Panel (b) restricts the sample to individuals with household incomes below 200% FPL, who are more likely to be eligible for means-tested public programs. Panel (c) plots the percent change per ten miles implied by β_k from equation (1) for the full population, separately for counties with higher and lower broadband access. We define counties with higher-broadband access as those in which at least 20% of households have download speeds of at least 25 mbps (splitting counties roughly in half), based on 2014 statistics collected from the FCC. Panel (d) plots each subgroup's estimate from equation (2) on the same scale against its predicted digital access, defined as the age group's internet access share (2013 ACS, <200% FPL) times its county broadband share. The dashed line is the precision-weighted line of best fit ($R^2 = 0.84$).

If digital access were indeed the binding constraint, then CommonHelp's effects should be larger in areas where internet access was more widespread. Consistent with this prediction, Panel (c) shows that effects are concentrated in counties with higher broadband access, where enrollment increased by a statistically significant 2.2% per ten miles.[10] Conversely,

10. We define counties with higher-broadband access as those in which at least 20% of households had download speeds of at least 25 megabits per second (mbps), based on 2014 Federal Communications Commission (FCC) statistics.

effects are largely absent in lower-broadband counties, where enrollment increased by an insignificant 0.3% per ten miles. Unsurprisingly, this contrast is sharpest among non-elderly adults (who have greater digital access), while effects for the elderly remain insignificant in both higher- and lower-broadband counties (Appendix Table A.4). That said, in higher-broadband counties where the elderly more often live with others, individuals aged 80 and above show larger—albeit imprecise—point estimates, potentially reflecting help from more digitally literate co-residents (Appendix Table A.5).

Panel (d) subjects this mechanism to a sharper test. We divide the population into six subgroups defined by age (children, non-elderly adults, and the elderly) and county broadband access (higher or lower). For each subgroup, we construct a measure of predicted digital access equal to the share of its members with internet access (among those in Virginia with incomes below 200% FPL), multiplied by the share of households in their county with broadband speeds above 25 mbps.[11] We then plot each subgroup's estimated effect against its predicted digital access. If differential digital access drives the heterogeneity in effects across ages, then effects should rise with predicted access. Not only do we find this to be the case, but the relationship is also close to linear. In fact, predicted digital access alone accounts for roughly 84% of the variation in effects across the six subgroups. Taken together, the age gradient in digital access, the concentration of effects in higher-broadband counties, and the tight relationship between predicted digital access and estimated effects all point to digital access as the key binding constraint.

5.2 Ruling Out Alternative Explanations

Besides digital access, we consider leading alternative explanations and find little evidence that any can account for the elderly's null response.

The elderly enrolled in programs less impacted by CommonHelp. Programs may have differed in how much CommonHelp lowered their barriers to enrollment. Because the elderly enrolled disproportionately in Medicaid and rarely in TANF, their muted response could have arisen if Medicaid was the program least affected by CommonHelp, perhaps because it carried lower baseline burdens (individuals can enroll through alternative channels such as hospitals). Under this hypothesis, the age heterogeneity in effects should disappear once we hold programs fixed. But when estimating program-specific effects, we instead

11. The internet access share is a weighted average of the age-specific shares in Panel (b) within each age group. The FCC reports county broadband connections in binned ranges rather than exact shares, so we assign each county the midpoint of its range: 0.1 for lower-broadband counties (0–200 connections per 1,000 housing units) and 0.6 for higher-broadband counties (200+ per 1,000 housing units).

find that elderly effects remain insignificant for SNAP and Medicaid separately, while non-elderly effects are positive, generally significant, and similar in magnitude across programs (Appendix Table A.6).[12] As a result, the age heterogeneity persists within programs.

The elderly qualified for smaller benefit amounts. Under this hypothesis, age differences in effects should track age differences in benefit amounts, with larger responses among those qualifying for more. A comparison of average benefits among 2013 enrollees in Virginia cuts against this in two ways.[13] First, the elderly received lower SNAP amounts than younger enrollees ($77 per person per month, versus $135 for children and $137 for non-elderly adults) but their Medicaid benefits were more than twice as high ($1,075 per person per month, versus $225 and $778). Yet, the age heterogeneity in effects is similar for both programs. Second, children and non-elderly adults responded similarly despite differences in their own benefit amounts. As a result, effect magnitudes do not appear to correlate with benefit amounts across ages.

The elderly placed lower value on the programs. The elderly may have faced greater information or stigma frictions or possessed lower socioeconomic need. If so, then elderly subgroups with higher revealed willingness to pay (WTP) should respond more strongly to CommonHelp. Yet we find no such pattern across any classification of higher WTP individuals (Appendix Table A.7). Elderly individuals who participated in SNAP, Medicaid, or TANF during 2007–2011 (prior to CommonHelp), who by definition demonstrated positive WTP and likely overcame any information or stigma barriers, show effects that are if anything more clearly null than those with no prior connection. The same holds for those connected to multiple programs and those whose prior participation had lapsed by September 2012, a group with more room to enroll in response to the portal. Nor do effects differ for elderly individuals aged 65–70 who earned low (but positive) wages between ages 55–61, a group that exhibited need before retirement but was unlikely to have been disabled or out of the labor force.[14] Finally, elderly individuals with a prior criminal record, who may have been more disadvantaged and thus likely valued program benefits more, also show insignificant

12. We do not report TANF effects for the elderly, since essentially no elderly individuals receive TANF.

13. Our administrative records contain program receipt indicators but not benefit amounts, so we use Quality Control records for SNAP and CMS records for Medicaid to characterize benefits by age. Because SNAP benefits are issued at the case level, we divide each case's benefit by the number of people on the case. Medicaid benefits are already reported at the person level. These figures reflect amounts received by enrollees, who may differ from the eligible non-enrollees that are most relevant for our estimates. We nonetheless view the comparison as informative, since benefit amounts among enrollees likely correlate with those a marginal applicant would receive.

14. We exclude those with zero wages between ages 55–61 since this may reflect disability, not high WTP.

effects like those without one.[15] The elderly have also been shown to respond to interventions delivered through other channels (e.g., Kauff et al. 2014; Finkelstein and Notowidigdo 2019), which is difficult to reconcile with low valuation of these programs.

The elderly were less likely to live with children. CommonHelp may have especially benefited households with children, who were more likely eligible for multiple programs and may have faced more involved applications. If so, then CommonHelp's effects should be concentrated among households with children. CommonHelp increased the number of programs received by 1.7% per ten miles for those in households with children and 1.4% for those without (Appendix Table A.2). The similarity of these two statistically significant estimates is the opposite of what the hypothesis predicts, and the elderly show no response whether or not children are present (Figure 2, Panel (c)).

The elderly were mostly retired and had fewer earnings to document. CommonHelp, and particularly its 2015 upgrade allowing the electronic upload of documents, reduced the burdens of submitting supporting documents. This may have benefited households with wage earners, especially those with volatile earnings that must be documented and verified more frequently. Because most elderly individuals are retired, this channel could generate the age heterogeneity in effects. But using linked UI wage records, we find similar effects across the not-employed, low-volatility employed, and high-volatility employed over the full analysis period (Appendix Table A.8). After the upgrade, effects were larger for the high-volatility employed than the low-volatility employed—consistent with reduced documentation burdens—but were just as strong as the not-employed, which is difficult to reconcile with the hypothesis (Appendix Table A.9). Alternative volatility measures based on the number of jobs held and job-switching behavior yield similarly mixed patterns. Thus, reduced documentation burdens may account for a small part of the upgrade-period effect but do not drive the overall age heterogeneity, and the elderly remain unresponsive even among those employed.

The elderly substituted toward other benefit programs. Programs targeting the elderly, such as Old-Age and Survivors Insurance (OASI) or Medicare, could crowd-out SNAP and Medicaid benefits. We assess this descriptively by checking whether monthly SNAP or Medicaid receipt fell as individuals reached eligibility for OASI at age 62 or Medicare at age 65, conditioning on receipt of SNAP or Medicaid 12–24 months before each age cutoff.

15. A criminal record did not generally affect SNAP or Medicaid eligibility in Virginia during our study period, although separate restrictions applied to some drug-related felony convictions.

We find no sharp discontinuities at either cutoff, although we detect a modest inflection in SNAP receipt after age 62 (Appendix Figure A.9). These changes are too small to account for CommonHelp's null effects among the elderly, especially in their Medicaid receipt.

6 Conclusion

Digital reforms are reshaping how people enroll in public programs, motivated by the promise of lowering burdens and expanding enrollment. We provide evidence from Virginia's CommonHelp portal that digitization can substantially increase enrollment for younger populations while leaving elderly individuals behind. The null result for the elderly is notable given that they have among the lowest participation rates and face geographic burdens similar to those of younger adults. The gap appears to be driven primarily by limited digital access among the elderly, rather than by differences in program mix, the value placed on benefits, household composition, or documentation burdens. Computer ownership and internet use fall steeply at older ages, especially among the low-income population. Consistent with this, the portal's effects are concentrated in counties with greater broadband access, and a simple measure of predicted digital access accounts for 84% of the variation in effects across groups.

Importantly, the digital divide we document is not specific to the setting we study. Age gaps in internet access persist through 2024 (Appendix Figure A.10) and appear nationally beyond Virginia (Appendix Figure A.11). They also extend to smartphone ownership, which might have been expected to narrow the digital access gap but show a similarly steep decline at older ages (Appendix Figure A.12). With nearly every state now operating online portals, similar constraints likely shape who is reached by those reforms.

Policy efforts to broadly increase program take-up should accordingly use digital portals to supplement—rather than replace—avenues that do not require independent internet use, such as staffed telephone applications, in-person assistance, and partnerships with community organizations and institutional providers like nursing homes that can apply on behalf of older adults. This will only become more pressing as public programs integrate AI tools in enrollment systems, which risk compounding digitization's uneven effects within the population.

References

Benzarti, Youssef. 2020. "How Taxing Is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs." *American Economic Journal: Economic Policy* 12 (4): 38–57. https://doi.org/10.1257/pol.20180664.

Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block FAFSA Experiment." *The Quarterly Journal of Economics* 127 (3): 1205–1242. https://doi.org/10.1093/qje/qjs017.

Bitler, Marianne, 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." Forthcoming, *Tax Policy and the Economy.*

Briscese, Guglielmo, Giulio Zanella, and Veronica Quinn. 2022. "Providing Government Assistance Online: A Field Experiment with the Unemployed." *Journal of Policy Analysis and Management* 41 (2): 579–602. https://doi.org/10.1002/pam.22368.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2024. *Difference-in-Differences with a Continuous Treatment.* Technical report. National Bureau of Economic Research.

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

Cholli, Neil A., and Derek Wu. 2026. "Multiple Program Participation in the Safety Net: Incidence, Impediments, and Implications." Working paper.

Code for America. 2025. *The Benefits Enrollment Field Guide.* https://codeforamerica.org/explore/benefits-enrollment-field-guide/.

Cody, Scott, and James C. Ohls. 2005. *Evaluation of the USDA Elderly Nutrition Demonstrations: Volume I, Evaluation Findings.* Contractor and Cooperator Report CCR-9-1. U.S. Department of Agriculture, Economic Research Service, July. https://www.ers.usda.gov/publications/pub-details?pubid=85968.

Conlan, Chris. 2024. "AI Could Help Clear Food Stamp Backlogs Plaguing Some States." Available at https://www.route-fifty.com/emerging-tech/2024/02/ai-could-help-clear-food-stamp-backlogs-plaguing-some-states/394590/, February.

Currie, Janet. 2006. "The Take-Up of Social Benefits." In *Public Policy and the Income Distribution,* edited by Alan J. Auerbach, David Card, and John M. Quigley, 80–148. New York: Russell Sage Foundation.

Dahlstrand, Amanda. Forthcoming. "Defying Distance? The Provision of Medical Services in the Digital Age." *American Economic Review.*

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. https://doi.org/10.1257/pol.20180076.

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.

Ericson, Keith Marzilli, Timothy J. Layton, Adrianna McIntyre, and Adam Sacarny. Forthcoming. “Reducing Administrative Barriers Increases Take-Up of Subsidized Health Insurance Coverage: Evidence from a Field Experiment.” *Review of Economics and Statistics.*

Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor.* New York: St. Martin’s Press.

Finkelstein, Amy, and Matthew J. Notowidigdo. 2019. “Take-Up and Targeting: Experimental Evidence from SNAP.” *Quarterly Journal of Economics* 134 (3): 1505–1556. https://doi.org/10.1093/qje/qjz013.

Gabor, Vivian, Susan Schreiber Williams, Hilary Bellamy, and Brooke Layne Hardison. 2002. *Seniors’ Views of the Food Stamp Program and Ways To Improve Participation: Focus Group Findings in Washington State.* Final Report E-FAN-02-012. U.S. Department of Agriculture, Economic Research Service, June.

Gallegos, Sebastian, Benjamin Roseth, Ana Cuesta, and Mario Sánchez. 2023. “Increasing the Take-Up of Public Health Services: An At-Scale Experiment on Digital Government.” *Journal of Public Economics* 227:104975. https://doi.org/10.1016/j.jpubeco.2023.104975.

Gosciak, Jennah, Eric Giannella, Zhaowen Guo, Michael Chen, and Allison Koenecke. 2026. *LLMs in Social Services: How Does Chatbot Accuracy Affect Human Accuracy?* arXiv: 2603.11213 [cs.CY]. https://arxiv.org/abs/2603.11213.

Gray, Colin. 2019. “Leaving Benefits on the Table: Evidence from SNAP.” *Journal of Public Economics* 179:104054. https://doi.org/10.1016/j.jpubeco.2019.104054.

Haider, Steven J., Alison Jacknowitz, and Robert F. Schoeni. 2003. “Food Stamps and the Elderly: Why Is Participation So Low?” *Journal of Human Resources* 38 (Supplement): 1080–1111. https://doi.org/10.2307/3558982.

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

Hicks, Jeffrey. 2024. “Distance, Information, and the Take-Up of Disability Insurance.” Working paper.

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

Hunsaker, Amanda, and Eszter Hargittai. 2018. “A Review of Internet Use Among Older Adults.” *New Media & Society* 20 (10): 3937–3954. https://doi.org/10.1177/1461444818787348.

Kauff, Jacqueline, Lisa Dragoset, Elizabeth Clary, Elizabeth Laird, Libby Makowsky, and Emily Sama-Miller. 2014. *Reaching the Underserved Elderly and Working Poor in SNAP: Evaluation Findings from the Fiscal Year 2009 Pilots.* Technical report. Washington, DC: Mathematica Policy Research.

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

Kopczuk, Wojciech, and Cristian Pop-Eleches. 2007. "Electronic Filing, Tax Preparers and Participation in the Earned Income Tax Credit." *Journal of Public Economics* 91 (7–8): 1351–1367. https://doi.org/10.1016/j.jpubeco.2006.11.003.

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

Pew Research Center. 2021. *Mobile Technology and Home Broadband 2021.* Technical report. Washington, DC: Pew Research Center.

Rossin-Slater, Maya. 2013. "WIC in Your Neighborhood: New Evidence on the Impacts of Geographic Access to Clinics." *Journal of Public Economics* 102:51–69. https://doi.org/10.1016/j.jpubeco.2013.03.009.

Sama-Miller, Emily, Libby Makowsky, Gretchen Rowe, Elizabeth Brown, Elizabeth Clary, Laura Castner, and Miki Satake. 2014. *Evaluation of the Effectiveness of Pilot Projects in Increasing Supplemental Nutrition Assistance Program (SNAP) Participation among Medicare's Extra Help Population.* Nutrition Assistance Program Report 06843.160. Prepared for the U.S. Department of Agriculture, Food and Nutrition Service, Office of Policy Support. Contract No. AG-3198-D-10-0079. Mathematica Policy Research, December. https://www.fns.usda.gov/research/snap/evaluation-effectiveness-pilot-projects-increasing-supplemental-nutrition-assistance-program-snap.

Schwabish, Jonathan A. 2012. "Downloading Benefits: The Impact of Online Food Stamp Applications on Participation." Working Paper.

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.

U.S. Department of Agriculture, Food and Nutrition Service. 2026. *Use of Advanced Automation in SNAP.* Available at https://www.fns.usda.gov/snap/advanced-automation.

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

Wu, Derek, and Bruce D. Meyer. Forthcoming. "Certification and Recertification in Welfare Programs: What Happens When Automation Goes Wrong?" *American Economic Journal: Economic Policy.*

Zuo, Dongmei, and Colleen M. Heflin. 2023. "Cognitive Impairment and Supplemental Nutrition Assistance Program Take-Up Among the Eligible Older Americans." *The Journals of Gerontology: Series B* 78 (1): 99–110. https://doi.org/10.1093/geronb/gbac111.

A Supplemental Figures and Tables

Figure A.1: Adoption of Integrated Digital Portals in the U.S. over Time

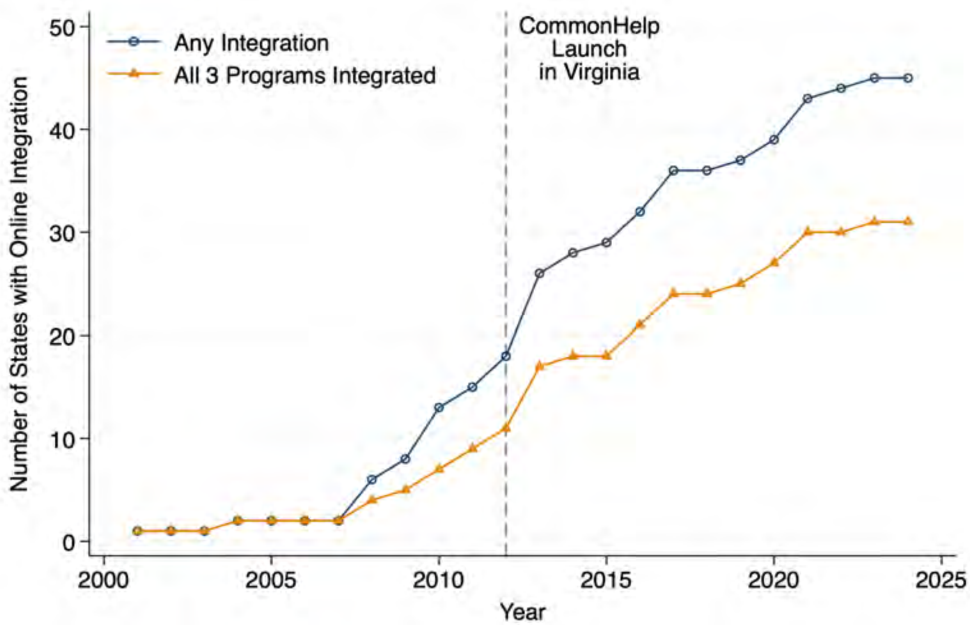

Notes: This figure shows the trend over time in the number of states adopting integrated digital portals for safety net programs. The count of states in 2024 is based on Code for America data, while counts for earlier years are based on manual searches for the start date of integration. We show separate 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.2: Share of Population with Income Below 200% FPL, 2007–2019

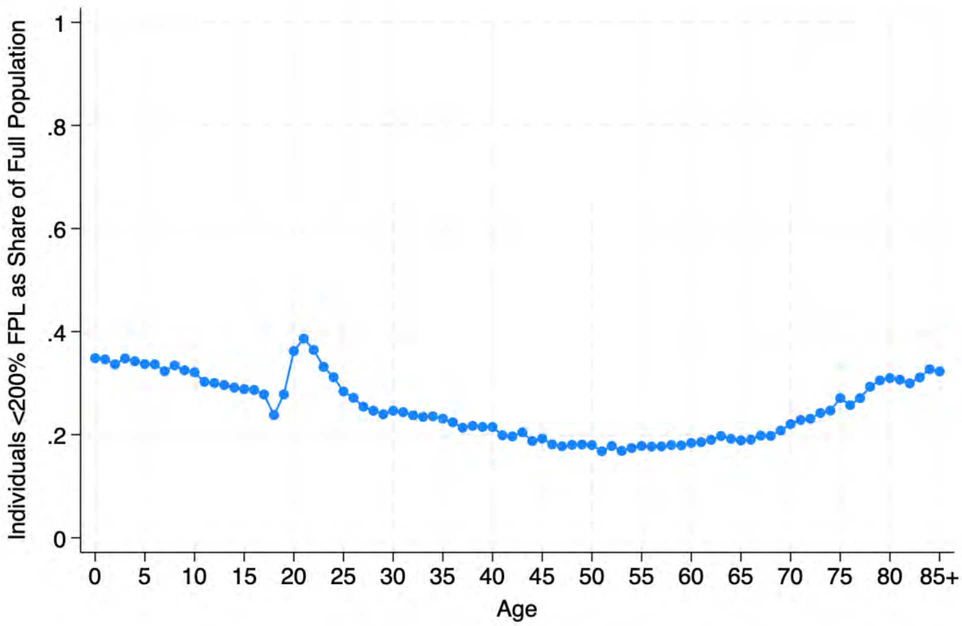

Notes: This figure plots the share of Virginia's population with household income below 200% of the federal poverty line (FPL) by each age, using data from the ACS and pooled across 2007–2019.

Figure A.3: Distance Gradient in Per Capita Enrollment, by Program and Age Group

(a) SNAP

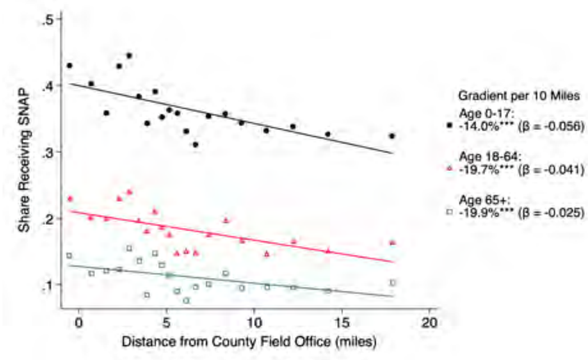

(b) Medicaid

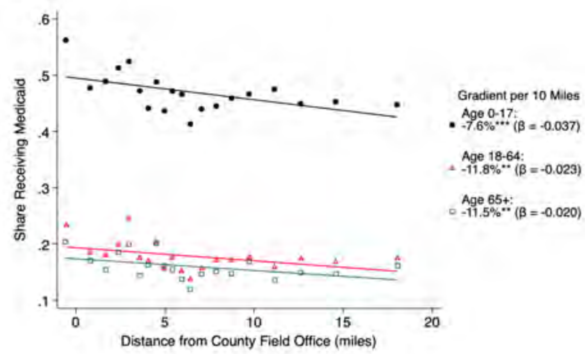

(c) TANF

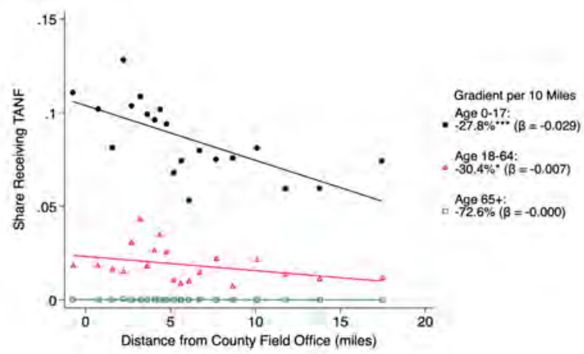

Notes: This figure plots binscatters of the per-capita share of the relevant age group enrolled in the indicated program (SNAP, Medicaid, or TANF) against the zip code's distance from the assigned county field office, pooling 2007–2022. See notes to Panel (c) of Figure 1 for more details.

Figure A.4: Distance Gradient in Number of Programs Per Capita, Robustness Checks

(a) No Covariates

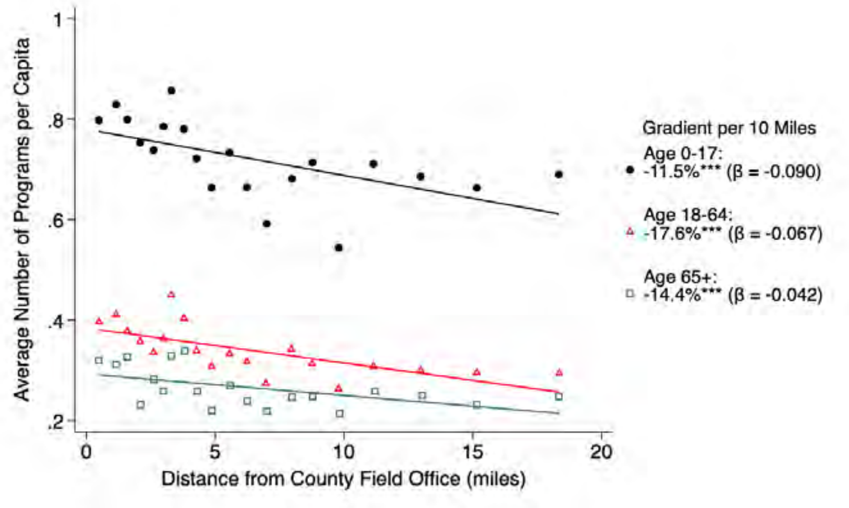

(b) Using min{Driving, Transit} Distance

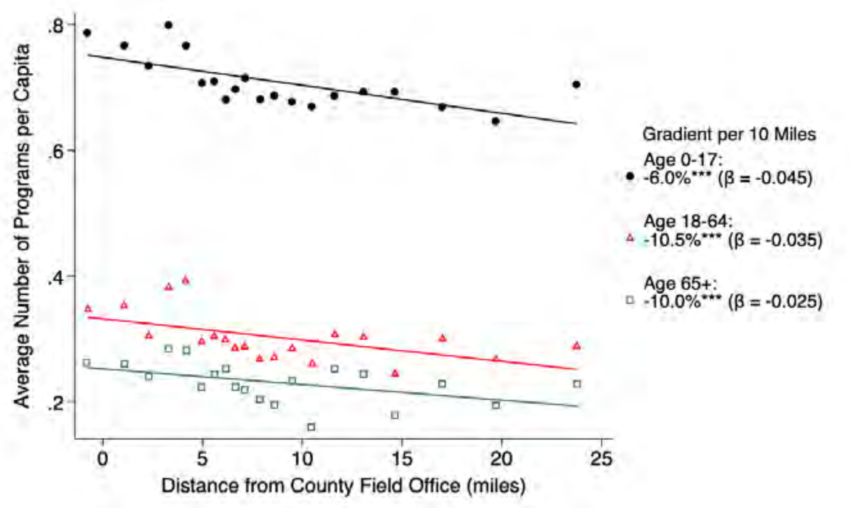

Notes: This figure plots alternative binscatters of the average number of programs received per capita by age group against the zip code's distance from the assigned county field office, pooling 2007–2022. Panel (a) does not include any zip- or time-varying covariates. Panel (b) measures distance using the minimum of driving and transit distance (as opposed to Haversine distance). See notes to Panel (c) of Figure 1 for more details.

Figure A.5: Raw Trends in Number of Programs Received in Distant vs. Close Zip Codes

(a) Full Population

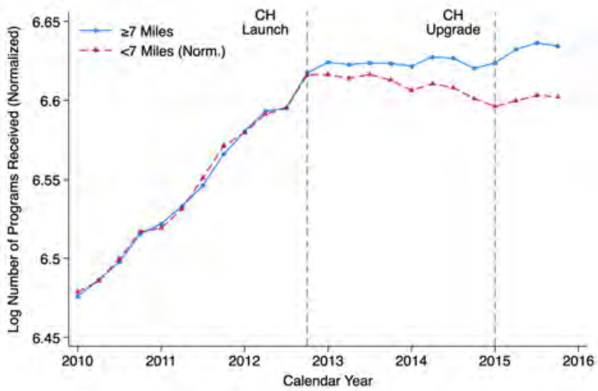

(b) Age 0–17

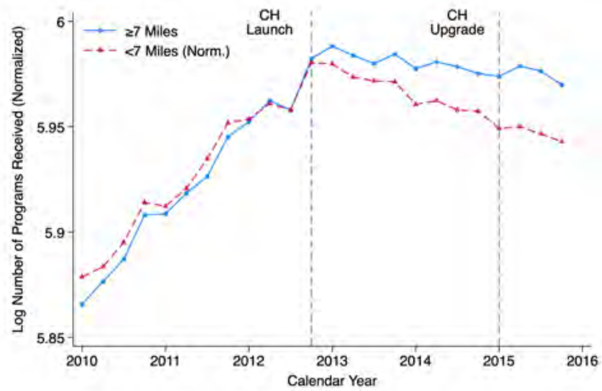

(c) Age 18–64

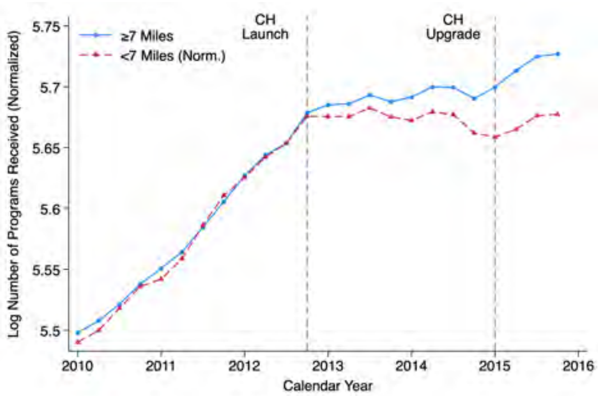

(d) Age 65+

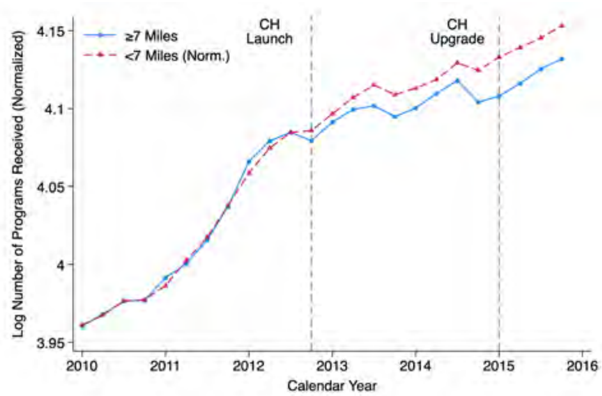

Notes: Each panel plots the log number of programs received in Virginia, by quarter, separately for zip codes located ≥ 7 miles from their county field office (blue solid line) and < 7 miles from their county field office (red dashed line, normalized to match the level of the ≥ 7 miles series at September 2012). Within each distance group and quarter, we take the unweighted average of programs received across zip codes and then take the log of that average. We present trends for the full population and separately by age subgroup. The vertical dashed lines mark CommonHelp's October 2012 launch and January 2015 upgrade.

Figure A.6: Event-Study Estimates, by More Granular Age and Household Type Subgroups

(a) Age 0–17: By Age Subgroups

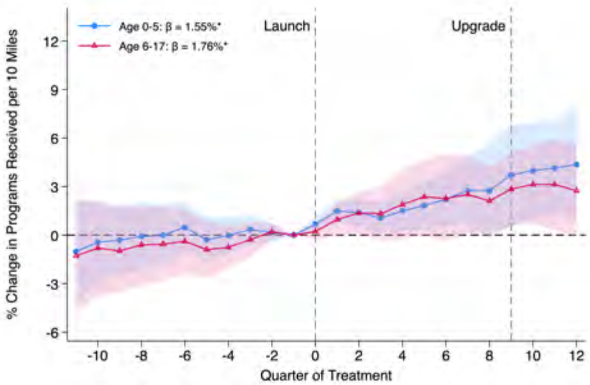

(b) Age 0–17: By HH Type Subgroups

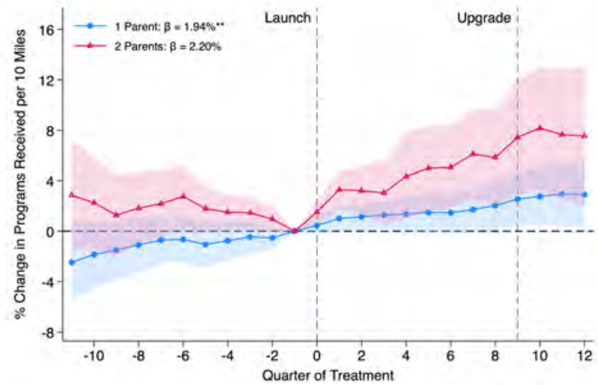

(c) Age 18–64: By Age Subgroups

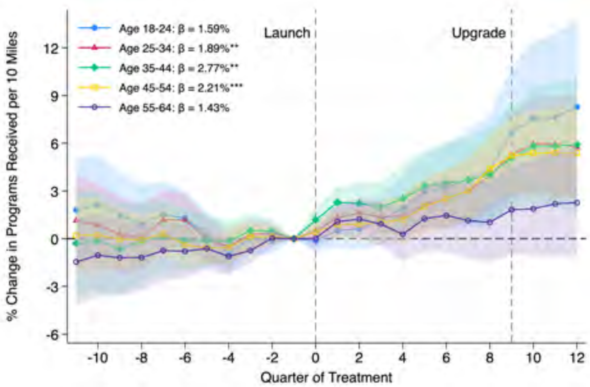

(d) Age 18–64: By HH Type Subgroups

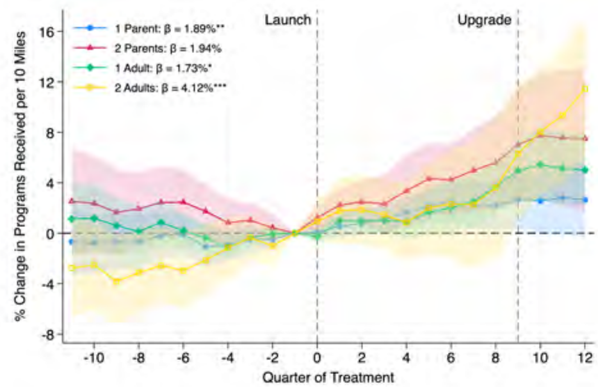

(e) Age 65+: By Age Subgroups

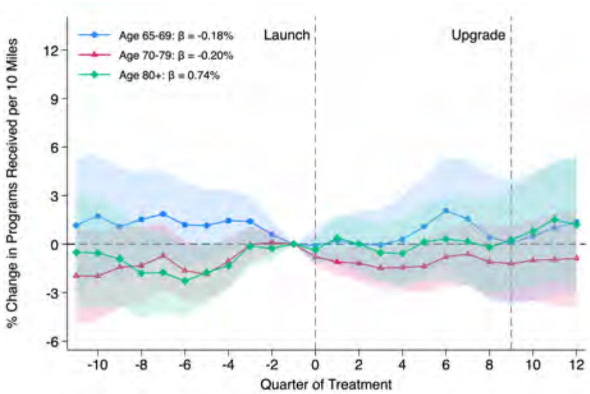

(f) Age 65+: By HH Type Subgroups

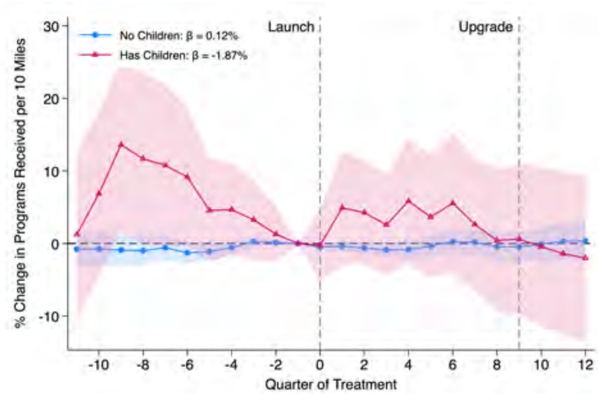

Notes: The figure plots the percent change per ten miles implied by β_k from equation (1) for ages 0–17, 18–64, and 65+ (and within each broad age bucket separately by granular age bins and by household type). Pooled estimates of β from equation (2) are reported in the legend on the same scale. See notes to Figure 2 for more details.

Figure A.7: CommonHelp's Impact on Enrollment Increases in Distance to Field Office

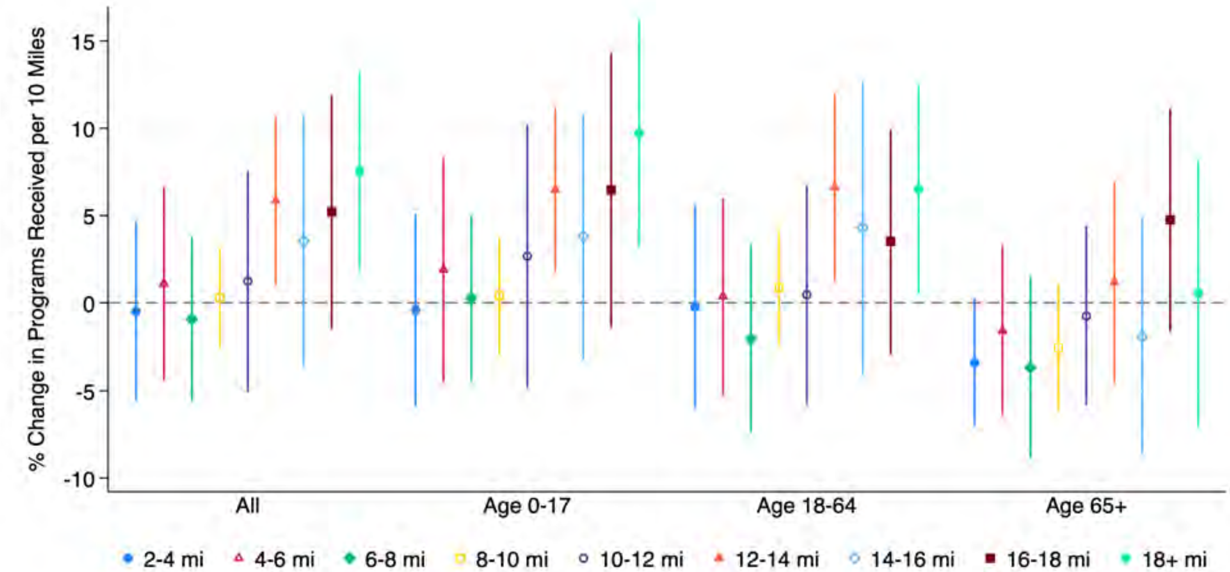

Notes: This figure reports estimates of β from equation (2) in which D_z is replaced with $\mathbb{1}\{D_z \in [d_0, d_1)\}$ that is estimated using zip codes with $D_z < 2$ or $D_z \in [d_0, d_1)$ for different choices of d_0, d_1. Estimates are reported as percent changes relative to zip codes within 2 miles of the field office, given by $100 \times [\exp(\beta) - 1]$, with standard errors obtained by the delta method. Vertical lines are 95% confidence intervals based on robust standard errors clustered at the county level. Zip codes are weighted by the average number of programs received between January 2010–September 2012.

Figure A.8: Functional Limitations Among Virginia Population by Age (2013)

(a) Full Population

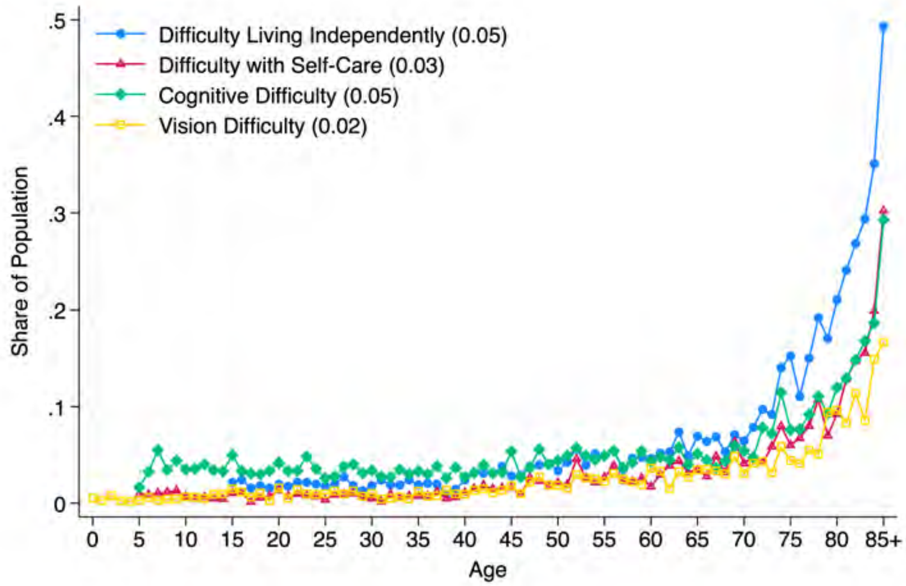

(b) Incomes <200% FPL

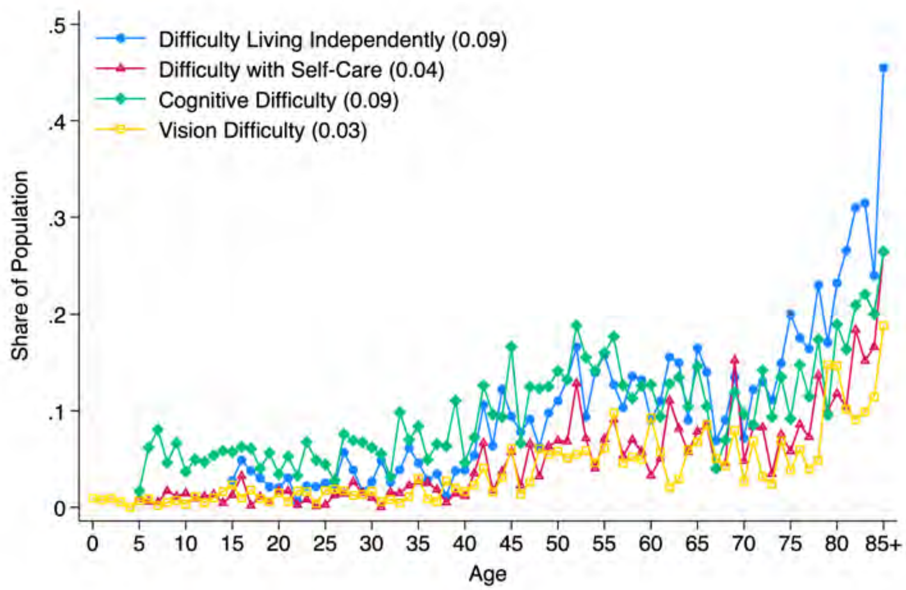

Notes: This figure reports the share of individuals with selected functional limitations by age using the 2013 American Community Survey (ACS) for Virginia. The functional limitations shown are difficulty living independently, difficulty with self-care, cognitive difficulty, and vision difficulty. Panel (a) reports estimates for the full population. Panel (b) restricts the sample to individuals with household incomes below 200% of the federal poverty level (FPL), who are more likely to be eligible for means-tested public programs. Population-weighted means are reported in the legend in parentheses.

Figure A.9: SNAP and Medicaid Receipt Around Eligibility for OASI and Medicare

(a) Around Age 62 (OASI)

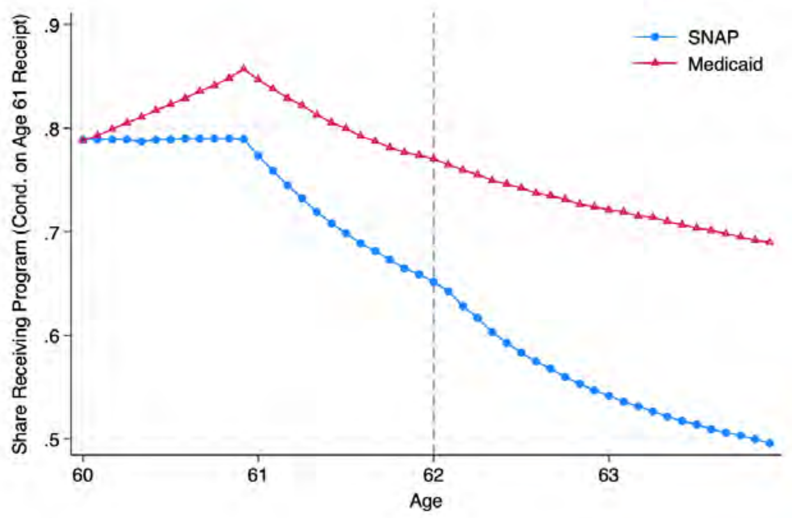

(b) Around Age 65 (Medicare)

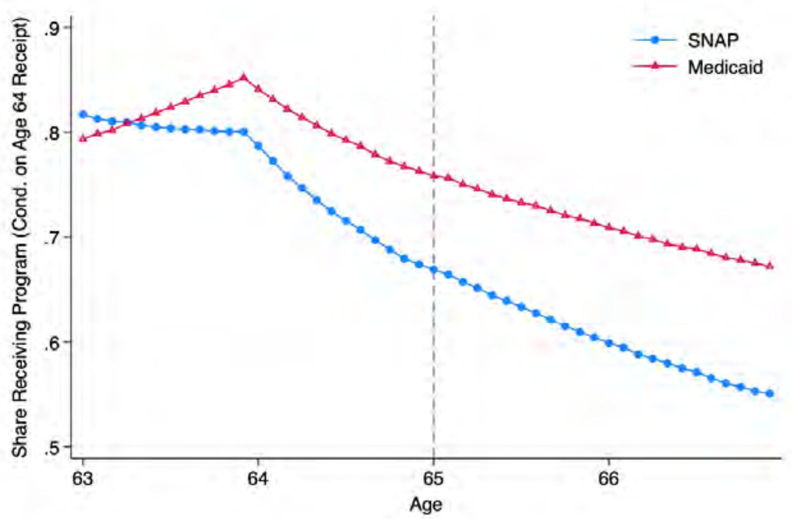

Notes: This figure descriptively examines whether SNAP and Medicaid receipt changes around age-based eligibility thresholds for potentially substitutable benefits. Panel (a) examines receipt around age 62, when individuals first become eligible to claim Old-Age and Survivors Insurance (OASI) benefits. Panel (b) examines receipt around age 65, when most individuals become eligible for Medicare. In each panel, the sample is restricted to individuals who received the corresponding VDSS program between 12 and 24 months before reaching the relevant age threshold. The figure reports the share of these individuals receiving each program by age. Vertical dashed lines denote the relevant eligibility thresholds. Data are from Virginia Department of Social Services (VDSS) administrative records.

Figure A.10: Internet Access Among Virginia Population by Age and Year

(a) Full Population

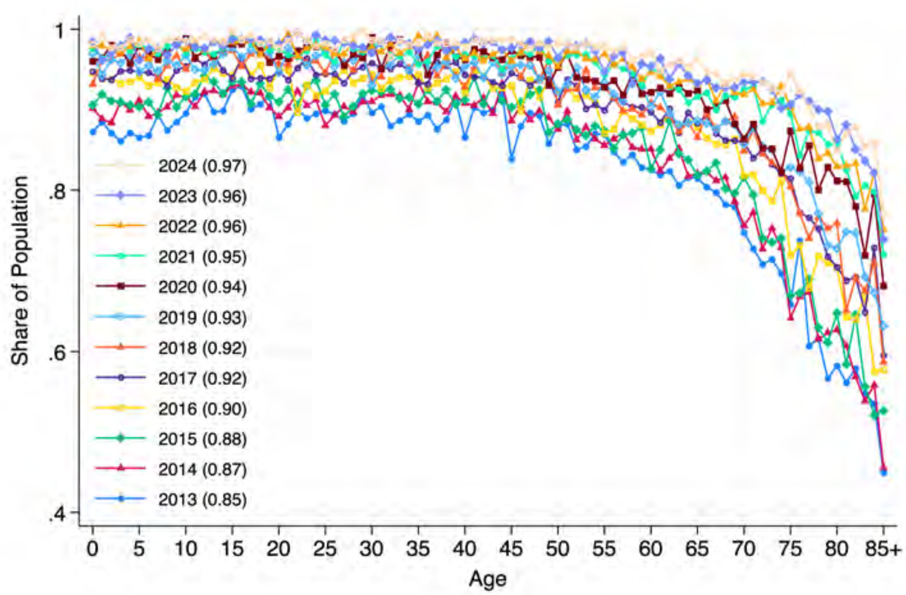

(b) Income <200% FPL

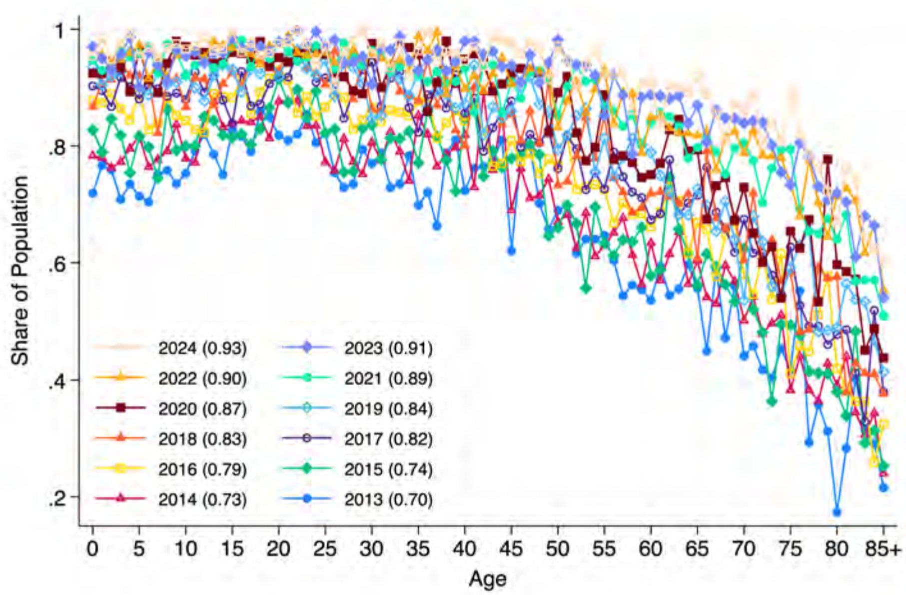

Notes: This figure reports the share of individuals with internet access by age and year using the 2013–2024 American Community Surveys (ACS) for Virginia. Panel (a) reports estimates for the full population. Panel (b) restricts the sample to individuals with household incomes below 200% of the federal poverty level (FPL), who are more likely to be eligible for means-tested public programs. Population-weighted means for each year are reported in the legend in parentheses.

Figure A.11: Internet Access Among Virginia vs U.S. Population by Age (2013 and 2019)

(a) Full Population

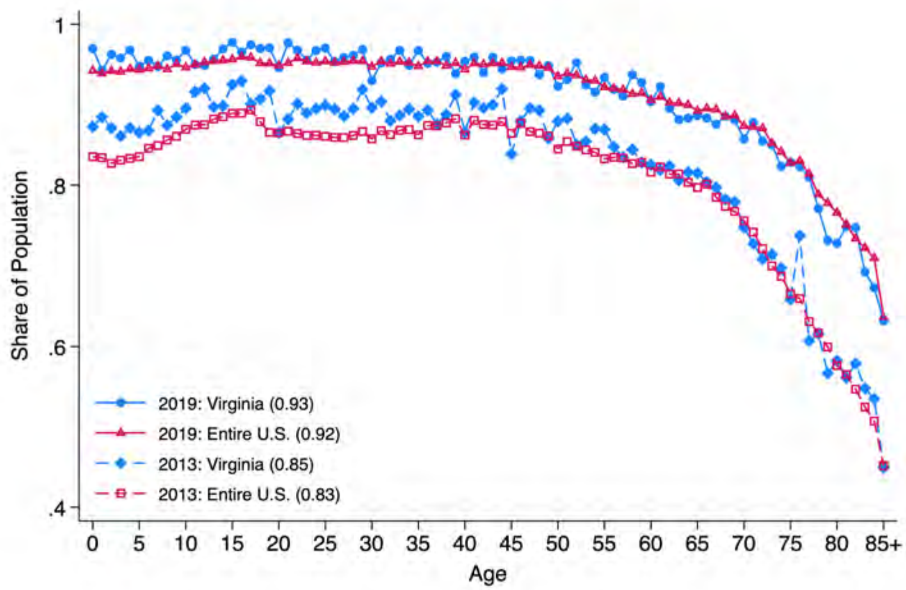

(b) Income <200% FPL

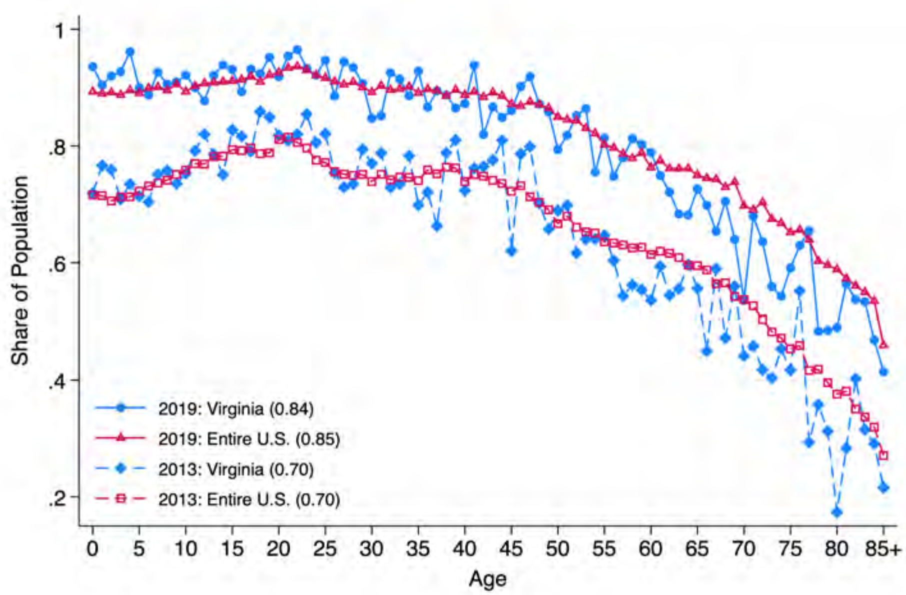

Notes: This figure reports the share of individuals with internet access by age using the 2013 and 2019 American Community Surveys (ACS) for Virginia versus the entire U.S. Panel (a) reports estimates for the full population. Panel (b) restricts the sample to individuals with family incomes below 200% of the federal poverty level (FPL), who are more likely to be eligible for means-tested public programs. Population-weighted means for each year are reported in the legend in parentheses.

Figure A.12: Smartphone Ownership Among Virginia Population by Age (2016)

(a) Full Population

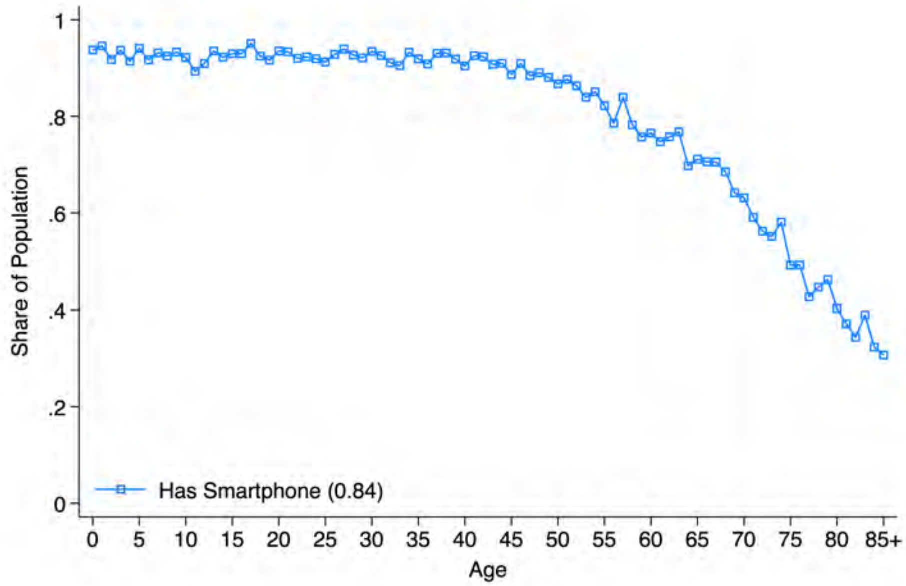

(b) Income <200% FPL

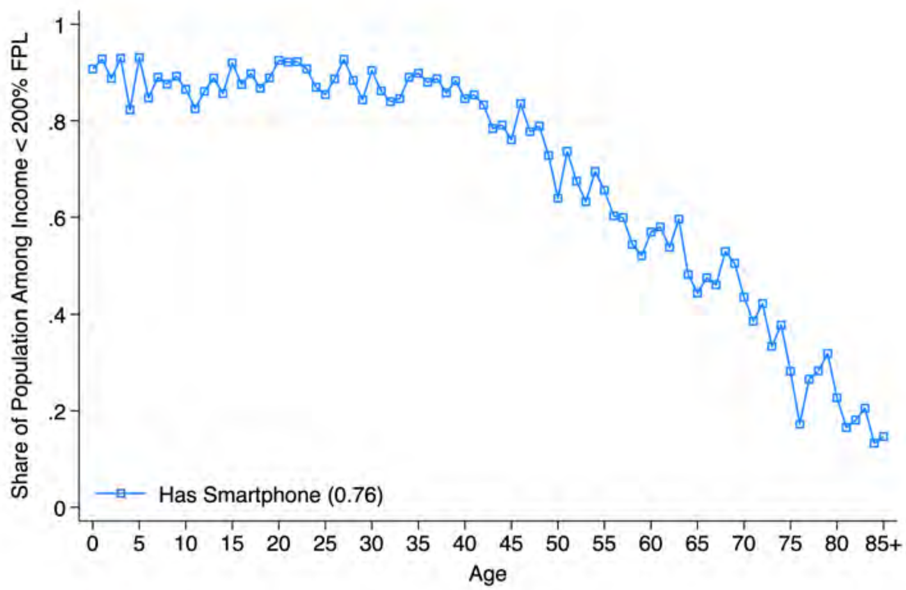

Notes: This figure reports the share of individuals owning a smartphone by age using the 2016 American Community Surveys (ACS) for Virginia. This is the first year in which the ACS asks about smartphone ownership. Panel (a) reports estimates for the full population. Panel (b) restricts the sample to individuals with family incomes below 200% of the federal poverty level (FPL), who are more likely to be eligible for means-tested public programs. Population-weighted means are reported in the legend in parentheses.

Table A.1: Heterogeneity in Effects by Age

	Programs Received (% Change per 10 Miles)		
	Pooled Estimates (1)	Launch Estimates (2)	Upgrade Estimates (3)
All	1.582**	1.086*	2.316***
	(0.725)	(0.649)	(0.704)
Ages 0–17	1.691**	1.321*	1.735**
	(0.786)	(0.748)	(0.800)
Ages 0–5	1.546*	1.087	2.244*
	(0.941)	(0.810)	(1.197)
Ages 6–17	1.759*	1.455	1.394**
	(1.047)	(1.025)	(0.642)
Ages 18–64	1.993**	1.261*	3.404***
	(0.838)	(0.725)	(0.890)
Ages 18–24	1.591	0.561	5.173***
	(1.123)	(0.970)	(1.483)
Ages 25–34	1.888**	1.086	3.739***
	(0.942)	(0.845)	(0.967)
Ages 35–44	2.774**	2.121*	2.906***
	(1.317)	(1.165)	(1.016)
Ages 45–54	2.208***	1.445*	3.581***
	(0.836)	(0.763)	(0.896)
Ages 55–64	1.431	1.137	1.298
	(1.141)	(1.021)	(0.940)
Ages 65+	0.078	-0.125	0.898
	(0.982)	(0.875)	(0.651)
Ages 65–69	-0.178	-0.382	0.914
	(1.572)	(1.451)	(0.988)
Ages 70–79	-0.197	-0.337	0.624
	(0.968)	(0.877)	(0.900)
Ages 80+	0.741	0.398	1.491
	(1.611)	(1.463)	(0.939)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$, with standard errors obtained by the delta method. The pooled column reports the overall pre-post difference-in-differences estimate. The launch column reports the effect associated with the initial CommonHelp launch, while the upgrade column reports the additional effect associated with the later system upgrade. Rows report estimates separately by age subgroup. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, county-specific linear time trends, and zip/time-varying covariates. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.2: Heterogeneity in Effects by Household Type

	Programs Received (% Change per 10 Miles)		
	Pooled Estimates (1)	Launch Estimates (2)	Upgrade Estimates (3)
All	1.582**	1.086*	2.316***
	(0.725)	(0.649)	(0.704)
By Presence of Children			
No Children	1.442*	0.745	3.353***
	(0.818)	(0.747)	(0.736)
Has Children	1.694**	1.273*	1.940**
	(0.806)	(0.751)	(0.771)
Ages 0–17			
1 Parent	1.943**	1.604*	1.619**
	(0.983)	(0.916)	(0.739)
2+ Parents	2.199	1.474	3.265***
	(1.584)	(1.419)	(1.118)
Ages 18–64			
1 Childless Adult	1.732*	0.964	3.755***
	(0.893)	(0.791)	(0.948)
2+ Childless Adults	4.116***	2.948*	6.069***
	(1.605)	(1.551)	(1.513)
1 Parent	1.890**	1.558*	1.495*
	(0.936)	(0.838)	(0.865)
2+ Parents	1.938	1.061	3.736***
	(1.550)	(1.351)	(1.333)
Ages 65+			
No Children	0.123	-0.097	0.973
	(0.998)	(0.892)	(0.656)
Has Children	-1.874	-1.363	-2.337
	(3.078)	(3.079)	(3.463)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$. The pooled column reports the overall pre-post difference-in-differences estimate. The launch column reports the effect associated with the initial CommonHelp launch, while the upgrade column reports the additional effect associated with the later system upgrade. Rows report estimates separately by age-by-household-type subgroup. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, county-specific linear time trends, and zip/time-varying covariates. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.3: Summary of Robustness Checks

Specification	Programs Received (% Change per 10 Miles) All (1)	Age 0–17 (2)	Age 18–64 (3)	Age 65+ (4)
Main Estimates	1.582**	1.691**	1.993**	0.078
	(0.725)	(0.786)	(0.838)	(0.982)
Exclude All Covariates	1.935**	2.097*	2.517**	-1.217
	(0.955)	(1.161)	(1.164)	(0.930)
Add Employment Controls	1.308*	1.431*	1.673**	-0.338
	(0.723)	(0.763)	(0.819)	(0.874)
$\log(Y+1)$	0.002***	0.002***	0.004***	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Inverse Hyperbolic Sine	0.003***	0.002***	0.004***	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Drop Zips >15 Miles	1.152	1.155	1.781	-0.925
	(1.002)	(1.058)	(1.171)	(1.241)
Drop Zips >20 Miles	1.708**	1.801**	2.170**	0.259
	(0.793)	(0.848)	(0.930)	(1.099)
Driving Distance	1.277**	1.417**	1.495**	0.193
	(0.530)	(0.564)	(0.619)	(0.730)
$\min\{\text{Driving}, \text{Transit}\}$ Distance	1.308**	1.460**	1.518**	0.218
	(0.539)	(0.576)	(0.627)	(0.731)

Notes: This table reports estimates of β from equation (2) corresponding to various specifications. Estimates from the Poisson specifications are reported as percent changes per ten miles, given by $100 \times [\exp(10\beta) - 1]$. Estimates from the $\log(Y+1)$ and inverse hyperbolic sine specifications are reported as 10β. Columns 1–4 report estimates for all individuals, children ages 0–17, adults ages 18–64, and adults ages 65 and older, respectively. Standard errors clustered at the county level are reported in parentheses. The main specification is estimated using Poisson pseudo-maximum likelihood and includes county-fixed effects, year-month fixed effects, county-specific linear time trends, and the baseline covariates. "Exclude All Covariates" removes the baseline covariates. "Add Employment Controls" adds county-level unemployment and labor-force non-participation rates (available only starting in 2011). "$\log(Y+1)$" and "IHS" estimate linear models using the log-transformed outcome or inverse hyperbolic sine, weighted by the number of programs received. "Drop Zips >15 Miles" and "Drop Zips >20 Miles" exclude ZIP codes more than 15 or 20 miles, respectively, from the nearest local department of social services office. "Driving Distance" replaces Haversine distance with driving distance, and "$\min\{\text{Driving}, \text{Transit}\}$ Distance" uses the minimum of driving and transit distance. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.4: Heterogeneity in Effects by County-Level Access to Broadband Internet

Broadband Access	Programs Received (% Change per 10 Miles)			
	All (1)	Age 0–17 (2)	Age 18–64 (3)	Age 65+ (4)
Lower Broadband Access	0.306 (1.140)	0.288 (1.278)	0.621 (1.122)	-1.261 (1.330)
Higher Broadband Access	2.268** (0.957)	2.406** (1.223)	2.856*** (1.084)	0.542 (1.363)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$, for subgroups defined by whether they live in a county with lower or higher broadband access. We define a higher-broadband county as having at least 20% of households with an internet download speed of at least 25 mbps (constituting slightly more than half of all counties). Columns 1–4 report estimates for all individuals, children ages 0–17, adults ages 18–64, and adults ages 65 and older, respectively. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, and county-specific linear time trends. We omit the zip/time-varying covariates in this table, given that the subgroups are themselves defined by a county-level characteristic that may absorb much of the variation these covariates capture. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.5: Heterogeneity in Elderly Effects by County-Level Access to Broadband Internet and Living Arrangement

	Programs Received (% Change per 10 Miles)	
Living Arrangement	Age 65+ (1)	Age 80+ (2)
High Broadband Access		
More Likely to Live with Others	-0.062	4.317
	(1.434)	(3.528)
Less Likely to Live with Others	1.653	-0.914
	(2.954)	(3.116)
Low Broadband Access		
More Likely to Live with Others	-1.116	-1.856
	(2.467)	(3.811)
Less Likely to Live with Others	-0.461	-0.268
	(1.101)	(2.978)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$. The upper and lower panels report estimates for counties with higher and lower broadband access, respectively. We define a higher-broadband county as one in which at least 20% of households have an internet download speed of at least 25 Mbps. Within each broadband category, counties are classified as more or less likely to have older adults living with others using a median split of the county-level share of residents ages 65 and older who live with others, as measured in the five-year ACS estimates for 2010–2014. These classifications therefore describe county-level living arrangements rather than an individual recipient's living arrangement. Columns 1 and 2 report estimates for adults ages 65 and older and ages 80 and older, respectively. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, and county-specific linear time trends. We omit the zip/time-varying covariates in this table, given that the subgroups are themselves defined by a county-level characteristic that may absorb much of the variation these covariates capture. Statistical significance is denoted by $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table A.6: Heterogeneity in Effects by Specific Program

Program	Programs Received (% Change per 10 Miles)			
	All	Age 0–17	Age 18–64	Age 65+
	(1)	(2)	(3)	(4)
SNAP	1.922**	1.973*	1.994**	0.621
	(0.911)	(1.020)	(0.954)	(0.978)
TANF	2.719	2.887	2.133	-12.360
	(1.945)	(1.990)	(1.982)	(42.902)
Medicaid	1.197	1.401*	1.920**	-0.244
	(0.741)	(0.797)	(0.957)	(1.080)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$, where the outcome is now a binary indicator for receipt of a particular program. Columns 1–4 report estimates for all individuals, children ages 0–17, adults ages 18–64, and adults ages 65 and older, respectively. We do not report TANF estimates for the elderly (aged 65+), given that essentially no elderly individuals receive TANF. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, county-specific linear time trends, and zip/time-varying covariates. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.7: Heterogeneity in Effects by Revealed Willingness-to-Pay (Age 65+ Subgroups)

	Programs Received (% Change per 10 Miles)	
Subgroup Classification	Low WTP (1)	High WTP (2)
Previous Connection to Programs		
Any Prior Program Participation	3.680*	-0.378
	(2.005)	(0.965)
Any Prior Program Participation (Lapsed)		1.722
		(5.048)
Multiple Prior Programs		-0.496
		(0.912)
Multiple Prior Programs (Lapsed)		-2.814
		(5.959)
Wages at Ages 55–61		
Wages Relative to 10th Percentile	2.185	-3.361
	(2.247)	(5.384)
Wages Relative to 25th Percentile	1.527	2.278
	(2.232)	(3.556)
Wages Relative to 50th Percentile	-0.529	3.717
	(2.350)	(2.931)
Criminal History		
Prior Criminal History	0.100	-1.655
	(0.973)	(5.521)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$, for subgroups of adults ages 65 and older classified as having relatively low or high revealed willingness to pay (WTP) for VDSS programs. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, county-specific linear time trends, and ZIP/time-varying covariates. The previous-program-connection classifications measure participation in VDSS programs during 2007–2011, prior to CommonHelp. Individuals with no prior program participation are classified as low WTP, while those with any prior program participation are classified as high WTP. Lapsed participation indicates that prior participants were no longer enrolled by September 2012. The wage-based classifications restrict the sample to adults ages 65–70 and measure their wages between ages 55 and 61. Individuals with wages below the indicated percentile are classified as high WTP, while those with wages at or above the indicated percentile are classified as low WTP. Individuals with no observed wages between ages 55 and 61 are excluded from these comparisons. The criminal-history classification defines individuals with recorded criminal involvement in 2011 or earlier as high WTP and those without recorded criminal involvement as low WTP. Blank cells indicate that no corresponding subgroup estimate is reported. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.8: Heterogeneity in Effects by Documentation Burdens (Pooled Estimates)

	Programs Received (% Change per 10 Miles)		
Documentation Burden	All (1)	Age 18–64 (2)	Age 65+ (3)
Wage Volatility			
No Wages, Last 4 Quarters	1.553**	1.872**	0.104
	(0.781)	(0.857)	(1.014)
Low-Volatility Employment	1.370	2.035*	-1.630
	(1.046)	(1.120)	(4.015)
High-Volatility Employment	1.833	2.084*	0.212
	(1.271)	(1.169)	(3.784)
Number of Jobs			
No Wages, Last Quarter	1.716**	2.033**	0.110
	(0.758)	(0.828)	(0.984)
Single Job, Last Quarter	1.406	1.944*	-2.725
	(1.145)	(1.157)	(3.515)
Multiple Jobs, Last Quarter	1.058	1.242	10.439
	(1.703)	(1.654)	(9.718)
Employment and Job Count Change			
Not Employed, No Job-Count Change	1.545**	1.862**	0.110
	(0.780)	(0.856)	(1.016)
Not Employed, Job-Count Change	2.808**	2.996***	0.640
	(1.221)	(1.129)	(4.578)
Employed, No Job-Count Change	0.667	1.662	-1.411
	(1.359)	(1.317)	(4.797)
Employed, Job-Count Change	1.641	1.800	-1.295
	(1.300)	(1.291)	(4.136)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$, for subgroups defined using measures of potential documentation burdens. Columns report pooled estimates for all individuals, individuals ages 18–64, and individuals ages 65 and older. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, county-specific linear time trends, and ZIP/time-varying covariates. Wage-volatility categories are constructed using household earnings observed in unemployment-insurance wage records during the preceding four quarters. Individuals classified as having no wages have no observed earnings during those quarters. Among individuals with observed earnings, low- and high-volatility employment indicate a coefficient of variation in quarterly earnings below or above the sample average, respectively. Number-of-jobs categories are based on the number of jobs observed during the preceding quarter. The multiple-jobs estimate for individuals ages 65 and older is omitted because almost no individuals in this age group hold multiple jobs. The final panel interacts current employment status with whether the individual's number of jobs changed during the preceding 12 months. A job-count change includes transitions such as zero to one job, one to two jobs, or the reverse. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A.9: Heterogeneity in Effects by Documentation Burdens (Launch vs Upgrade)

	Programs Received (% Change per 10 Miles)					
	Launch Estimates			Upgrade Estimates		
Documentation Burden	All (1)	Age 18–64 (2)	Age 65+ (3)	All (4)	Age 18–64 (5)	Age 65+ (6)
Wage Volatility						
No Wages, Last 4 Quarters	0.962	1.092	-0.113	2.794***	3.721***	0.961
	(0.677)	(0.737)	(0.911)	(0.880)	(0.958)	(0.641)
Low-Volatility Employment	1.164	1.559	-1.297	0.942	2.107*	-1.562
	(0.944)	(0.957)	(4.037)	(0.967)	(1.257)	(4.144)
High-Volatility Employment	1.172	1.158	0.338	3.036**	4.281***	-0.578
	(1.193)	(1.087)	(3.707)	(1.199)	(1.270)	(3.597)
Number of Jobs						
No Wages, Last Quarter	1.127*	1.241*	-0.096	2.796***	3.800***	0.911
	(0.676)	(0.724)	(0.879)	(0.838)	(0.918)	(0.654)
Single Job, Last Quarter	0.916	1.194	-2.531	2.266***	3.406***	-0.914
	(1.042)	(1.015)	(3.394)	(0.861)	(1.084)	(3.132)
Multiple Jobs, Last Quarter	0.800	0.620	9.334	1.095	2.590	4.194
	(1.431)	(1.413)	(9.384)	(1.744)	(1.831)	(10.812)
Employment and Job Count Change						
Not Employed, No Job-Count Change	0.955	1.083	-0.106	2.787***	3.715***	0.953
	(0.677)	(0.736)	(0.912)	(0.883)	(0.959)	(0.640)
Not Employed, Job-Count Change	2.201*	2.126*	0.992	2.979***	4.329***	-1.659
	(1.221)	(1.116)	(4.615)	(1.103)	(1.194)	(4.758)
Employed, No Job-Count Change	0.276	1.076	-1.425	1.775*	2.553*	0.066
	(1.276)	(1.174)	(4.813)	(1.077)	(1.416)	(3.859)
Employed, Job-Count Change	1.213	1.050	-1.365	1.924	3.351**	0.332
	(1.092)	(1.097)	(4.012)	(1.360)	(1.381)	(4.176)

Notes: This table reports estimates of β from equation (2), expressed as percent changes per ten miles and given by $100 \times [\exp(10\beta) - 1]$, for subgroups defined using measures of potential documentation burdens. The launch columns report the effect associated with the initial CommonHelp launch, while the upgrade columns report the additional effect associated with the later system upgrade. Within each set of columns, estimates are reported for all individuals, individuals ages 18–64, and individuals ages 65 and older. Standard errors clustered at the county level are reported in parentheses. All specifications include county fixed effects, year-month fixed effects, county-specific linear time trends, and ZIP/time-varying covariates. Wage-volatility categories are constructed using household earnings observed in unemployment-insurance wage records during the preceding four quarters. Individuals classified as having no wages have no observed earnings during those quarters. Among individuals with observed earnings, low- and high-volatility employment indicate a coefficient of variation in quarterly earnings below or above the sample average, respectively. Number-of-jobs categories are based on the number of jobs observed during the preceding quarter. The final panel interacts current employment status with whether the individual's number of jobs changed during the preceding 12 months. A job-count change includes transitions such as zero to one job, one to two jobs, or the reverse. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.