

Take-up and Targeting: Experimental Evidence from SNAP

Amy Finkelstein and Matthew J. Notowidigdo*

May 2018

Abstract

This paper develops a framework for evaluating the welfare impact of various interventions designed to increase take-up of social safety net programs in the presence of potential behavioral biases. We calibrate the key parameters using a randomized field experiment in which 30,000 elderly individuals not enrolled in – but likely eligible for – the Supplemental Nutrition Assistance Program (SNAP) are either provided with information that they are likely eligible, provided with this information and also offered assistance in applying, or are in a “status quo” control group. Only 6 percent of the control group enrolls in SNAP over the next 9 months, compared to 11 percent of the Information Only group and 18 percent of the Information Plus Assistance group. The individuals who apply or enroll in response to either intervention receive lower benefits and are less sick than the average enrollee in the control group. The results are consistent with the existence of optimization frictions that are greater for needier individuals, suggesting that the poor targeting properties of the interventions reduce their welfare gains.

JEL codes: C93; H53; I38

Keywords: SNAP; Food Stamps; Take-Up; Targeting; Welfare.

*MIT and Northwestern (respectively), National Bureau of Economic Research, and J-PAL North America. We are grateful to Martin Aragonese, Aileen Devlin, Carolyn Stein, John Tebes and Ting Wang for excellent research assistance and to Laura Feeney for superb research management. We thank our excellent partners at Benefits Data Trust, and particularly Rachel Cahill and Matt Stevens who worked tirelessly and patiently to address our innumerable requests and questions. We thank Abhijit Banerjee, Stefano DellaVigna, Manasi Deshpande, Paul Goldsmith-Pinkham, Colin Gray, Nathan Hendren, Henrik Kleven, Kory Kroft, Elira Kuka, Ben Olken, Jesse Shapiro, Chris Udry, and numerous seminar participants for helpful comments. The experiment reported in this study is listed in the AEA RCT Registry (#0000902). We gratefully acknowledge financial support from the Alfred P. Sloan Foundation.

1 Introduction

Enrollment in U.S. social safety net programs is not automatic: individuals must apply and demonstrate eligibility. Often, eligibility rules are complicated, application forms long, and documentation requirements substantial. Perhaps as a result, incomplete take-up is pervasive (Currie, 2006). The main explanations typically offered are: lack of knowledge about eligibility and transaction costs associated with enrollment.¹

Various public policy efforts try to increase take-up by increasing awareness of eligibility and simplifying application processes. For example, for the Supplemental Assistance Nutrition Program (SNAP), New York City Mayor Bill de Blasio proposed an enrollment campaign that contacted Medicare recipients about their SNAP eligibility and improved online services (Hu, 2014), the state of Texas simplified the application process (Aaronson, 2011), and Congress provided funding to study various models for facilitating access to SNAP among the elderly (Kauff et al., 2014).

However, incomplete take-up may represent a constrained social optimum. For many programs, eligibility is not measurable in passively-collected administrative data. SNAP, Social Security Disability Insurance (SSDI), Supplemental Security Income (SSI), and Medicaid are some examples. As a result, eligibility cannot be assessed without the government incurring administrative costs; these represent a negative externality which can drive a wedge between privately optimal and socially optimal application decisions (e.g., Kleven and Kopczuk 2011). In addition, neoclassical theory has emphasized that informational or transactional barriers to enrollment may serve as useful screens, allowing a given amount of public spending to be directed to individuals with higher marginal utility from enrollment (or higher social welfare weights); these theories suggest a trade-off between productive efficiency and allocative efficiency (e.g., Nichols et al., 1971, Nichols and Zeckhauser 1982, Besley and Coate 1992).

By contrast, recent work in behavioral economics has conjectured that these barriers may have exactly the opposite targeting effect to what the neoclassical theory assumes, and discourage the applicants with the highest marginal utility of enrollment (e.g., Bertrand et al. 2004, Mani et al. 2013, Mullainathan and Shafir 2013). For example, in their book “Scarcity”, Mullainathan and Shafir (2013) argue that poverty imposes a “bandwidth tax” that makes poor individuals more likely to fail to undertake high-net-value activities, such as taking chronic disease medication, paying regular bills on time, or enrolling in a public benefit program for which one is eligible. Ultimately, the targeting properties of these barriers and their welfare implications are empirical questions.

This paper formalizes an approach to analyzing the normative consequences of how information and transaction costs affect the number and type of people who apply and enroll in social safety net programs, and applies it to a randomized evaluation of interventions aimed at elderly non-participants in the SNAP program. SNAP - also known as food stamps - is one of the most important social safety net programs in the United States. It is the only benefit that is virtually universally available to low-income households. During the Great Recession, as many as one in

¹A third common explanation - stigma associated with program participation - can be modeled as a form of a transaction cost (Moffitt 1983; Currie 2006).

seven individuals received SNAP (Ganong and Liebman, 2013). In 2015, public expenditures on SNAP were about \$70 billion, roughly the same amount as the Earned Income Tax Credit (EITC) and higher than the \$60 billion spent on SSI or the \$30 billion on cash welfare (TANF).² Take-up of SNAP is especially low among the elderly, who are the focus of our study; in 2012, only 42 percent of eligible elderly enrolled in SNAP, compared to 83 percent overall (Eslami, 2014). And the stakes associated with non-participation are non-trivial for the elderly; average annual SNAP benefits are about \$1,500, or about 15 percent of household income among the eligible (Center on Budget and Policy Priorities, 2017).

To explore barriers to enrollment and the types of individuals deterred by these barriers, we partnered with Benefits Data Trust (BDT), a national not-for-profit organization committed to transforming how individuals in need access public benefits. BDT in turn has partnered with the Pennsylvania state government to receive administrative data on applications and enrollment in a variety of public benefits, including SNAP and Medicaid. BDT uses linked data on SNAP enrollment and on Medicaid enrollment among elderly individuals (ages 60 and over) to identify individuals not enrolled in but likely eligible for SNAP. These individuals form the basis of our study population.

We selected about 30,000 elderly individuals from this population in late 2015 and randomized them into three equally-sized groups: an Information Only treatment, an Information Plus Assistance treatment, and a status quo control group. The treatments took place in the first half of 2016. Study participants in the Information Only treatment received a mailing - and a follow-up reminder postcard - from the Secretary of Pennsylvania's Department of Human Services (DHS), informing them of their likely eligibility for SNAP, and providing them a phone number at DHS to call to apply. Study participants in the Information Plus Assistance arm received a virtually identical letter and reminder postcard, with one key change: they were provided with a phone number at the PA Benefits Center (the local name of BDT) to call to apply. Callers in this arm received phone-based application assistance from one of BDT's Benefits Outreach Specialists; these BDT employees asked a series of questions that allowed them to inform the caller of their potential eligibility and likely benefit amount, to fill out the application, to assist the applicant in collecting necessary verification documents, to submit the application, and to assist with any follow-up questions that arose from DHS. Both intervention arms included sub-treatments that varied the content of the letter and, in one case, whether or not the reminder postcard was sent; we describe these in more detail below, although we focus primarily on the main treatments. We tracked calls from study participants to both BDT and DHS, and received administrative data from DHS on SNAP applications, enrollments, and benefit amounts after the intervention; we obtained additional demographic and health data pre-intervention from the study participants' Medicaid records.

The experiment itself yields two main empirical findings. First, information alone increases enrollment, while information combined with assistance increases enrollment even more, but at a

²US Department of Agriculture 2016, US Department of Health and Human Services 2016, US Internal Revenue Service 2015, US Social Security Administration (2016)

higher cost per enrollee. Nine months after the intervention - at which point the initial impact appears to be fully in place - enrollment is 6 percentage points in the control arm compared to 11 percentage points in the Information Only arm and 18 percentage points in the Information Plus Assistance arm; these enrollment rates are all statistically distinguishable ($p < 0.001$). A rough calculation suggests the intervention cost per additional enrollee is lower in the Information Only treatment: about \$20 per enrollee compared to about \$60 per enrollee in the Information Plus Assistance treatment. We also find that a sub-treatment of the Information Only intervention, which omits the reminder postcard, reduces its impact by about 20 percent; this suggests a role for inattention in explaining at least some of the impact of the Information Only intervention.

We observe intervention effects at several intermediate stages. About 30 percent of the participants in each intervention arm call in response to the outreach materials, suggesting a likely ceiling for the impact of the interventions on enrollment. Similar call-in rates in the two interventions also suggest that the larger enrollment effects of Information plus Assistance relative to Information Only are likely due to the assistance per se, rather than the anticipation of assistance. Each intervention increases applications proportionally to its effect on enrollment; the success rate of applications is about 75 percent in all three arms.

Our second main empirical finding is that the interventions decrease targeting. Seemingly contrary to the “behavioral” hypothesis that information barriers and transaction costs deter the neediest eligible individuals, we find that marginal applicants and enrollees in either intervention are less needy than average applicants or enrollees in the control group. They receive lower benefits from the progressive benefit formula (if they enroll) and are less sick (as measured by pre-intervention rates of hospital visits and chronic diseases). Additionally, they are more likely to be white and more likely to have English as their primary language, suggesting that they may be less socioeconomically disadvantaged than the control group applicants and enrollees. These targeting results are broadly similar across the intervention arms; the characteristics of the marginal applicants and enrollees across intervention arms are statistically and substantively indistinguishable.

We develop a simple model that allows us to assess the normative implications of interventions that inform individuals about their likely eligibility (“information interventions”) or reduce the private costs of applying (“assistance interventions”). Beyond our particular study, such interventions are common in enacted and proposed policies (as described above), as well as in previous academic studies (e.g., Alatas et al. 2016, Bettinger et al. 2012, Bhargava and Manoli 2015). However, their welfare properties have not been examined in detail. In particular, analyses of targeting have focused primarily on observable characteristics of affected individuals that are related to poverty but, as we show, are not necessarily related to the welfare properties of targeting.

A key feature of the model is a fiscal externality on the government from the program, which creates the standard wedge between private and socially optimal application choices; we model this as the public costs of processing each application and paying benefits for the marginal enrollees. In addition, we allow for the possibility that individuals may not make privately optimal application decisions; specifically, we assume they may misperceive the expected benefits from applying.

Without such misperceptions, interventions that increase applications have no first-order effect on marginal enrollees' private welfare (by the envelope theorem); the information intervention will strictly decrease social welfare due to the fiscal externality if it increases applications, while the assistance intervention's effect on social welfare is *a priori* ambiguous. Moreover, in this "neoclassical benchmark" with no misperceptions, the targeting properties of the interventions are irrelevant for private welfare; and contrary to the "folk wisdom" that targeting poorer individuals is desirable, interventions that improve targeting to poorer individuals will actually *decrease* social welfare if the fiscal externality from the intervention is larger for poorer individuals (as would be the case, for example, with a progressive benefit formula and public costs of processing applications that are uniform across applicants).

If, however, individuals under-estimate expected benefits from applying - which is presumably the premise of an information intervention - then marginal applicants who apply as a result of either intervention have positive private welfare gains, with these gains increasing in the size of the misperception. A sufficient condition for interventions that increase targeting to be more likely to increase private welfare is that the under-estimation of expected benefits is (weakly) greater in percentage terms for the targeted individuals. However, even in this case, the social welfare impacts of increased targeting remains ambiguous. Indeed, we show that the targeting properties of the intervention have no general relationship to its social welfare impact; analysis of social welfare requires information not only on how misperceptions vary across individuals, but also how the size of the fiscal externality varies across these individuals.

We use the empirical findings from the experiment to calibrate the model. The evidence in our setting is consistent with the "behavioral" hypothesis that individuals underestimate their expected benefits from applying. This suggests potential private and social welfare gains from each intervention. Our estimates also suggest that underestimation of expected benefits is greater for needier individuals, again consistent with leading behavioral theories (e.g., Bertrand et al. 2004, Mani et al. 2013, Mullainathan and Shafir 2013). However, in contrast to these models and consistent with neoclassical theory (e.g., Nichols et al., 1971, Nichols and Zeckhauser 1982, Besley and Coate 1992), we find that our interventions to reduce transaction costs or improve information target *less* needy individuals. This bodes poorly for their welfare effects. Indeed, our calibrated model suggests that if - counterfactually - our intervention had been perfectly targeting, the social welfare benefits would have been substantially higher. Despite the poor targeting properties of the interventions, however, our rough, back-of-the-envelope calculation suggests that they may still be a cost-effective way of redistributing to low-income individuals relative to other safety net programs.

The rest of the paper proceeds as follows. Section 2 briefly describes the related literature. Section 3 presents our theoretical framework. Section 4 provides background information on the SNAP program and the application and enrollment process. Section 5 describes our experimental design and data sources. Section 6 presents the experimental results. Section 7 uses the results to calibrate the model from Section 3 and perform welfare analysis of the interventions. The last section concludes.

2 Related Literature

Our paper relates to two strands of literature: analysis of barriers to take-up and analysis of how barriers to take-up affect the *characteristics* of applicants and enrollees. Studies of barriers to take-up have, with the exception of Bettinger et al. (2012), focused on either informational barriers or transaction costs, rather than analyzing them together as we do here.³ This literature has also focused primarily on the descriptive, with little of the normative analysis we add here.

Barriers to Take-up

Reductions in informational barriers have been found to be quantitatively important in generating take-up in some contexts but not others. In a recent series of randomized interventions aimed at increasing take-up of the EITC among likely eligible individuals, Day Manoli and co-authors have found that take-up is highly sensitive to both the frequency and nature of reminder letters sent by the IRS, although the effects of the reminder do not persist into the following year when the individuals would have to sign up again (Bhargava and Manoli 2015, Manoli and Turner 2014, Guyton et al. 2016). Quasi-experimental studies have also found that information is an important barrier to take up of SSDI (Armour, forthcoming) and post-secondary enrollment among unemployment insurance recipients (Barr and Turner, forthcoming). Several of these studies conclude, as we do, that the results are consistent with misperceptions by individuals (see, e.g., Bhargava and Manoli 2015, Armour forthcoming). However, Alcott and Greenstone’s (2017) randomized evaluation finds that informational interventions do not affect take-up of home energy efficiency audits, concluding that lack of awareness is not a contributor to low take-up; likewise, Bettinger et al.’s (2012) randomized evaluation finds that providing low-income families with information about financial aid eligibility and nearby colleges had no effect on applications to college.

In the SNAP context, a survey of likely eligible SNAP non-participants found that about half reported that they were not aware of their eligibility (Bartlett et al. 2004). And in an early and innovative small randomized trial in 1993 in Pennsylvania, Daponte et al (1999) found suggestive evidence that informing non-participating, eligible households about their SNAP eligibility affected SNAP applications; however, small sample sizes (32 households were in the treatment arm and responded to the follow-up survey) as well as loss to follow-up made definitive conclusions difficult.

Reductions in transactional barriers have been found to be important for increasing enrollment in several different programs. Bettinger et al. (2012) found that while information alone was ineffective, combining information with assistance in completing a streamlined application process increased aid applications and ultimately college attendance and persistence by low-income individuals. Our findings in the SNAP context suggest, by contrast, that information alone can have

³The literature has paid comparatively less attention to the role of stigma, but the limited evidence does not point to a large role for stigma (Currie 2006). Recent efforts at “stigma” interventions have proven less successful at increasing take-up than informational interventions such as reminders or simplification (Bhargava and Manoli 2015). In the specific context of SNAP, Currie (2003) describes several pieces of survey evidence consistent with both lack of awareness and transaction costs in reducing SNAP take-up, but concludes that it does not appear from the existing survey evidence that stigma is a major deterrent to SNAP enrollment.

an effect, but that pairing it with assistance doubles the impact. In addition, Deshpande and Li (2017) find that the closing of local field offices where SSDI and SSI applications can be submitted substantially reduces both applications and enrollment, and Rossin-Slater (2013) finds that openings and closings of Women, Infants and Children (WIC) local program office affect program participation. Alatas et al. (2016) present evidence from a randomized evaluation across Indonesian villages that increasing the transaction cost of applying for a conditional cash transfer program reduces enrollment. At the extreme of reducing transaction barriers, defaulting to enrollment has been found to have substantial effects on outcomes such as participation in tax-subsidized 401(k) savings plans (Madrian and Shea 2001).

In the SNAP context, Schanzenbach (2009) provides evidence from one California county of a randomized evaluation in which a low-income tax preparer provided assistance to likely SNAP-eligible individuals. It found that, among those who expressed interest in learning more about SNAP, those in offices randomized into full assistance (in which the tax preparer went through a detailed interview with the client and then filled out and filed the application on the client’s behalf), were more likely to file an application than those who received help filling out the application but had to file it themselves, or those who only received a blank application (which might be viewed as analogous to our “Information Only” intervention).

Targeting

Our paper also relates to a second strand of the literature that investigates how barriers to enrollment affect the *characteristics* of applicants and enrollees. The existing “targeting” literature has been primarily descriptive, focusing on the observable characteristics of individuals affected by different barriers. Our theoretical framework, however, suggests that there is no general relationship between this targeting on observables and the impact of the intervention on either private or social welfare. We provide additional conditions that need to be examined empirically in order for an intervention’s targeting properties to yield normative implications.

To our knowledge, our study is the first to examine the targeting properties of both an information intervention and an assistance intervention. From prior information interventions, there is evidence that complexity disproportionately deters EITC enrollment of lower income potential recipients (Bhargava and Manoli 2015), and that lower income employees are more likely to choose dominated health insurance plans, due at least in part to a lack of insurance literacy (Bhargava et al., 2017). Our findings, by contrast, suggest that information about eligibility disproportionately encourages enrollment among less needy applicants.

Prior studies have tended to find that transaction costs increase targeting on some but not all dimensions. Alatas et al. (2016) find that introducing transaction costs by requiring individuals to apply for a conditional cash transfer in Indonesian villages rather than have the government automatically screen the individuals for eligibility improves targeting; specifically, it results in substantially poorer enrollees. However, marginally increasing the transaction costs does not further affect the characteristics of enrollees. Deshpande and Li (2017) find that increasing transaction costs

in U.S. disability programs (SSDI and SSI) worsens targeting among applicants - by increasing the share of applicants with only moderately severe disabilities - but increases targeting among enrollees, decreasing the share of enrollees with the least severe disabilities (conditional on being severe enough to be eligible); however they also find that the increased transaction costs reduce the share of enrollees with low education levels and low pre-application earnings, suggesting a reduction in targeting. In our context, by contrast, we find that reducing transaction costs decrease targeting on all dimensions, and at all stages (application and enrollment).

3 Conceptual framework

We develop a simple model for analyzing effects of interventions that provide information and/or application assistance. The model is useful for informing the welfare consequences of the impacts of the interventions, and also for providing sufficient conditions for when a finding that an intervention increases (or decreases) “targeting” translates into higher (or lower) welfare benefits from the intervention.

We analyze a policy whose purpose is redistribution. The two key assumptions of the model are that both the potential applicant and the state are uncertain about eligibility and benefit amounts unless the individual actively applies with the required information, and that the application - regardless of outcome - imposes costs on both the applicant and the government. These assumptions correspond to many U.S. transfer programs, such as SSDI, SSI, Medicaid and - as we discuss in more detail in Section 4 - the SNAP program we study empirically.

3.1 Set-up

Individuals decide whether to apply for benefits in a safety net program. For simplicity, we assume there are three possible outcomes from applying for the safety net program: (1) enrolling in the program and receiving a high level of benefits from program, b_H , (2) enrolling in the program and receiving a low level of benefits, $b_L (< b_H)$, and (3) being rejected after applying. We assume that there is ex-ante uncertainty about benefit eligibility both on the part of the potential applicant and on the part of the government.

For the individuals, we assume for simplicity that there are two types ($j \in \{h, l\}$) in the population who are potentially affected by the treatment. Individuals of type j know their income y_j for sure. Note that we use lower case j to refer to individual types and upper case J to denote benefit amounts, so an individual of type j can either receive benefits b_H or b_L if his application is accepted. We denote expected benefits, if the individual applies and is accepted, by $B_j = b(y_j)$. Uncertainty about benefits can come from several sources: uncertainty about the inputs into the benefit function, uncertainty about the function $b(\cdot)$ that maps inputs to benefits, and the potential for unexpected mistakes in the application process.

We parameterize the uncertainty about benefits for each type j as follows: type h individuals receive b_H with probability π_{hH} and b_L with probability π_{hL} , while type l individuals receive b_H

with probability π_{lH} and receive b_L with probability π_{lL} (with $\pi_{hH} > \pi_{lH}$ and $\pi_{lL} > \pi_{hL}$ and $y_l > y_h$).⁴ Thus, an application of a type j individual is accepted with probability $\pi_j = \pi_{jH} + \pi_{jL}$ and rejected with probability $1 - \pi_j$. Conditional on an application being accepted, a type j individual receives expected benefits $B_j = \frac{\pi_{jH}b_H + \pi_{jL}b_L}{\pi_j}$; expected benefits for type j from applying are therefore given by $\pi_j B_j$. We assume that $\pi_h > \pi_l$ so that type h individuals are both more likely to have application accepted and are more likely to receive high benefits (conditional on acceptance).

Applying involves both a private and public cost. We denote by c the private application cost, which captures private costs such as the time and effort spent compiling documents, filling out forms, and participating in the interview. Within each type j , individuals vary in their private cost of applying c , which is distributed $f_j(c)$ for each type. The government can determine eligibility and benefit amounts only through reviewing an application, which costs the government g_j for each application processed (whether it is rejected or not). Applications thus generate a fiscal externality on the government.

Finally, we allow for some possibility of individual “mistakes” - such as misperceived expected benefits from applying, or inattention. For our baseline model of “mistakes”, we assume that each type may misperceive their true probability of their application being accepted by ϵ_j , which changes the perceived expected benefit of applying for benefits for type j individuals by changing the perceived probability of acceptance from π_j to $(1 + \epsilon_j)\pi_j$.⁵ We refer to the special case of no misperceptions ($\epsilon_j = 0$) for $j \in \{h, l\}$ as the “neoclassical” benchmark case. While it is possible to have misperceptions in either direction, we focus on the case of under-estimation of the probability of acceptance ($\epsilon_j < 0$); this is both a natural motivation for an information intervention of the type we study, and - we show later - consistent with the empirical results. With $\epsilon_j < 0$, misperception reduces the perceived expected benefit of applying.

Individual application decision and private welfare

In deciding whether to apply for benefits, we assume that individuals maximize their private utility, given their (possibly incorrect) perceptions of the expected benefits from applying. Individuals of type j therefore apply if the expected utility of applying is greater than the utility of not applying.

⁴Note that there could also be individuals who are likely to be ineligible, but we exclude them from model because the experiment targeted individuals who are likely to be eligible.

⁵For simplicity, we assume that the misperception is “proportional” so that it changes the perceived probability of acceptance but does not alter the expected benefits conditional on an application being accepted. Our baseline model also assumes that any misperceptions are the same across individuals within a type; we discuss relaxations of this in Appendix E. Finally, while our baseline model focuses on misperceived probabilities of acceptance as the departure from the neoclassical benchmark, in Appendix E we show that our main results are robust to alternative departures based on misperceptions in cost of applying (rather than in probability of application acceptance) or based on inattention. Non-behavioral frictions could also exist, such as (in the case of non-elderly applications) agency issues within the family between the parent (who bears the cost of applying) and the child (who receives some of the benefits).

For an individual with private application cost c this condition is defined as follows:

$$(1 + \epsilon_j)\pi_j u(y_j + B_j - c) + (1 - (1 + \epsilon_j)\pi_j)u(y_j - c) > u(y_j).$$

If we take a first-order Taylor approximation around $u(y_j)$, the decision to apply for the program simplifies to:

$$(1 + \epsilon_j)\pi_j B_j > c. \tag{1}$$

Equation 1 defines the individual's application decision. We allow for the possibility that the individual decision is not privately optimal by allowing for the possibility of misperception, specifically under-estimation of the probability of acceptance ($\epsilon < 0$).

We can use the private application decision rule in equation (1) and integrate across the distribution of private costs to get the total private welfare of type j individuals:

$$\begin{aligned} V_j &= E[u()|apply] + E[u()|\neg apply] \\ &= \int_0^{(1+\epsilon_j)\pi_j B_j} (\pi_j u(y_j + B_j - c) + (1 - \pi_j)u(y_j - c))f_j(c)dc + \int_{(1+\epsilon_j)\pi_j B_j}^{\infty} u(y_j)f_j(c)dc \\ &\approx u(y_j) + \int_0^{(1+\epsilon_j)\pi_j B_j} u'(y_j)(\pi_j B_j - c)f_j(c)dc \end{aligned}$$

where the last line again uses a first-order Taylor approximation around $u(y_j)$. Note in the above expression that the ϵ_j term affects the individual application decision but not realized utility, since the ϵ_j term only changes perceived probability of acceptance.

Social welfare

We consider a redistributive social welfare function, which is natural given the redistributive purpose of the program. Specifically, we consider a utilitarian social welfare function, although we could easily accommodate alternative individualistic social welfare functions at the cost of additional notation for the social welfare weights on different types. Because the individual's private decision does not internalize the fiscal externality each additional application imposes on the government, even a privately optimal decision may be socially sub-optimal. This reflects the classic tension between private and social optimum. In our setting, it stems from the public costs of processing the application g_j as well as the expected benefit payout $\pi_j B_j$. In other settings there may be alternative or additional fiscal externalities, such as those that arise if individuals distort their labor supply in the presence of means-tested benefits, which in turn reduces the tax revenue for the government. In such settings, the non-benefit component of public costs g_j may vary across individuals. We therefore keep this generality in the model, even though in our specific application we suspect that the costs of processing the application do not vary meaningfully with the individual's type.

We define a utilitarian social welfare function W , which is the sum of total (private) welfare of both types of individuals minus the public cost of processing applications and paying benefits:

$$W = \underbrace{V_l + V_h}_{\text{Private Welfare}} - \underbrace{[(\pi_l B_l + g_l)A_l + (\pi_h B_h + g_h)A_h]}_{\text{Public Costs (Benefits, Application Costs)}} \quad (2)$$

where A_j is the expected number of applications from type j individuals; this is equal to $A_j = 1 - F_j((1 + \epsilon_j)\pi_j B_j)$ from the take-up decision in equation (1). The public cost of an application by type j is $g_j + \pi_j B_j$ - the cost of processing the application (which is incurred regardless of success) plus the expected benefit payout $\pi_j B_j$.⁶ For ease of exposition in what follows, we denote the expected number of enrollees receiving high benefits (b_H) as $E_H = \pi_{hH}A_h + \pi_{lH}A_l$, and the expected number of enrollees receiving low benefits (b_L) is $E_L = \pi_{hL}A_h + \pi_{lL}A_l$. Thus, the total expected benefits (ignoring application costs) can be written as $b_H E_H + b_L E_L$, which is equal to $\pi_l B_l A_l + \pi_h B_h A_h$ but is easier to measure empirically since we do not necessarily directly observe types in our data, but we observe the number of individuals receiving different benefit levels.⁷

Interventions and their targeting properties

We now model two alternative interventions to motivate the design of our experiment. In the **Information Only treatment**, the treatment increases the perceived probability of the application being accepted ($d\epsilon$), which reduces misperceptions if individuals previously had under-estimated the probability of acceptance (i.e., $\epsilon < 0$ initially). In the **Information Plus Assistance treatment**, both the perceived probability and the actual private cost of applying are affected ($d\epsilon, -dc$). The benefit formula is progressive - i.e., $b(y_j)$ is decreasing in y ; therefore we define “targeting” as the share of enrollees who are high-benefit enrollees (relative to low-benefit enrollees); i.e., $e = E_H/(E_H + E_L)$. We say that a treatment T **increases targeting** if $de/dT > 0$.⁸

3.2 Welfare Impacts of Interventions

We begin with an analysis of the private and social welfare effects of the two interventions defined above. The treatments affect application decisions of each type of individual, which in turn affect enrollments. For simplicity, we assume the interventions have zero marginal cost; this would be the case, for example, for interventions that reduce the complexity of the program in a way that increases the perceived probability of acceptance, or that simplify the application process to reduce

⁶We assume this cost is born by someone with the average marginal utility of consumption in society; implicitly, our W expression in equation 2 is thus a “money metric” social welfare expression, normalized by the average marginal utility of consumption in the population.

⁷Given these definitions, the total public costs can be re-written as $b_H E_H + b_L E_L + g_h A_h + g_l A_l$, where $A_h + A_l$ is the total applications across all individuals.

⁸Another type of targeting that we would be natural to analyze in the context of the model is acceptance rate targeting - i.e., the expected share of applicants who are accepted. We do not develop this aspect of the model since, as we will see, neither of our treatments has an effect on acceptance rate targeting. This is consistent with other recent findings of the (non-)impact of reductions in transaction costs on acceptance rates (e.g., Alatas et al., 2016; Deshpande and Li 2017, Armour forthcoming).

application costs. When we use the model to quantify the normative implications of the experimental results in Section 7 below, we consider both costless interventions such as we model here as well as interventions with the actual costs of the ones we implemented.

For notational ease we introduce the following definition:

Definition. Define $\mu_j \equiv -u'(y_j)(\pi_j B_j)\epsilon_j$.

Thus μ_j denotes (for type j) marginal utility of income times the expected benefit from applying times the misperception of the application acceptance probability. Note that if individuals underestimate the probability of acceptance (which is the premise of the information interventions), then $\epsilon_j < 0$ and thus $\mu_j > 0$. The effects of each treatment on social welfare can then be summarized by the following proposition:

Proposition 1. *The effect of the Information Only treatment on welfare is given by the following expression:*

$$\frac{dW^{Information\ Only}}{dT} = \underbrace{\mu_l \frac{dA_l}{dT} + \mu_h \frac{dA_h}{dT}}_{\text{Change in Private Welfare}} - \underbrace{\left[(\pi_l B_l + g_l) \frac{dA_l}{dT} + (\pi_h B_h + g_h) \frac{dA_h}{dT} \right]}_{\text{Change in Public Cost}} \quad (3)$$

And the effect of the Information Plus Assistance treatment on welfare is given by the following expression:

$$\frac{dW^{Information\ Plus\ Assistance}}{dT} = \underbrace{\mu_l \frac{dA_l}{dT} + \mu_h \frac{dA_h}{dT} + u'(y_l)A_l + u'(y_h)A_h}_{\text{Change in Private Welfare}} - \underbrace{\left[(\pi_l B_l + g_l) \frac{dA_l}{dT} + (\pi_h B_h + g_h) \frac{dA_h}{dT} \right]}_{\text{Change in Public Cost}}. \quad (4)$$

Proof: See Appendix E.

To develop intuition for these two expressions, consider first the specific case where individuals have accurate beliefs (i.e., $\epsilon_j = 0$). In this case, the Information Only treatment, which affects the perceived probability of acceptance ($d\epsilon$)⁹, has no effect on private welfare (since $\mu_j = 0$). But since it increases applications, it increases public costs through both processing costs (g_j) and expected benefit payouts ($\pi_j B_j$). The intervention therefore unambiguously reduces social welfare in this case. This is a stark result, but it can be intuitively understood as a consequence of the envelope theorem: individuals have accurate beliefs and are already optimizing, so the marginal applicants who apply as a result of the intervention were close to indifferent to applying before the intervention.¹⁰ Moreover, since there are public costs of processing applications and paying benefits,

⁹Even with accurate beliefs, the Information Only treatment may induce individuals to update their beliefs in response to the information.

¹⁰While this result is obtained in a model where everyone has correct beliefs, in Appendix E we work through an extension where individuals have heterogeneous (but mean unbiased) beliefs, and we prove a similar proposition in this extended model. To do so we need several additional technical assumptions: (1) misperceptions are independent

any treatment which increases applications raises social costs and therefore reduces social welfare. Intuitively, individuals do not consider public costs when making their application decisions, so encouraging more applications for individuals who are (privately) indifferent to applying will reduce social welfare. Typically benefit payments would “net out” because they would show up both as a cost and a benefit; here they show up as a pure cost, since the individual has to pay a private cost c to obtain these benefits, and the marginal individual receives zero benefit (net of private application costs) from enrolling.

For the Information Plus Assistance treatment - which affects both the perceived probability of acceptance ($d\epsilon$) and the private cost of applying ($-dc$) - the overall change in welfare is similar to the Information Only treatment but with two additional terms ($u'(y_l)A_l + u'(y_h)A_h$) that represent the additional change in private welfare from reducing costs for the infra-marginal applicants of both types. In this case, the intervention raises private welfare because it reduces infra-marginal applicants’ application costs. However, the social welfare effect of the Information Plus Assistance treatment is ambiguous and depends on the increased public costs from encouraging more applicants relative to the increase in private welfare from reducing infra-marginal applicants’ private costs of applying.

These results make sense in light of a model where individuals privately optimize with accurate beliefs, and there is a fiscal externality due to the public costs of processing applications and paying benefits. The standard solution to this problem would be to charge a Pigouvian “application tax” to each applicant equal to the public costs. In the absence of such a policy instrument, both changing perceptions of probability of application acceptance as well as increasing “hassle costs” to applicants may serve as alternative instruments for addressing the fiscal externality. Changing perceptions has no welfare effect on the infra-marginal applicants, so one solution would be to have a policy that induces under-estimation of expected benefits such that the marginal applicant has an expected welfare gain from applying equal to the public costs from that application. This is very similar to the Pigouvian solution (in terms of the actual costs of the marginal applicant). By contrast, using hassles instead of “manipulating perceptions” imposes additional costs on all of the infra-marginal applicants. As a result, the second-best solution using “hassle costs” may not be as effective.

If we move away from the neoclassical benchmark and assume that individuals under-estimate the probability of acceptance before the intervention (i.e., $\epsilon_j < 0$), there are now two additional terms to consider in the welfare expressions: $\mu_j \frac{dA_j}{dT}$ for $j \in \{h, l\}$. These two additional terms represent the change in private welfare for the marginal applicants. Recall that without misper-

of private application costs; (2) private application costs are symmetrically distributed around the level of expected benefits; and (3) the heterogeneous misperceptions are symmetrically distributed around 0. Each of these assumptions is needed for the heterogeneous (but mean unbiased) misperceptions to “cancel out” in the expression for the change in welfare from Information Only treatment. We also show in the Appendix E that if beliefs are heterogeneous and mean unbiased and the treatment reduces the variance in beliefs that this will increase private welfare. The intuition for this result is that with heterogeneity in misperceptions, some individuals are over-estimating benefit of applying and some are under-estimating benefit of applying. In both cases, this leads to some marginal individuals making “mistakes” (relative to perfect information benchmark). Reducing variance in beliefs is thus equivalent to reducing magnitude of these mistakes.

ceptions (i.e., $\epsilon_j = 0$), these terms were zero due to the envelope theorem. However, if individuals under-estimate the probability of acceptance before the intervention, marginal applicants who apply experience a positive change in private welfare. The magnitude of the expected welfare change for marginal individuals is increasing in the magnitude of the under-estimation before the intervention.¹¹ Specifically, the marginal enrollee has a non-zero expected private welfare change equal to his marginal utility of income times the size of his under-estimation of the probability of acceptance times the benefits he is expected to receive conditional on applying. However, the social welfare impact of the intervention is ambiguous, given that any private benefits to marginal applicants (and to infra-marginal ones in the case of the Information Plus Assistance intervention) must still be balanced against the increased public costs from processing additional applications and paying benefits.

3.3 Relationship between targeting impacts and changes in welfare

As discussed in the Introduction, a theoretical literature has considered how various barriers to enrollment may potentially deter those with either higher or lower marginal utility from enrollment, and two recent empirical studies (Alatas et al. 2016 and Deshpande and Li 2017) have examined how interventions that increase hassle costs affect the characteristics of applicants and enrollees. These studies - like ours - examined observable characteristics of the applicants and enrollees that are likely correlated with marginal utility of consumption - Alatas et al. (2016) directly measure consumption while Deshpande and Li (2017) analyze health and socioeconomic status. We will follow in this tradition, examining how our information and assistance interventions affect targeting, which we have defined as the share of enrollees who are high-benefit (i.e., low net resource and therefore presumably a higher marginal utility of consumption).

Such analyses rest on the intuition that an improvement in targeting along some observable dimension due to the intervention will, all else equal, increase the social welfare benefits from that intervention. We show here that this is not unconditionally true. We then derive sufficient conditions under which an improvement in targeting due to the intervention increases the intervention’s impact on social welfare.

To see the basic idea, consider the expressions in Proposition 1 for the impact of the intervention on social welfare, and assume that $\mu_l > 0$ but that $\mu_h = 0$. In this case, an intervention that increased applications and enrolled only high-benefit enrollees would unambiguously improve targeting but would have a negative effect on welfare (because there would be no change in private welfare and social welfare would be reduced due to the increased public costs of processing applications and paying benefits). Similarly, in this same set up, for μ_l “large enough”, even an intervention that reduced targeting by only enrolling low-benefit enrollees could still improve welfare since these individuals could still experience change in their private welfare well in excess of

¹¹This result is somewhat similar to the recent analysis of optimal Unemployment Insurance in the presence of biased beliefs in Spinnewijn (2015). In the standard optimal unemployment insurance model, unbiased beliefs and agent optimization imply that behavioral responses will only matter to the extent that they affect the government budget (Baily 1978; Chetty 2006).

the public costs. As a result, there is no general relationship between changes in targeting and changes in welfare. However, our model can be used to describe specific situations where changes in targeting and changes in social welfare go hand-in-hand. This is summarized in the following result:

Proposition 2. *Holding constant the change in applications due to an intervention, the change in social welfare in response to an improvement in targeting ($de/dT > 0$) from either intervention (Information Only or Information Plus Assistance treatment) is given by the following expression:*

$$\frac{\partial}{(de/dT)} \left(\frac{dW}{dT} \right) \Big|_{\frac{dA}{dT}} = [(\mu_h - \mu_l) - ((\pi_h B_h + g_h) - (\pi_l B_l + g_l))] (E_H + E_L) \frac{(E_H + E_L)}{E_H(\pi_{lL} - \pi_{hL}) + E_L(\pi_{hH} - \pi_{lH})} \quad (5)$$

Proof: See Appendix E.

This result shows that the ceteris paribus (i.e., holding constant the change in applications) change in social welfare from a change in targeting is increasing in $(\mu_h - \mu_l)$. The intuition for why μ_h enters positively but μ_l enters negatively is because the change in total applications is held constant in the proposition. As a result, to increase targeting while holding the effect of interventions on applications constant, an application from a type l individual is essentially “swapped out” for an application from a type h individual. For the improvement in targeting holding applications constant to increase social welfare, it must be the case that $(\mu_h - \mu_l)$ - the welfare gain from “swapping” an l individual’s application for an h individual’s application - exceeds the increased public cost of that swap, $(\pi_h B_h + g_h) - (\pi_l B_l + g_l)$.

Recall that $\mu_j \equiv -u'(y_j)\epsilon_j(\pi_j B_j)$. Therefore, in the neoclassical benchmark case ($\epsilon_j = 0$ and thus $\mu_j = 0$ for $j = \{h, l\}$), an improvement in targeting does nothing for private welfare. The lack of private welfare consequences from the intervention’s targeting property follows directly from the envelope theorem: the enrollees who are marginal to the intervention have no change in private welfare and thus their type is irrelevant for private welfare. The impact of an improvement in targeting on social welfare is ambiguous, and depends on how the fiscal externality varies across types. If (as is the case in our application) benefits are progressive ($B_h > B_l$) and non-benefit public costs are constant across types ($g_h = g_l$), then an increase in targeting *reduces* social welfare. This follows directly from the progressive benefit formula: increasing targeting brings in more high-benefit enrollees relative to low-benefit enrollees at higher fiscal cost but with no change in private welfare. This is the exact opposite of the standard intuition that social welfare increases from an intervention that increases targeting on observables that are correlated with the marginal utility of consumption.¹²

¹²The distinction between this result and the “standard intuition” from papers such as Nichols and Zeckhauser (1982) that ordeals can improve social welfare by disproportionately screening out the high ability types is precisely that in their setting, enrollment by high ability types comes with a larger fiscal cost on the government (via the greater loss in tax revenue from their higher labor market earnings). In both settings, if the public cost from enrollment were constant across individuals, the targeting property of the intervention would be irrelevant for the impact of the intervention on social welfare.

Once we move away from the neoclassical benchmark and allow for the possibility that individuals underestimate the probability of application acceptance ($\epsilon_j < 0$), then μ_j is increasing in three type-specific factors: the marginal utility of consumption ($u'(y_j)$), the public cost or fiscal externality (i.e., the expected benefit level plus the application costs) for each type of applicant ($\pi_j B_j + g_j$), and the magnitude of the under-estimation ($-\epsilon_j$). If the expected benefits are higher for h types (i.e., a progressive benefit formula) and the application cost is the same for each type, then the first two type-specific factors are higher for h types than l types. This presumably drives the standard intuition that the change in welfare due to the intervention would be higher with increased targeting on observables correlated with the marginal utility of consumption. (Moreover, with a non-utilitarian social welfare function that puts higher social welfare weights on h , this would add yet another type-specific factor that would be higher for h than l .) However, these are not sufficient conditions because the size of under-estimation is also important in determining relative magnitude of μ_h and μ_l . Assuming that the h types have a higher marginal utility of consumption (or a higher social marginal utility of consumption), then a sufficient condition for an increase in targeting holding constant the change in applications to increase social welfare is that under-estimation is non-zero for at least one type and weakly higher (in absolute value) for the h type (i.e., $\epsilon_h \leq \epsilon_l \leq 0$, with at least one inequality strict). This includes the case where both types under-estimate the probability of application acceptance by same amount (in a proportional sense), so that $\epsilon_h = \epsilon_l < 0$.¹³

Overall, Proposition 2 provides guidance for specific situations in which improved targeting properties of the intervention increase the likelihood the intervention is social welfare improving. In our empirical work below we will show that both the Information Only intervention and the Information Plus Assistance intervention decrease targeting. We will also present evidence consistent with a departure from the neoclassical benchmark ($\epsilon_j \neq 0$) and suggestive of $\epsilon_h < \epsilon_l < 0$. And as noted, we will also assume that g_j is limited to the public application costs, and does not vary across types. Proposition 2 tells us that, under these conditions, our findings that the interventions decrease targeting suggest that they are less likely to increase social welfare.

However the proposition also highlights that there is no general relationship between targeting properties and the likelihood the interventions improve social welfare, and that therefore additional empirical estimates are required in order to draw normative inferences from targeting results. One way to see this is to note that the change in social welfare from the intervention depends on the “net” benefits of individuals affected by the treatment, not the “gross” benefits to individuals (i.e., the amount of the benefits times their marginal utility of consumption). If, for example, the l types under-estimate the probability of acceptance by much more than the h types, the l types could have a higher net private welfare gain from the intervention and thus the welfare of the treatment could be larger the more l types brought in, even though they get lower benefits and have lower marginal

¹³This discussion implicitly assumes that $u'(y_j) > 1$ for $j \in \{H, L\}$, so that the marginal utility of consumption from $\pi_j B_j$ exceeds its social cost. This would follow if, for example, we normalized our expression for $\frac{dW}{dT}$ by the average marginal utility of the population, and both eligible types $j \in \{H, L\}$ have higher marginal utility than the average in the population, as would be expected from any means-tested benefit.

utility of consumption. This is reminiscent of the distinction between gross and net benefits to schooling (Card 2001). In that setting, individuals differ both in their return to schooling and their cost of schooling. As a result, interventions which reduce the costs of schooling will attract students whose return (net of costs) is close to zero, but the “gross” returns to schooling for these individuals is not clear; it may or may not be similar to the average (gross) return to schooling in the population. Similarly, in our model without misperceptions, reduction in costs of applying will attract individuals with very low net benefit of applying, but their gross benefit could be large or small.

Thus far we have considered “marginal” interventions and made heavy use of the envelope theorem. In Appendix E, we develop an extension that considers “non-marginal” changes.¹⁴ We describe how even in this case there is no general relationship between changes in targeting and changes in welfare. Intuitively, in the non-marginal case the increase in private welfare from a given change in enrollees depends on the shape of the type-specific cost distribution $f_j(c)$. If most of the individuals induced to apply were close to indifferent before the non-marginal change in costs, then a non-marginal change in costs can have a non-marginal change in private welfare; however if most of the individuals induced to apply were close to indifferent to applying *after* the non-marginal change, then the non-marginal change in costs will have a negligible effect on their private welfare, since they are close to indifferent to applying after the intervention. Thus, the cost distribution functions introduce another “free parameter” that can affect the relationship between improvements in targeting and changes in social welfare, much as the misperception terms did when we departed from the neoclassical benchmark.

4 Setting and Background

4.1 SNAP Eligibility and Benefits

SNAP is the second-largest means-tested program in the United States after Medicaid (US Congressional Budget Office 2013). It is a household-level benefit designed to ensure a minimum level of food consumption for low-income families (Hoynes and Schanzenbach 2016). Our study focuses on elderly households – i.e., households with an individual aged 60 or over - in Pennsylvania (PA) in 2016. Appendix G provides more detail on the eligibility and benefit rules; we summarize a few key features here.

Eligibility may be categorical - if the individual receives a qualifying benefit such as SSI or TANF - or based on means testing which depends on gross income, assets, and, in some cases, information on particular types of income and expenditures. Once deemed eligible, an elderly household is certified to receive SNAP benefits for 24 months, although there are exceptions that require earlier re-certification.¹⁵

¹⁴ This extension is similar to Kleven (2018), which extends sufficient statistics analysis to discrete policy changes.

¹⁵ At the time of the intervention, households were required to submit an annual reporting form. Additionally, these households were required to report certain changes, such as when gross monthly income exceeds 130% of FPIG.

If eligible, the progressive benefit formula is a decreasing function of net income - gross income minus certain exempt income and deductions for certain expenses - subject to a minimum and maximum. Benefits are designed so that - subject to a minimum and maximum - households spend approximately 30% of their net income (i.e., gross income minus deductions and exemptions) on food. During our study period, the minimum monthly benefit was \$0 or \$16 depending on household type, and the maximum monthly benefit was \$194 for a household size of 1, \$357 for a household size of 2, and \$511 for a household size of 3. In practice, as we will see in our data, there are distinct modes of benefit distribution at the minimum and maximum, which influenced our decision to model a discrete type space in Section 3.

Average monthly benefits per SNAP recipient in elderly households in Pennsylvania were \$136 in fiscal year 2013 (Center on Budget and Policy Priorities, 2015). SNAP benefits are a substantial source of potential income for eligible households. About two-thirds of elderly households in Pennsylvania receiving SNAP had household incomes below the federal poverty line; successful enrollment entitles the household to benefits for at least 24 months, boosting annual income by about 15 percent (Center on Budget and Policy Priorities, 2017).

4.2 Application and Enrollment Process

For an elderly household to successfully enroll in SNAP in PA during our study period, the household must complete an application, provide the necessary documents verifying household circumstances, and participate in an interview (phone or in person). The applicant must provide identifying information (such as Social Security Number, name and address), information about each household member, information on resources and income, and information on various household expenses such as medical expenses, rent and utilities. They must provide documentation verifying their identity, proof of residency, and proof of earnings, income, resources and expenses. Applications can be submitted by mail, fax, in person at the County Assistance Office, or on line. The on-line information and application system in Pennsylvania is considered one of the better state designs (Center on Budget and Policy Priorities, 2016). The interview is used, among other things, to make sure the application is complete and, if not, to request additional documentation or request answers to additional questions. The state has 30 calendar days to process an application.¹⁶ Successful applicants access their benefits electronically, using plastic cards that can be used to buy food in authorized food stores.

Given the SNAP program rules, both the individual and state's determination of eligibility and benefit amounts require the individual to actively apply with the required information. From the individual's perspective, there is uncertainty about the benefit function, the inputs into it (e.g., various shelter and medical expenses that serve as deductions to income and affect benefits), and the potential for mistakes in the process (e.g., not showing up for the interview, not filing the appropriate documentation of expenses, etc.) which cause an otherwise eligible application to be

¹⁶Households who - by virtue of extreme need - qualify for expedited review must have their application reviewed within 5 calendar days of application.

rejected or assigned a lower benefit amount. From the government’s perspective, the needed information cannot be passively obtained, even if it had access to data on the individual from tax returns and other public benefit programs. In particular, three specific types of information are not available from other sources. The first is the definition of a household, which is a SNAP-specific definition: people who “live together and customarily prepare food together” (Gray et al., 2016). The household unit is required both to assess eligibility and to determine benefit amounts. Second, the resource limit that is applied to all non-categorically eligible households requires information on resources like bank accounts and second properties that are not readily available in other administrative data. Third, the calculation of net income – which is required in some cases to determine eligibility and in all cases to determine benefits – likewise can be affected by information not otherwise available (like excess out-of-pocket medical expenses and shelter expenses), although of course one could provide less information here and receive commensurately lower benefits. Underlining the difficulty of circumventing the active application process is the experience of the tax preparer Intuit (TurboTax), which in 2015 tried - through a program called Benefits Assist - to submit applications for SNAP on behalf of their low-income clients, using the information that had been provided on their tax returns. States encountered substantially increased administrative burden in response to the noticeable increase in applications, and it appeared that many of these applications were incomplete and could not be approved as filed.¹⁷

The application imposes costs on both the applicant and the state. Survey evidence from the late 1990s suggests that the average application takes about five hours to complete, including two trips to the SNAP office or other places, and average out-of-pocket costs were about \$10, primarily for transportation (Ponza et al. 1999). However regulatory changes enacted since the time of that survey were designed to reduce applicant costs by, for example, allowing a phone interview in lieu of an in person interview (Hoynes and Schanzenbach 2016).

In addition to the costs on the applicant, the application imposes costs on the government. The state must process applications to determine eligibility, including verifying self-reported information in various available administrative data systems. Estimates from Isaacs (2008) suggest that annualized state administrative certification costs are about 10-15 percent of annual benefits. This is a substantially higher share of benefits than administrative costs for the Earned Income Tax Credit - where there is concern about ineligible people enrolling; this highlights empirically the trade-off that Kleven and Kopczuk (2011) analyze theoretically between greater administrative costs and balancing false rejections and false acceptances.

¹⁷See, e.g., <https://fns-prod.azureedge.net/sites/default/files/snap/State-Guidance-on-Intuit-SNAP-Applications.pdf>; http://www.macssa.org/memberlogin/15minutes/selfsufficiency_dec15.pdf; and <https://benkallos.com/press-release/memorandum-automatic-benefits-using-government-data-deliver-better-citizen-services-le>.

5 Empirical Design and Data

5.1 Design of Interventions

We partnered with Benefits Data Trust (BDT), a national not-for-profit organization based in Philadelphia that strives to be a “one-stop shop” for benefits access, screening individuals for benefit eligibility, and providing application assistance. Since its inception in 2005, BDT has submitted over 500,000 benefit applications across multiple states, resulting in approximately \$5 billion in benefits delivered to low-income individuals and families (Benefits Data Trust 2016). Through existing data share agreements with the Pennsylvania Department of Human Services and other state agencies, BDT receives application and enrollment data for a variety of public benefits, including SNAP and Medicaid. An observational study by Mathematica of six different SNAP outreach and enrollment approaches nationwide concluded that the BDT’s intervention for the elderly in Pennsylvania was the lowest cost per enrollment of any of the methods studied (Kauff et al. 2014), although the 2009 program studied there was somewhat different than BDT’s 2016 approach, which is what we study here.

For our study, as with past BDT SNAP enrollment efforts, the state of Pennsylvania supplied BDT with administrative data that allowed them to identify likely eligible SNAP non-participants. Specifically, BDT received data on individuals aged 60 and older who were enrolled in Medicaid but not in SNAP. Such individuals are likely income eligible for SNAP, since Medicaid tends to have income criteria similar to that of SNAP.

We randomized our study population of approximately 30,000 elderly individuals enrolled in Medicaid into three equally-sized arms. Individuals in the control group received no intervention. Individuals in the Information Only intervention received outreach materials informing them of their likely eligibility for SNAP and the benefits they might receive, and providing them with information on how to call the Department of Human Services to apply. Individuals in the Information Plus Assistance intervention received similar outreach materials but with information on how to call BDT to apply; if they called they then received application assistance. We did not design the Information Plus Assistance intervention; it follows BDT’s current practices for helping to enroll individuals in SNAP.¹⁸ Estimating the cost of these interventions is difficult because there are both fixed costs (determining a set of likely eligible individuals, designing letters and setting up a mailing system, hiring and training staff etc) and variable costs. BDT estimates the variable costs to be about \$1 per attempt for the Information Only intervention; the costs consist primarily of the mailing costs, but there are also printing costs. BDT estimates the variable cost of the Information Plus Assistance intervention to be about \$7 per attempt, with the higher costs due to the labor costs of the BDT staff who provide the assistance.

¹⁸Loosely speaking, our information treatment resembles the control arm in Schanzenbach (2009) - in which likely eligible individuals who expressed interest were handed a blank Food Stamps application with the address of the county agency where it could be submitted - while our Information Plus Assistance arm lies somewhere in between the “lighter touch” and “heavier touch” assistance provided in Schanzenbach (2009).

Information Plus Assistance

BDT conducts a series of outreach services to inform these individuals of their likely eligibility, and assist them in applying for benefits. This outreach has two components: information and assistance. The information component consists of proactively reaching out by mail to individuals whom they have identified as likely eligible for SNAP, and following up with a postcard after 8 weeks if the individual has not called BDT. Letters and postcards inform individuals of their likely SNAP eligibility (“Good news! You may qualify for help paying groceries through the Supplemental Nutrition Assistance Program (SNAP)”) and typical benefits (“Thousands of older Pennsylvanians already get an average of \$119 a month to buy healthy food”), and provide information on how to apply (“We want to help you apply for SNAP!”), offering a number at BDT to call (“Please call the PA Benefits Center today. It could save you hundreds of dollars each year”). These materials are written in simple, clear language for a 4th to 6th grade reading level and are sent from the Secretary of the Pennsylvania Department of Human Services. Appendix Figure 1 shows these standard outreach materials. In the framework of Section 3, we think of this intervention as increasing the perceived probability of an application’s acceptance ($d\epsilon$).

The assistance component begins if, in response to these outreach materials, the individual calls the BDT number. BDT then provides assistance with the application process. This includes asking them questions so that BDT staff can populate an application and submit it on their behalf, advising them of what documents they need to submit, offering to review and submit documents on their behalf, and assisting with post-submission requests or questions regarding the application. BDT also tries to ensure that the individual receives the maximum benefit for which they are eligible by collecting detailed information on income and expenses (the latter contributing to potential deductions). Appendix A provides more detail on the nature of BDT’s assistance. In the framework of Section 3, we think of this intervention as reducing the private costs of applying ($-dc$). In our setting, those costs are born by BDT, so there is not obviously a reduction in the total (private + BDT) cost of applying. However one could imagine changes in the application process that produced a net reduction in costs.

Data from our intervention indicate that BDT submitted about 70 percent of applications made by individuals in the Information Plus Assistance intervention, and provided their full set of services (including document review) for about two-thirds of the applications it submitted.¹⁹ For callers who end up applying, BDT has about 2.1 calls, averaging (in total) about 47 minutes across the calls. For callers who end up not applying, the average time on the phone is about 30 minutes.

Information Only

Our “Information Only” intervention contains only the letters and follow-up postcards to non-respondents sent as part of the outreach materials. They are designed to be as similar as possible

¹⁹Given that, as we will see in the results below, we estimate that about one-third of applicants are always takers, this suggests that BDT submits applications for the vast majority of compliers, and provides their full set of services for about three-quarters of these compliers.

to the information content of the Information Plus Assistance intervention: both are sent from the Secretary of the Pennsylvania Department of Human Services (DHS) and include virtually identical language and layout. Some minor differences were naturally unavoidable. In particular, the Information Plus Assistance materials direct individuals to call the PA Benefits Center (the local name of BDT) while the Information Only materials direct them to call the Department of Human Services (“Please call the Department of Human Services today. It could save you hundreds of dollars each year”). In addition, the hours of operation for DHS (8:45am-4:45pm) listed on the Information Only outreach materials differed slightly from the BDT hours (9:00am-5:00pm) listed on the Information Plus Assistance outreach materials. Appendix A provides more details, and Appendix Figure 2 shows the outreach materials in the Information Only arm.

Sub-treatments

Within each treatment, we created additional sub-treatments in the presentation and frequency with which the information was presented. In practice, most of these sub-treatments had little or no impact and therefore in most of our analysis we pool them. Appendix A provides more detail of the sub-treatments and how they were distributed across arms. One sub-treatment we will present separately in the main text is the one where we found substantial effects: the elimination of the postcard follow-up in the standard Information Only intervention.

5.2 Study Population

Our study population consists of individuals aged 60 and older who are enrolled in Medicaid but not SNAP. They are considered likely income eligible for SNAP based on their enrollment (and hence eligibility) for Medicaid. This is, of course, an imperfect proxy of SNAP eligibility. This is by necessity; as described in Section 4, exact assessment of SNAP eligibility requires non-income information that must be actively supplied on an application; eligibility cannot be passively determined through existing administrative data.

Our study population thus consists of individuals already enrolled in at least one public benefit program: Medicaid. This is a particular subset of people eligible for but not enrolled in SNAP. For example, our analysis in the pooled 2010-2015 Consumer Expenditure Survey (CEX) suggests that only about 20 percent of individuals aged 60 and over who are not enrolled in SNAP but have income less than 200 percent of FPL (a rough proxy for potential SNAP eligibility) are enrolled in Medicaid. Caution is always warranted in generalizing findings beyond the specific study population. In this particular case, one might be concerned that enrollment in another public benefit program could be indicative of the study population’s general knowledge about benefit eligibility, or interest and ability to sign up for government services. This particular issue, however, may not be a major concern. Many individuals do not actively choose to enroll in Medicaid themselves but rather are enrolled in Medicaid by social workers at hospitals when they arrive uninsured and ill – a fact that has led researchers to refer to many of those eligible for Medicaid but not currently enrolled as “conditionally covered” (Cutler and Gruber, 1996).

An ancillary benefit of using Medicaid enrollment as a proxy for likely eligibility is that it enables us to measure a number of demographic characteristics of our study population, pre-randomization. In particular, the Medicaid file contains birth date, gender, address, and primary language, as well as which Medicaid program the individual is enrolled in. Additional characteristics of our study population come from their 2015 Medicaid enrollment and claims data, which DHS also supplied for everyone on the outreach list. The 2015 Medicaid enrollment file allows us to measure the individual’s race, and the individual’s history of prior Medicaid enrollment spells. In addition, the linked 2015 Medicaid claims file allows us to construct detailed healthcare utilization and health measures in the year prior to the 2016 intervention.²⁰

Summary statistics

To construct the study population, DHS supplied BDT with a list of approximately 230,000 individuals aged 60 and older who were enrolled in Medicaid as of October 31, 2015; DHS also merged on a flag for whether the individual was currently enrolled in SNAP. Table 1 illustrates the construction of our study population and the pre-randomization characteristics of the sample. Column 1 shows the initial outreach list of 229,584 individuals aged 60 and over enrolled in Medicaid as of October 31, 2015. In column 2 we exclude individuals enrolled in the Long-Term Care Medicaid program (N= 47,729) and individuals with an address in Philadelphia (N= 37,932). Of the resulting 143,923 individuals on the outreach list after these exclusions, column 3 shows the 84,038 (about 60 percent) who were enrolled in SNAP or living with someone enrolled in SNAP, while column 4 shows the 59,885 who were not enrolled in SNAP and not living with anyone in SNAP; recall that SNAP is a household-level benefit. Column 5 shows our final study population of 31,888 individuals. These are a subset of the individuals not enrolled in SNAP in column 4. From column 4, we randomly select one individual from each “household” (this excludes 1,842 individuals)²¹, and excluded all individuals to whom BDT had previously sent any outreach materials (N=26,155).²²

There is no clear demographic gradient between Medicaid enrollees who do and do not enroll in SNAP. Columns 3 and 4, respectively, describe some characteristics of each group. Those not on SNAP (column 4) are older, with similar gender, racial, and language makeup than those not on SNAP (column 3).

An additional complication is that only about three-quarters of our study population was enrolled in Medicaid for the entirety of 2015. We therefore annualized all of the health care utilization measures by dividing by the number of days enrolled out of 365. Needless to say this is an imper-

²⁰Medicaid in PA is provided either fee-for-service or managed care, determined in large part based on geography. Our “claims” data are therefore a mix of encounter data from Medicaid Managed Care and Fee for Service claims. Although there are well-known measurement issues with encounter data - and comparability issues with fee for service claims data (e.g., Lewin Group 2012) - such measurement issues should not bias our comparisons of these measures across randomly assigned arms.

²¹There is no household identifier in the Medicaid outreach file; BDT therefore created a pseudo-“household” ID to identify individuals on the outreach list sharing the same last name and address.

²²BDT has comprehensive data on outreach efforts since 2012, and limited data on outreach back to 2007. BDT started conducting SNAP outreach in 2008.

fect approach, since utilization during a partial coverage year may be disproportionately higher (or lower) than it would be if coverage existed for the full year. However, we are not unduly concerned given that this adjustment will affect enrollees in randomly assigned arms equivalently, and we confirm this below.

On some dimensions they appear sicker - they have more hospital days and Skilled Nursing Facility (SNF) days, but on other dimensions they appear less sick - such as fewer chronic conditions. One notable difference is that those not on SNAP have been on Medicaid for less time (i.e., only one third had their last enrollment spell starting before 2011, compared to about one half of those on SNAP).²³

5.3 Randomization

Our interventions randomize the information and application assistance provided to the study population. Specifically, the 31,888 individuals in our study population were randomly split into three equally sized groups: Information Only treatment, Information Plus Assistance treatment, and control (no intervention). There were separate sub-treatments within each treatment: one-quarter of each treatment was randomized into an arm with a variant of the outreach letters and postcards designed to attract clients by using a “marketing” approach that borrowed language and graphics from credit card solicitations; in the Information Plus Assistance treatment the remaining three-quarters received the standard outreach (“standard”); in the Information Only treatment, one-quarter received the standard outreach, while another one-quarter received the standard letter but no follow-up postcard (“no postcard”) and another one-quarter received a letter that varied the description of the expected benefit amounts (“framing”). See Appendix A for more detail.

For practical reasons, the outreach letters were randomly distributed across 11 separate, equally-sized weekly mailing batches. The first batch was sent on January 6, 2016, and the last on March 16, 2016; follow-up postcards were sent eight weeks following each mailing, with the last postcards scheduled to be sent on May 11, 2016.²⁴ Appendix Figure 4 provides more detail on the timing of the mailings.

We wrote the computer code that assigned individuals to these different treatments and treatment mailing batches by simple random assignment according to the share we wanted in each arm; this code also randomly assigned the control individuals to (non-) mailing weekly batches, so that outcomes for all individuals in our study can be measured relative to an initial “mail date”. Imple-

²³We also used national survey data to compare consumption and income for individuals on Medicaid who either were or were not on SNAP; measurement of Medicaid and SNAP enrollment likely has more error in such survey data than in administrative data (Meyer and Mittag 2015). We limited our analysis to individuals 60 and over on Medicaid in households below 200 percent of the poverty line (a proxy for likely SNAP eligibility). In the pooled 2010-2015 Consumer Expenditure Survey, those on Medicaid but not on SNAP have higher (mean or median) per capita consumption across a variety of consumption measures. In the 2016 March CPS, those on Medicaid but not on SNAP have higher (mean or median) per capital household income; for example, mean per capita annual household income is about \$10,500 among those enrolled in SNAP compared to about \$12,700 among those not enrolled. However mean (or median) total household income is virtually identical between the two groups.

²⁴Due to an implementation error, postcards for the January 27 and February 3 batches were not mailed when scheduled and instead were sent on May 26 and May 27, respectively.

mentation of the code on the actual, identified data was done by our BDT partner who had access to these data and oversaw the physical mailings. BDT staff also performed a series of quality assurance tests that we programmed to ensure fidelity of the randomization protocol and the quality of the de-identified data that we received. Appendix Table A2 shows balance of the characteristics of our study population across the arms, as would be expected based on our randomized design.

All study materials, including letters, postcards, and envelopes, were approved by BDT and the Department of Human Services (DHS) before the study was launched. MIT’s Institutional review board (IRB) approved this research as well as the data sharing outlined in Appendix B (Protocol: 1506106206; FWA: 00004881).²⁵ The trial was registered on the AEA RCT Registry (AEA RCTR-0000902) in October 2015, prior to our launch - at which point we pre-specified our primary and secondary outcomes.²⁶ We updated the registry to specify additional detail - such as a 9 month time frame for the outcomes - and to post the more detailed analysis plan in March 2016, prior to receiving any data on applications or enrollment.²⁷

5.4 Outcomes data

Applications, Enrollment and Benefit Amounts DHS provided data on SNAP applications from March 2008 through February 2018. The application data also include disposition codes and dates, which enable us to determine if and when the application was approved; we use this to measure enrollment. Our enrollment measure is therefore a flow measure (“did the individual enroll within n months after the initial mail date”) rather than a stock measure of whether the individual is enrolled as of a given date. We also observe whether and when an application was rejected, as well as the reason for rejection. Our main analysis focuses on application and enrollment within 9 months after the intervention; however we also examine the time pattern of effects through 23 months post intervention.

DHS also provided us with monthly benefit amounts for enrolled individuals. We measure the monthly benefit amount in months enrolled in the 9 months post outreach. The monthly benefit amount will serve as one of the key measures of enrollee characteristics.

²⁵Northwestern University’s IRB (FWA: 00001549) ceded approval to MIT’s IRB through an IRB Authorization Agreement. The IRB of the National Bureau of Economics (NBER) judged the protocol to be exempt (IRB Ref#15_129; FWA: 00003692).

²⁶Specifically, at that time we wrote: “Primary outcome: number of SNAP enrollees. Secondary outcomes: baseline characteristics of enrollees (e.g., demographics, measures of economic well-being, measures of health etc); number of SNAP applications; baseline characteristics of applicants; number of responses to outreach letters (i.e., phone calls to the number listed on the outreach letter). Outcomes (explanation) We are interested in measuring characteristics of the enrollees, for example measures of economic well-being, demographics and health status. Which characteristics we measure and how we measure them will depend largely on the quality and availability of data.”

²⁷Our analysis hews closely to the analysis plan in terms of the take-up outcomes analyzed (calls, applications, and enrollment) and the analysis of enrollee benefits and enrollee and applicant demographic and health characteristics. The exact analysis of study participant characteristics was still left unspecified at that point due to uncertainty on data availability. However, we were unable to execute on our aspirations to analyze additional characteristics like earnings and credit report outcomes due to lack of the relevant data.

Call-in data BDT tracks all calls it receives, which allows us to measure call-ins to the BDT number given in response to the outreach letters in the Information Plus Assistance treatment. In order to capture comparable information on which individuals call in to DHS in response to the Information Only treatment, we contracted with a call forwarding service, and the information-only outreach letters provided the 1-800 numbers of the call forwarding service, with a different call-in number in each sub-treatment arm. Call receptionists were asked to record the individual’s unique identification number (printed on the outreach materials) before forwarding the call to DHS. Appendix C provides more details on the call-in data, the call forwarding service, and the script for the receptionists, which was provided in English and in Spanish. The use of the call forwarding service allows us to measure for each individual in the Information Only treatment whether (and when) they called in in response to the outreach. It also allowed BDT to send follow-up postcards to non-callers in the Information Only intervention, as in the Information Plus Assistance intervention.

We have caller data from January 7, 2016 through October 14, 2016. We use these data to measure calls in the seven months after the initial mail date. We report the “raw” call-in rates in each study arm. Because the call forwarding service is not as good at determining the identity of callers as our BDT partner, the Information Only treatment has a non-trivial number of callers without a valid study ID. We therefore also report an “adjusted” call-in rate for the Information Only treatment, which adjusts the measured call-in rate to account for our estimate of the rate of unrecorded callers. Appendix C provides more details on this adjustment procedure.

6 Results

Our main analysis compares across three groups: the (pooled, equally-weighted) standard and marketing treatments in the Information Only arm (5,314), the (pooled, equally-weighted) “standard” and “marketing” treatments in the Information Plus Assistance arm (10,629), and the control (10,630) (see Appendix Figure A3.) In Appendix Tables A4, A5, A12 and A13 we present the full set of results separately for each sub-treatment; in general these sub-treatments had little or no impact, except the “no reminder postcard” sub-treatment which we discuss below.

6.1 Behavioral Responses to Intervention

Our primary measure of behavioral response is SNAP enrollment. However, we also examine intermediate steps toward enrollment: calls in response to outreach and applications submitted. We measure all of these outcomes over the first nine months after the mail date. As a result, our outcomes data span the period January 6, 2016 (the date of the first mailing) through December 16, 2016 (nine months after our last mailing). This was chosen to be a sufficiently long window to capture the full impact of the intervention on these outcomes.²⁸ Someone who enrolls in SNAP over

²⁸Based on their prior outreach efforts, BDT estimated that initial responses (i.e., calls) if they are going to occur will happen within the first three months. We wanted to allow sufficient time for applying and for the state’s 30-day decision time and erred on the side therefore of long time window to make sure we had the full impact.

these nine months likely remains enrolled for about 24 months, after which they need to re-certify their eligibility, although in rare cases they may need to re-certify earlier. Our current data allow us to extend the analysis out to 23 months; this is not yet long enough to study re-enrollment.

Enrollment, applications, and calls

Table 2 presents the main take-up results of the experiment by intervention arm. All outcomes are measured in the nine months after the initial mail date. The first row shows results for enrollment. In the control group, about 6 percent enroll. The Information Only intervention increases enrollment by 5 percentage points, or 83 percent relative to the control. Information Plus Assistance increases enrollment by 12 percentage points, or 200 percent relative to the control; the impacts of the intervention are statistically different from the control and from each other ($p < 0.001$).²⁹

Figure 1 shows the time pattern of intervention impacts on enrollment by month, for the 23-month period after the initial mail date. The time pattern is similar for both interventions: over 85 percent of the 9-month enrollment effect is present by 4 months, and the impact has clearly leveled off before 9 months (our baseline time window). The impacts of the intervention appear to largely persist, as least through the 23 months we can observe post-intervention: About 90 percent of the 9-month enrollment effect is present by 23 months. This suggests that the interventions are primarily generating new enrollment, as opposed to merely “moving forward” in time enrollment that would otherwise happen.

The next two rows of Table 2 show the impact of the interventions on applications, and on their failure (rejection) rate. The impacts on applications are roughly proportional to the increase in enrollment. About 22 percent of applications in each arm are rejected; differences across arms are substantively and statistically indistinguishable.³⁰ This suggests that assistance affects enrollment (over and above information alone) primarily by affecting individuals’ willingness to apply, rather than by increasing the success (i.e., approval) rate of a given application. Of course, since assistance may also change the composition of applicants (including their latent success probability), it is not possible to directly identify these two separate channels.

In Appendix Table A8 we briefly explored the nature of the “reasons” given by DHS for the rejections. Naturally these are not always straightforward to interpret. Nonetheless, it appears

²⁹For some perspective on these numbers, we considered how they compared to other take-up interventions, bearing in mind that these were different interventions conducted on different programs and populations. In the context of encouraging low-income high school seniors to apply for aid and attend college, Bettinger et al. (2012) found that providing information about aid eligibility and nearby colleges had no detectable effect, but combining the information with assistance in completing a streamlined application process increased college enrollment by 8 percentage points or about 25 percent relative to the control. In the context of informing low-income tax filers about their likely eligibility for the EITC, Bhargava and Manoli (2015) found that their average informational outreach increased EITC filing by 22 percentage points (or about 50 percent above baseline).

³⁰In other contexts, changes in transaction costs have similarly had a small effect on rejection rates of applicants. Deshpande and Li (2017) find that increasing transaction costs (via closings of Social Security field offices) results in about a two percentage point increase in rejection rates for SSDI applicants, while Alatas et al. (2017) find that increasing transaction costs (via requiring individuals to apply for the benefit as opposed to the government automatically screening potential eligibles) decreases rejection rates by about two percentage points in a conditional cash transfer program in Indonesia.

that relative to the control, the share of rejections in the Information Plus Assistance arm is higher for reasons that looks like “insufficient interest” on the part of the applicant - e.g., withdrew or didn’t show up for an appointment - and lower for reasons that look like ineligibility after review - e.g., failure to meet citizenship or residency requirement. This is consistent with assistance reducing the error rate on applications, but also pushing marginally motivated individuals to start the application process.

The last six rows of Table 2 examine call-in rates. A caller is defined as someone calling the number provided on the outreach material; the caller rate is therefore mechanically zero for those in the control arm.³¹ We show the raw caller rate, as well as the adjusted caller rate designed to account for the lower measurement of callers in the Information Only arm (see Section 5.4 above). The raw rates show a 30 percent caller rate in response to the Information Plus Assistance outreach letters, and a 27 percent caller rate in response to the Information Only outreach letters; with the adjustment, the adjusted caller rate rises to 29 percent in the Information Only intervention, statistically indistinguishable from the response to the Information Plus Assistance.³² The similar call-in rate is not surprising given the (deliberate) similarity of the outreach materials (recall Appendix A and Appendix Figures 1 and 2). It suggests that any difference in applications and enrollment between the Information Only and Information Plus Assistance interventions is attributable to the assistance itself, rather than to the expectation of assistance.³³ Conditional on calling, we find the average caller made 1.8 calls in the Information Plus Assistance arm and 1.6 calls in the Information Only arm (results not shown); these differences are statistically distinguishable ($p < 0.001$).

The results in the last four rows suggest that all marginal applicants affected by the interventions call in response to the outreach materials: the share of people who apply without calling is the same in all three arms. Such individuals presumably call the state directly (without being routed through BDT or our tracking service), and/or apply on-line or in person. Caller rates therefore provide a likely ceiling for the impact of the interventions: less than one-third of individuals appear to notice and respond to the outreach materials. The other 70 percent likely received the outreach materials, since less than 1 percent were returned to sender due to bad addresses. It is possible that they did not open or read the materials, or did so but were not moved by the materials to apply for SNAP benefits. Presumably at least one-quarter of non-callers are actually ineligible for SNAP, given that one-quarter of applications are rejected; perhaps an even larger share of non-callers

³¹Appendix Table A7 shows callers from each intervention arm into each possible number (with a different number for the Information Plus Assistance arm and for each sub-treatment in the Information Only arms). There is virtually no cross-contamination.

³²Appendix Figure A7 shows the monthly pattern of callers and applications post intervention. The time patterns are similar to what we saw for enrollment in Figure 1.

³³As described in more detail in Appendix A, callers receive assistance from a Benefit Outreach Specialist (BOS). BDT assigns callers to Benefit Outreach Specialists using a rotation system that is similar to how cases are often assigned to judges (i.e., a rotation system). As a result, we can plausibly assume that the assignment of a BOS to a caller is quasi-random. Using this assumption, we estimate treatment effect heterogeneity across Benefit Outreach Specialists using a non-parametric Empirical Bayes approach, and (perhaps surprisingly) we do not find statistically significant evidence of treatment effect heterogeneity - i.e., we cannot reject that the BOS fixed effect estimates are all equal to each other.

believe themselves (potentially correctly) to be ineligible.³⁴

If we interpret calling as a sign of interest, the results show that, conditional on interest, the application rate is twice as high when assistance is provided (about 60 percent) than when only information is provided (about 30 percent). Likewise, enrollment rates (conditional on interest) are about 45 percent when information and assistance is provided compared to 23 percent when only information is provided.

All of the results shown in Table 2 are based on comparisons of mean outcomes by intervention arm. No covariates are needed given the simple random assignment. For completeness however, we show in Appendix Table A14 that all of the results in Table 2 are robust to controlling for baseline demographic and health characteristics of the individuals, as well as for the date of their mail batch.

Cost effectiveness approximation

A rough back of the envelope calculation suggests that the Information Only intervention was about two-thirds cheaper per additional enrollee than the Information Plus Assistance intervention. Separating out fixed and marginal costs of the intervention is difficult, but BDT has estimated the marginal cost of the Information Plus Assistance intervention at about \$7 per individual who is sent outreach materials, and the marginal cost of the Information Only treatment was about \$1 per individual who was sent outreach materials; the cost of the latter is primarily composed of the cost of mailing a first class letter (\$0.49 at the time of our intervention) plus the cost of the follow up postcard (\$0.34 at the time of our intervention), plus the costs of printing and assembling the mailings. These numbers suggest that the cost per additional enrollee is \$20 in the Information Only treatment, compared to \$60 in the Information Plus Assistance treatment. Naturally there are additional costs to the applicants from the time spent applying and to the government from processing applications and paying benefits.

Because SNAP benefits are financed by the federal government, our results suggest that the state benefits financially from encouraging take-up, even if it bears the whole intervention cost as well as the processing costs. As we will see below, new enrollees receive, on average, about \$1,300 per year in annual SNAP benefits. This is paid for by the federal government. Isaacs (2008) estimated that the annualized administrative costs of the SNAP program (including certification costs as well as subsequent administrative costs) are about \$178 per receipt, or about \$134 per application given our estimate of a 75% acceptance rate; this is paid for by the state government. Thus, were the state to finance the marginal costs of either the Information Only intervention (\$20 per enrollee) or the Information Plus Assistance intervention that BDT currently undertakes (\$60

³⁴We generated a predicted eligibility measure in which we predict eligibility based on the relationship between application approval and the pre-randomization demographic and health characteristics shown in Table 1. We estimated this prediction using the sub-sample of the study population that applied, and then used the estimates to predict eligibility rates among non-callers and callers separately. The results suggest that - at least based on the available observable characteristics - non-callers had only a few percentage points lower predicted eligibility rates than callers.

per enrollee) as well as the administrative costs of processing the applications, these would still be less than 25 percent of the new federal benefits received by state residents, and presumably spent largely at local retail outlets. As a result, both of these interventions seem to pass a very simple cost-benefit test from a state public finance perspective. However, note that this conclusion would generate a different conclusion if virtually all of new enrollees received minimum benefit level (\$16 per month or \$192 per year); this would be similar to the state’s administrative costs.

Effects of reminders

We conducted a number of sub-treatments that varied the presentation and frequency of the information sent. Table 3 shows results for two of the information-only sub-treatments. Specifically it shows results for the standard treatment, which includes an initial letter and a reminder postcard 8 weeks later if the individual has not yet called in (see Appendix Figure A2), and a “no reminder postcard” sub-treatment in which the follow-up postcard is not sent.³⁵

The results indicate that reminders matter: all behavioral responses decrease by about 20 percent without the reminder postcard. Specifically, the standard Information Only intervention (with the reminder postcard) had a 30 percent call rate, a 15 percent application rate and an 11 percent enrollment rate. The lack of a postcard reminder reduced the caller rate by 7 percentage points ($p < 0.001$), the application rate by 3 percentage points ($p = 0.001$) and the enrollment rate by 2 percentage points ($p = 0.016$). Given the 2 percentage point increase in enrollment with the reminder postcard, and its marginal cost of roughly \$0.35, cost per additional enrollee is similar with and without the reminder postcard.

The non-trivial impact of a reminder postcard is similar to Bhargava and Manoli’s (2015) finding that a similar second reminder letter, sent just months after the first, increased EITC take-up. They interpret the effect of the reminder as indicative of it combating low program awareness, inattention, or forgetfulness; they present additional survey evidence consistent with low program awareness among those eligible for the EITC. A similar interpretation seems warranted in our context, where we estimate that less than 3 percent of our study population had applied for or enrolled in SNAP in the 10 years prior to our intervention. In addition, surveys suggest that about half of likely eligible, non-participants in SNAP reported that they were not aware of their eligibility (Bartlett et al. 2004). In our framework in Section 3, this is modeled as under-estimating the probability of eligibility (i.e., $\epsilon < 0$). An alternative explanation - that the impact reflects a high rate of non-delivered mail - does not seem warranted; less than 1 percent of outreach materials were returned to sender.

³⁵The results for the Information Only treatment results shown in Table 2 pool the results from the standard treatment and a “marketing” sub-treatment that varied the content of the outreach letters (see Appendix A and Appendix Figure A5 for more details); these two sub-treatments are pooled in the same proportions in the Information Plus Assistance treatment results shown in Table 2.

6.2 Characteristics of Marginal Applicants and Enrollees

We examined the characteristics of the marginal applicant or enrollee whose behavior is affected by the intervention. To do this, we define the outcome in each arm as the average of a specific characteristic among those who (endogenously) apply or enroll. For example, we compare the average monthly benefits among those who enroll in each arm, or the average healthcare spending in 2015 of individuals who enroll (or, alternatively, apply) in each arm. Differences in the average characteristics of enrollees or applicants in a given treatment arm relative to the control group reveals how the characteristic of the marginal individual who apply or enroll due to a given intervention differs from the average applicant or enrollee who would enroll absent the intervention. This approach to analyzing the characteristics of the marginal person affected by an intervention is analogous to approaches taken in prior work by Gruber et al. (1999) and Einav et al. (2010).

We have already seen that one characteristic of marginal applicants - their probability of rejection - was similar to control applicants (see Table 2). Here, we present evidence that marginal applicants and enrollees in either intervention arm are less needy than the average applicants and enrollees who apply in the absence of the intervention. For brevity, we focus the discussion on a comparison of characteristics of enrollees in the control group relative to enrollees in either intervention. However, the tables also show that characteristics tend to be similar between the two intervention arms, and within each intervention arm, between applicants and enrollees (and also callers - see Appendix Table A11). This suggests that the interventions generated interest (calls) in different types of individuals than the average applicant or enrollee, but that, conditional on inducing interest, there was no further differentiation in the characteristics of those who applied or who successfully enrolled.

Monthly benefits among enrollees

Table 4 shows monthly benefits for individuals who enrolled in the 9 months after the initial mail date, by study arm. Because the SNAP benefit formula is progressive, a lower benefit amount implies an enrollee with higher net resources. This is why in Section 3 we defined “targeting” as the share of enrollees who are high benefit (relative to the share that are low benefit). Average monthly benefits for enrollees in the intervention arms are 20 to 30 percent lower than for enrollees in the control arm. Average monthly benefits are \$146 in the control compared to \$115 in the Information Only intervention and \$101 in the Information Plus Assistance intervention; average benefits in each intervention arm are statistically different from those in the control ($p < 0.001$) as well as from each other ($p = 0.013$).³⁶

There are clear modes in the distribution of benefits received, corresponding to minimum and maximum benefit amounts. Among the controls, 18 percent receive \$16 (the minimum monthly

³⁶Differences in the average characteristics of enrollees in an intervention arm relative to the control arm reflect differences between the average characteristics of infra-marginal enrollees (or “always takers”) relative to marginal enrollees (or “compliers”). As another way of presenting the same information, Appendix Table A6 reports the average characteristics for always takers and compliers; estimation of these objects is standard (see, e.g., Abadie 2002, Abadie 2003, or Angrist and Pischke 2009) and we describe it in more detail in Appendix F.

benefit for a household of size 1 or 2 who are categorically eligible) and another 19 percent receive \$194 (the maximum monthly benefit for a household of size 1); see also Appendix Figure A8. Table 4 shows that the interventions increased the share of enrollees receiving the minimum benefit and decreased the share of enrollees receiving the maximum benefit.

In principle, we should observe benefits for all individuals whose applications have been approved during our nine-month observation window (our measure of “enrollee”). In practice, we are missing such information for about 4 percent of enrollees, and this missing rate is not balanced across arms. As shown in Table 4, 7.3 percent of control enrollees are missing benefit information, compared to 4.3 percent of the Information Only enrollees and 2.8 percent of the Information Plus Assistance enrollees; differences in missing benefit rates are statistically significantly different between either intervention arm and the control group. Such non-random attrition could bias our comparison of enrollee benefits across arms.

Therefore, we also generated a predicted benefit measure in which we predict the benefit amounts based on the relationship between benefits and the pre-randomization demographic and health characteristics shown in Table 1; Appendix D provides more detail on the prediction algorithm which follows a standard algorithm in machine learning (Rifkin and Klautau 2004). Table 4 shows that predicted benefits show the same pattern across arms as actual benefits, both among enrollees with non-missing benefit amounts (second to last row) and among all enrollees (last row).

Another potential concern is that benefits increase in household size. If the interventions disproportionately encourage smaller households to apply, this will lower enrollee benefits without necessarily reflecting higher *per capita* resources. Indeed, the penultimate row of Table 4 shows that the interventions increase the share of enrollees who are in a household size of 1. However, the bottom row of Table 4 shows that if we limit our analysis to households of size 1, average benefits for these households are still statistically significantly lower in each intervention arm relative to the control. An additional attraction of limiting to households with only a single individual is that we have essentially no missing benefits for such households.

Demographics and health of applicants and enrollees

Table 5 shows the demographic and health characteristics of applicants and enrollees. The first four columns show these characteristics for applicants, and the last four for enrollees. For a variety of demographic and health characteristics, marginal applicants and enrollees from the intervention appear less needy than the average applicant or enrollee in the control group.

Panel A shows results for a summary measure: predicted benefits, where the prediction is based on the underlying pre-randomization demographic and health information and higher predicted benefits correspond to lower net resources given the progressive SNAP benefit formula. The results for enrollees were already seen in Table 4. Results for applicants are similar: applicants in either intervention arm have lower predicted benefits than applicants in the control arm ($p < 0.001$).

Panel B shows results for health and healthcare, measured in the calendar year prior to the intervention. We measure health care utilization in three different ways: total medical spending,

total number of visits or days (summed across emergency room (ER) visits, doctor visits, hospital days, and skilled nursing facility (SNF) days), and weighted number of visits or days, where the weights are set based on the average cost per encounter.³⁷ Total medical spending is noisy - due to the well-known high variance of medical spending - and conflates variation in utilization with variation in recorded prices. Our total number of days or visits measures attempt to circumvent both problems by creating a utilization-based measure. The weighted utilization measure is designed to account for the fact that a hospital day is substantially more expensive than a SNF day or a doctor visit. Relative to the control arm, applicants and enrollees in the intervention arms use less health care.³⁸ For all three health care utilization measures, applicants and enrollees in the intervention arms have lower pre-randomization health care use than in the control arm, although these differences are not always statistically different from the control. However, when we pool across both intervention arms, the total number of visits and days and the weighted total number of visits and days are statistically different from the control arm ($p < 0.05$) for both applicants and enrollees (not shown).

In the final row of Panel B we show that the number of measured chronic conditions is also lower in both intervention arms relative to the control arm for both applicants and enrollees, with most of these differences statistically significant at conventional levels. A smaller number of chronic conditions could reflect better underlying health. It could also - partly or entirely - reflect lower health care utilization; since chronic conditions are only measured if the individuals use the relevant health care, they are a joint measure of underlying health and health care utilization (Song et al., 2010; Finkelstein et al., 2017).

Panel C shows results for demographic characteristics measured pre-randomization. Relative to the control group, applicants and enrollees in either intervention arm are statistically significantly ($p < 0.001$) older, more likely to be white, and more likely to have their primary language be English. For example, 71 percent of control enrollees are white, compared to 78 percent in either intervention arm.

7 Normative implications

We use the framework from Section 3 to analyze the welfare implications of our empirical findings.

³⁷Specifically, we sum up the total number of encounters of a given type and the total spending on those encounters across our study population and divide total spending by total encounters to get a per encounter average “cost”. The results are: \$1,607 for a hospital day, \$197 for an ED visit, \$147 for a SNF day, and \$79 for a doctor visit.

³⁸As discussed above, many of these health measures are annualized to account for the fact that not everyone was enrolled in Medicaid for the full year in 2015. The share enrolled for the full year is (as expected) balanced across control and intervention arms (see Appendix Table A2). Therefore, not surprisingly, we find in Appendix Table A15 that if we limit the analysis to the subset of study participants enrolled in Medicaid for the full year in 2015, the results remain qualitatively the same (although precision worsens).

7.1 Parameterizing the model

To simplify the parameterization of the model, we collapse the distribution of benefits to be only one of two possible levels: either the minimum benefit of \$16 per month, or \$178 per month (the mean benefit for the approximately 80 percent of control group enrollees who do not receive the minimum). We assume these two benefit levels correspond to the l and h types in the model. In other words, conditional on their application being accepted, l types receive the minimum benefit and h types receive the higher benefit level.³⁹ These assumptions imply that type h enrollees receive \$4,272 during the first 24 months of enrollment, while type l enrollees receive \$384 over 24 months. Recall that after 24 months, the individual must re-certify their eligibility; average lifetime benefits are therefore presumably greater than the 24-month amount, but may not extend indefinitely; moreover, additional private costs must be incurred to maintain them. For simplicity, we assume benefits last only 24 months; this is a conservative assumption since, as we will see, higher expected benefits among enrollees imply larger misperceptions about the probability of successfully enrolling.

To calculate expected benefits from applying, we assume that the probability of rejection is 0.25 for both types (the rejection rate for the controls). We treat this rejection rate as exogenous to the intervention, given that we found no evidence of an effect of either intervention on rejection rates of application. Thus, expected benefits conditional on applying ($\pi_j B_j$) are about \$3,200 for the h types and about \$290 for l types. This calculation assumes that SNAP benefits are valued dollar for dollar by beneficiaries.⁴⁰

The model in Section 3 underscores that a key determinant of the welfare analysis will be whether the neoclassical benchmark ($\epsilon = 0$) is a reasonable assumption. It is difficult to definitively reject this neoclassical benchmark. Given that applying takes an estimated five hours (Ponza et al. 1999), if we (generously) assume the value of time for this low-income elderly population is roughly twice the minimum wage of \$7.25 per hour, this implies the private (time) cost of applying is about \$75. With no misperceptions, rationalizing the decision not to apply therefore requires a non-time cost of applying of roughly \$3,100 for an h type. If we model stigma as a participation cost (Moffitt 1983), one way to rationalize the decision of non-applicants is to say that they experience stigma costs of participation that are about forty times larger than their transactional costs of applying. For an l type with no misperception of the probability an application is accepted, the implied non-time cost of applying is roughly \$200. Under this rational benchmark, the Information Only intervention overall unambiguously reduces private and social welfare (Proposition 1); in addition, the poor targeting properties of the intervention are irrelevant for private welfare and increase social welfare for a given change in application rate, due to reduction in the average fiscal cost of enrollees in our setting (Proposition 2, with g_j constant across types).

However, our reading of the evidence suggests that individuals under-estimate the probability

³⁹In other words, $\pi_{lL} = \pi_{hH} = 0.75$ and $\pi_{lH} = \pi_{hL} = 0$, so there is no “type uncertainty” in this setup, only uncertainty over application acceptance probability.

⁴⁰While recent evidence by Hastings and Shapiro (2017) calls into question the standard assumption that SNAP benefits are fungible with cash for large majority of SNAP-eligible households, it is not immediately clear whether this implies that SNAP benefits are valued more or less than cash at the margin.

their application is accepted (i.e., $\epsilon < 0$) and hence expected benefits from applying. As noted previously, existing survey evidence suggests that lack of awareness of expected benefits - e.g., underestimating expected benefits - is a primary barrier to participation among eligible non-participants (Bartlet et al., 2004); one interpretation of our “Information Only” intervention is that it reduces such misperception. In addition, the substantial increase in applications and enrollment from a reminder postcard in the Information Only intervention suggests some form of inattention, lack of awareness, or forgetfulness; i.e., individual application decisions may not be privately optimal, as implied by the neoclassical benchmark.

To calibrate the magnitude of the misperceptions, we assume that the time cost is the only cost of application. To rationalize non-participation with the time cost estimates above requires $\epsilon_h = -0.98$ and $\epsilon_l = -0.75$. In other words, for a type h individuals with a 75 percent chance of enrolling after applying, the only way to rationalize their not applying for benefits is that their misperceptions are so great that they perceive virtually no chance (less than 2 percent) of enrolling in program, or alternatively that they are completely ignorant of the program. For type l individuals with a 75 percent chance of enrolling after applying not to apply, they must perceive a 25 percent or lower chance of success. Thus $\epsilon_h < \epsilon_l < 0$.

7.2 Normative Findings

Proposition 2 indicates that with $\epsilon_h < \epsilon_l < 0$, a progressive benefit formula, and constant non-benefit fiscal externalities (i.e., g_j constant across types), then our finding that the interventions decrease targeting bodes poorly for their welfare impacts. However, this is merely a qualitative comparative static result. Even with $\epsilon_h < \epsilon_l < 0$, the targeting effects of the intervention are neither necessary nor sufficient to sign the overall social welfare impact of the intervention. Consider an intervention that improves targeting (with $\epsilon_h < \epsilon_l < 0$). Proposition 2 tells us that, all else equal, this improvement in targeting is good for the social welfare effects of the intervention. Still, the overall social welfare effect may be negative, for the negative externality from the public application processing costs and expenditures on benefits may outweigh the private welfare gains to individuals with misperceptions (see equation (2)). Likewise, if the intervention decreases targeting - thereby reducing the social welfare benefits from the intervention under our sufficient conditions - it may still increase social welfare overall if the private welfare gains to individuals with misperceptions outweigh the public costs.

Proposition 1 tells us that to make quantitative statements about the social welfare impact of the intervention - i.e., $\frac{dW}{dT}$ - we need estimates of g_j , $\pi_j B_j$, $\frac{dA_j}{dT}$ and $\mu_j \equiv -u'(y_j)\epsilon_j(\pi_j B_j)$ (for $j = \{h, l\}$). We assume that the fiscal externality is the same across types (so that $g_h = g_l = g$), and using Isaacs (2008) we estimate $g \sim \$267$ (see section 6.1 for more details on this calculation). Lastly, we parameterize $\pi_h B_h \sim \$3,200$, $\pi_l B_l \sim \$290$, $\epsilon_h \sim -0.98$, and $\epsilon_l \sim -0.75$.

The impact of the intervention on applications of each type $\frac{dA_j}{dT}$ comes directly from the experiment. Table 2 shows directly the increase in applications - for the Information Only intervention, $\frac{dA}{dT} = 0.07$ and for the Information Plus Assistance intervention, $\frac{dA}{dT} = 0.16$. Appendix Table A6

shows that, for each intervention, 44% of the marginal enrollees are l types (i.e., 44% of the compliers receive the minimum benefit level of \$16); this represents a decrease in targeting relative to the inframarginal enrollees (i.e., the always takers) for whom, Table 2 shows, only about 20% are l type. Given our assumption of a common, 25 percent, rejection rate for both types, this suggests that for the Information Only intervention, $\frac{dA_l}{dT} = .03$ and $\frac{dA_h}{dT} = .04$, and for the Information Plus Assistance intervention, $\frac{dA_l}{dT} = .07$ and $\frac{dA_h}{dT} = .09$.

Finally, we still require an estimate of the marginal utility of consumption for recipients (as well as our assumption of a particular set of social welfare weights on them - in our case, utilitarian). Consumption is notoriously difficult to measure in a low-income population, and the marginal utility of that consumption is an even more challenging object, as are social welfare weights. Given this difficulty, we instead follow Hendren (2016), who offers a way to assess redistributive programs without having to make assumptions about either individual utility functions or social welfare functions. He defines the marginal value of public funds ($MVPF$) from a marginal expansion of a program (or in our case, an intervention into that program) as the ratio of marginal benefits to marginal costs. We adapt Hendren (2016) to our setting and derive:

$$MVPF^{Information\ Only} = \frac{-\epsilon_h(\pi_h B_h) \frac{dA_h}{dT} - \epsilon_l(\pi_l B_l) \frac{dA_l}{dT}}{(\pi_h B_h + g_h) \frac{dA_h}{dT} + (\pi_l B_l + g_l) \frac{dA_l}{dT}}.$$

This formula bears a strong resemblance to $\frac{dW}{dT}^{Information\ Only}$ (see equation (3)). Beyond the fact that one is expressed as a ratio and the other as a difference, the key distinction is that the private welfare changes (the numerator) are expressed as a money metric, rather than multiplied by the marginal utility of consumption. We therefore have rough estimates of all the elements we need to evaluate this expression, and these suggest:

$$MVPF^{Information\ Only} = \frac{0.98(\$3,200)0.04 + 0.75(\$290)0.03}{(\$3,200 + \$267)0.04 + (\$290 + \$267)0.03} = 0.85$$

An $MVPF$ estimate of 0.85 suggests that for every dollar spent on the intervention (in the form of benefits and processing costs), low income recipients receive about 85 cents of benefits.⁴¹ An $MVPF$ below 1 is to be expected for a redistributive policy such as SNAP; redistribution inevitably involves some resource cost (Okun, 1975). A more natural benchmark is to compare this estimate to the $MVPF$ of other redistributive programs. Although we know of no elderly-specific estimates, it is interesting to see that this estimate is comparable to the estimate of the $MVPF$ for the Earned Income Tax Credit, which Hendren (2014) estimates to be about 0.9; this is higher than the likely $MVPF$ of public subsidies for health insurance for low-income adults (Finkelstein, Hendren and Shepard 2017), as well as the $MVPF$ of the SNAP program for the non-elderly, which Hendren (2014) estimates has an $MVPF$ of roughly 0.5-0.7. In other words, an information intervention about likely eligibility for SNAP among an elderly population transfers more resources to low-

⁴¹This calculation assumes that the information intervention is itself costless. Accounting for the intervention costs (\$1 per outreach, or approximately \$7 for the 15 percent of the intervention arm who applied) in the denominator, however, has no noticeable effect on the calculation.

income beneficiaries per dollar of public expenditure than the SNAP program for the non-elderly or subsidies for public health insurance for low-income adults. Of course, this calculation assumes away some potential behavioral responses to the intervention (such as decreased labor force participation) which could increase the fiscal externality on the government. This assumption seems empirically reasonable in the context of our experimental interventions, but of course might differ if these interventions were implemented at scale.

We can perform a similar analysis for the Information Plus Assistance intervention using the following extended formula:

$$MVPF^{Information\ Plus\ Assistance} = \frac{-\epsilon_h(\pi_h B_h) \frac{dA_h}{dT} - \epsilon_l(\pi_l B_l) \frac{dA_l}{dT} - (A_h + A_l + \frac{dA_h}{dT} + \frac{dA_l}{dT}) \frac{dc}{dT}}{(\pi_h B_h + g_h) \frac{dA_h}{dT} + (\pi_l B_l + g_l) \frac{dA_l}{dT}}.$$

The MVPF for the Information Plus Assistance intervention is the same as for the Information Only intervention, plus one additional term in the numerator, representing the welfare gain from reducing application costs for both the infra-marginal and marginal applicants. The term dc/dT is the (money-metric) change in application costs from the intervention, and it is scaled by the number of total applicants (both infra-marginal and marginal) of either type (i.e., this is the overall application rate in this treatment arm). The money metric term dc/dT replaces the $u'(y_j)$ terms multiplying the infra-marginal applicants in the expression for $\frac{dW}{dT}^{Information\ Plus\ Assistance}$ (see equation (3)).

Assuming that the application costs are costlessly reduced - which would correspond to removing some pre-existing barrier or ordeal - the MVPF is unambiguously higher for the Information Plus intervention than the Information Only one. If the intervention costlessly eliminated private application costs (i.e., reducing them from \$75 per application to zero), this would increase the MVPF from 0.85 in the Information Only intervention to 0.88. If we allow for BDT's cost per application estimate of \$45 (\$60 per enrollee, adjusted for the acceptance rate), the MVPF for the Information Plus Assistance Intervention would fall to 0.87.

To see the role that targeting plays in affecting MVPF, we calculation the MVPF in the Information Only intervention separately for each type:

$$MVPF_h^{Information\ Only} = \frac{-\epsilon_h(\pi_h B_h) \frac{dA_h}{dT}}{(\pi_h B_h + g_h) \frac{dA_h}{dT}} = \frac{0.98(\$3,200)0.04}{(\$3,200 + \$267)0.04} = 0.90$$

$$MVPF_l^{Information\ Only} = \frac{-\epsilon_l(\pi_l B_l) \frac{dA_l}{dT}}{(\pi_l B_l + g_l) \frac{dA_l}{dT}} = \frac{0.75(\$290)0.03}{(\$290 + \$267)0.03} = 0.39$$

As Proposition 2 predicts, given our estimate of $\epsilon_h < \epsilon_l < 0$, the MVPF of the intervention is larger for h types. The difference is substantial, highlighting the potential welfare gains in our setting from policies that are especially effective at targeting high-benefit types. Policies that primarily enroll low-benefit types appear to have quite low MVPF (~ 0.4). In other words, if those deterred

by barriers were exclusively the less needy, our interventions would have looked substantially worse.

8 Conclusion

Policymakers often advocate - and academics often study - interventions to increase take-up of public benefits. We provide a framework for analyzing the welfare impacts of such interventions. We apply this framework to the results of a randomized field experiment of interventions designed to increase SNAP take-up. The interventions were designed to reduce potential information barriers to enrollment as well as potential transaction cost barriers.

We found that both information and transaction costs are barriers to take-up. In the 9 months following the intervention, the Information Only intervention increased enrollment by 5 percentage points (or 83 percent relative to the enrollment rate among controls), while the Information Plus Assistance increased enrollment by 12 percentage points (a 200 percent increase relative to the controls). The impact of the treatments appears to be fully present by about 6 months; the time pattern of effects out to 23 months suggests that the treatments primarily generate new enrollment, rather than merely moving forward in time enrollment that would have happened anyway. A back of the envelope calculation suggests that the Information Only treatment may be more “cost effective”, with an intervention cost of about \$20 per new enrollee, compared to about \$60 per new enrollee for the Information Plus Assistance intervention.

We also find that reducing informational or transactional barriers decreases targeting: the marginal applicants and enrollees from either intervention are less needy than the average enrollees in the control group. The average monthly SNAP benefit (which is based on a progressive formula) is 20 to 30 percent lower among enrollees in either intervention arm relative to enrollees in the control group. In addition, relative to the control group, applicants and enrollees in either intervention arm are in better health, more likely to be white, and more likely to have English as their primary language. The finding that barriers to take-up deter relatively less needy individuals from enrolling is consistent with neoclassical theories of ordeal mechanisms (e.g., Nichols et al., 1971, Nichols and Zeckhauser 1982, Besley and Coate 1992).

The framework we developed highlights that normative implications depend critically on whether individuals have accurate beliefs about the expected benefits from applying, as well as what types of individuals have greater misperceptions. We present several pieces of evidence that are consistent with standard behavioral models (e.g., Mullainathan and Shafir 2013) in which individuals under-estimate expected benefits from applying, with this under-estimation greater among needier individuals. Under the assumptions in our setting, this is a sufficient condition for a decrease in targeting to decrease the social welfare gains from intervention. Nevertheless, given the magnitude of the misperceptions we estimate (particularly for the neediest individuals), our estimates suggest that increasing awareness and reducing transaction costs for elderly individuals eligible for SNAP are both relatively cost-effective ways to redistribute to this low-income population.

The framework we developed also clarifies conditions under which the targeting properties of

an intervention based on observable characteristics such as poverty may be informative about the likely welfare impact of the intervention. These conditions suggest the importance of measuring additional empirical objects - specifically, the size of any misperceptions across individuals with different observable characteristics as well as the size of the fiscal externality across these individuals - in order to draw normative inferences from targeting results. This should hopefully be useful for analyzing the welfare impacts of other interventions designed to increase take-up of social benefits.

9 References

Alberto Abadie, 2002. "Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models", *Journal of the American Statistical Association* 97(457), pp. 284-292.

Alberto Abadie, 2003. "Semiparametric instrumental variable estimation of treatment response models", *Journal of Econometrics* 113, pp. 231-263.

Aaronson, B., 2011. Number of Texans Receiving Food Stamps Up Sharply Amid Recession [WWW Document]. *Tex. Trib.* URL <http://www.texastribune.org/library/data/texas-supplemental-nutrition-assistance-program/> (accessed 2.3.15).

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

Alcott, Hunt and Michael Greenstone. 2017. "Measuring The Welfare Effects of Residential Energy Efficiency Programs." NBER Working Paper 23386.

Armour, Philip. forthcoming. "The Role of Information in Disability Insurance Application: An Analysis of the Social Security Statement Phase-In." *American Economic Journal: Economic Policy*.

Angrist, Joshua D. and Jörn-Steffen Pischke, 2009. "Mostly Harmless Econometrics: An Empiricist's Companion.", Princeton University Press.

Baily, Martin N. 1978. "Some Aspects of Optimal Unemployment Insurance", *Journal of Public Economics* 10(3), pp. 379-402.

Barr, Andrew and Sarah Turner. "A Letter and Encouragement: Does Information Increase Post-Secondary Enrollment of UI Recipients?" *American Economic Journal: Economic Policy*.

Benefits Data Trust, 2016. About: Our Model.

Bartlett, Susan, Nancy Burstein, and William Hamilton. 2004. "Food Stamp Access Study: Final Report". United States Department of Agriculture, Economic Research Service, November. <https://www.ers.usda.gov/publications/pub-details/?pubid=43407> (last accessed July 31, 2017).

Bertrand, M., Mullainathan, S., Shafir, E., 2004. A behavioral-economics view of poverty. *Am. Econ. Rev.* 419-423.

Besley, T., Coate, S., n.d. Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs. *Am. Econ. Rev.* 82, 249-61.

Bettinger, E.P., Long, B.T., Oreopoulos, P., Sanbonmatsu, L., 2012. The role of application

assistance and information in college decisions: Results from the H&R Block FAFSA experiment. *Q. J. Econ.* 127, 1205–1242.

Bhargava, S., Manoli, D., 2015. Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment. *Am. Econ. Rev.* 105, 3489–3529. doi:10.1257/aer.20121493

Bhargava, S, Loewenstein, G and J. Sydnor. 2017. “Choose to Lose: Health Plan Choices from a Menu with Dominated Options.” *Quarterly Journal of Economics* August. Volume 132 (3): 1319-1372.

Cutler, D. and J. Gruber. 1996. “Does Public Insurance Crowd out Private Insurance?” *Quarterly Journal of Economics* 111(2): 391-430.

Center on Budget and Policy Priorities. 2015. “Pennsylvania: Supplemental Nutrition Assistance Program, 2015.”

Center on Budget and Policy Priorities. 2016. “SNAP Online: A review of State Government SNAP Websites.” <http://www.cbpp.org/research/food-assistance/snap-online-a-review-of-state-government-snap-websites>. Last accessed, June 29 2017.

Center on Budget and Policy Priorities. 2017. “SNAP Helps Millions of Low-Income Seniors.” <https://www.cbpp.org/research/food-assistance/snap-helps-millions-of-low-income-seniors>

Chetty, Raj. 2006. "A General Formula for the Optimal Level of Social Insurance", *Journal of Public Economics* 90(10-11), pp. 1879-1901.

Cunningham, K., 2010. State Trends in Supplemental Nutrition Assistance Program Eligibility and Participation Among Elderly Individuals. *Mathematica Policy Research*.

Congressional Budget Office. 2013. “Federal Means-Tested Programs and Tax Credits – Infographic.” <https://www.cbo.gov/publication/43935> Last accessed, June 29, 2017.

Currie, J., 2006. The Take Up of Social Benefits. In Alan Auerbach, David Card, And John Quigley (Eds). *Poverty, the Distribution of Income, and Public Policy* (New York: Russell Sage), 80-148.

Currie, J., 2003. US food and nutrition programs, in: *Means-Tested Transfer Programs in the United States*. University of Chicago Press, pp. 199–290.

Daponte, B.O., Sanders, S., Taylor, L., 1999. Why do low-income households not use food stamps? Evidence from an experiment. *J. Hum. Resour.* 612–628.

Deshpande, Manasi and Yue Li. 2017. “Who is Screened Out? Application Costs and the Targeting of Disability Programs.” Unpublished Miemo. <https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbntxtZGVzaHBhbmRlZWVbnxneDo3ZTU1MGYyNTdhMWFhNmY3>

Einav, L., Finkelstein, A. and Cullen, M.R. 2010. Estimating Welfare In Insurance Markets Using Variation in Prices. *Q. J. Econ.*

Eslami, E., 2014. Trends in Supplemental Nutrition Assistance Program Participation Rates: Fiscal Year 2010 to Fiscal Year 2012.

Finkelstein A, Gentzkow M, Williams H. 2016. Sources of geographic variation in health care: evidence from patient migration. *Q J Econ*;131:1681-726.

Finkelstein A, Hendren N, Shepard M. 2017. Subsidizing health insurance for low income adults: evidence from Massachusetts. NBER Working Paper 23668.

Ganong, P., Liebman, J.B., 2013. The Decline, Rebound, and Further Rise in SNAP Enrollment: Disentangling Business Cycle Fluctuations and Policy Changes (Working Paper No. 19363). National Bureau of Economic Research.

Gray, K., Fisher, S., Lauffer, S., 2016. Characteristics of Supplemental Nutrition Assistance Program Households: Fiscal Year 2015 (No. Report No. SNAP-16-CHAR). United States Department of Agriculture.

Gruber, J., Levine, P., Staiger, D., 1999. Abortion Legalization and Child Living Circumstances: Who is the Marginal Child? *Q. J. Econ.* 114, 263–292.

Guyton, John, Dayanand Manoli, Brenda Schafer and Michael Sebastiani. 2016. “Reminders & Recidivism: Evidence from Tax Filing & EITC Participation Among Low-Income Nonfilers.” NBER Working Paper 21904.

Hastings, Justine and Jesse Shapiro. 2017. “How Are SNAP Benefits Spent? Evidence from a Retail Panel.” NBER Working Paper 23112.

Hendren, N. (2014). The inequality deflator: Interpersonal comparisons without a social welfare function. Working Paper .

Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy* 30 (1), 51–89.

Hoynes and Schanzenbach. 2016. U.S. Food and Nutrition Programs. in Robert Moffitt (ed) “Means Tested Transfer Programs in the United States, Volume II.” University of Chicago Press.

Hu, W., 2014. Rations Reduced as Demand Grows for Soup Kitchens. *N. Y. Times*.

Isaacs, Julia. 2008. “The Costs of Benefit Delivery in the Food Stamp Program.” USDA Contractor and Cooperator Report No. 39. https://www.brookings.edu/wp-content/uploads/2016/06/03_food_stamp_isaacs.pdf. Last accessed June 29, 2017

Kauff, J., Dragoset, L., Clary, E., Laird, E., Makowsky, L., Samaa-Miller, E., 2014. Reaching the Underserved Elderly and Working Poor in SNAP: Evaluation Findings from the Fiscal Year 2009 Pilots. Mathematica Policy Research.

Kleven, H and W. Kopczuk. 2011. “Transfer Program Complexity and the Take Up of Social Benefits.” *American Economic Journal: Economic Policy* 3(1), 54-90.

Kleven, H. 2018. “Sufficient Statistics Revisited.” Working paper. https://www.henrikkleven.com/uploads/3/7/3/1/37310663/kleven_sufficientstats_march2018.pdf

Lewin Group. 2012. “Evaluating Encounter Data Completeness.” <https://www.ccwdata.org/documents/10280/19002254/evaluating-encounter-data-completeness.pdf>

Madrian, Brigitte and Dennis Shea. 2001. “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior” *Quarterly Journal of Economics*. CXVI (4): 1149-1187.

Mani, A., Shafir, E., Mullainathan, S., Zhao, J., 2013. Poverty impedes cognitive function. *Science* 341, 976–980.

Manoli, Dayand and Nicholas Turner. 2014. “Nudges and Learning: Evidence from Informational Interventions for Low-Income Taxpayers.” NBER Working Paper 20718.

Meyer, Bruce and Nikolas Mittag. 2015. "Using Linked Survey and Administrative Data to Better Measure Income: Implications for Poverty, Program Effectiveness, and Holes in the Safety Net." NBER Working Paper 21676

Mills, Gregory, Tracy Vericker, Heather Koball, Kyle Lippold, Laura Wheaton, and Sam Elkin, 2014. "Understanding the Rates, Causes, and Costs of Churning in the Supplemental Nutrition Assistance Program (SNAP)," United States Department of Agriculture, (November).

Moffitt, Robert. 1983. "An Economic Model of Welfare Stigma." *American Economic Review* 73 (5): 1023-1035.

Mullainathan, S., Shafir, E., 2013. *Scarcity: Why having too little means so much*. Macmillan.

Nichols, A.L., Zeckhauser, R.J., 1982. Targeting transfers through restrictions on recipients. *Am. Econ. Rev.* 372-377.

Nichols, D., Smolensky, E., Tideman, T., 1971. Discrimination by Waiting Time in Merit Goods. *Am. Econ. Rev.* 61, 312-323.

Okun, A. 1975. *Equality and Efficiency*. Brookings Institution Press.

Pennsylvania Department of Human Services. Online. "Supplemental Nutrition Assistance Program Policy Manual". <http://services.dpw.state.pa.us/oimpolicymanuals/snap/> (last accessed June 28, 2017).

Ponza, Michael, James C. Ohls, Lorenzo Moreno, Amy Zambrowski, and Rhoda Cohen. 1999. "Customer Service in the Food Stamp Program." *Mathematica Policy Research Inc.*, Reference Number 8243-140.

Rifkin, Ryan and Aldebaro Klautau. 2004. "In Defense of One-Vs-All Classification", *Journal of Machine Learning Research*, 5 (January):101-141.

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

Schanzenbach, Diane. 2009. "Experimental Estimates of the Barriers to Food Stamp Enrollment." *Institutet for Research on Poverty, University of Wisconsin-Madison*, Discussion Paper no. 1367-09.

Song Y, Skinner J, Bynum J, Sutherland J, Wennberg JE, Fisher ES. 2010. Regional variations in diagnostic practices. *New Engl J Med*;363:45-53.

Spinnewijn, J. 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs." *Journal of the European Economic Association* 13(1): 130-167. February.

U.S. Congressional Budget Office, 2013. Growth in means-tested programs and tax credits for low-income households. U.S. Department of Agriculture: Food and Nutrition Service, 2015. Supplemental Nutrition Assistance Program (SNAP): Eligibility [WWW Document]. URL <http://www.fns.usda.gov/snap/eligibility> (accessed 2.1.15). U.S. Department of Health and Human Services, 2012. Fiscal year 2011 TANF financial data.


U.S. Department of Agriculture. 2016. The food assistance landscape: FY 2015 annual report.

U.S. Department of Health and Human Services. 2016. FY 2015 Federal TANF & state MOE financial data.

U.S. Internal Revenue Service. 2016. EITC calendar year report.

U.S. Social Security Administration. 2016. Social Security Income program 2016 technical materials. February.

Figure 1: Time pattern of enrollment responses



NOTE: Figure shows, by month, the (cumulative) estimated treatment effects on enrollment (relative to the control) for the Information Only arm and the Information Plus Assistance arm. 95 percent confidence intervals on these estimates are shown in the dashed light gray lines.

Table 1: Description of Study Population

	Original Outreach List		After Exclusions		Study Population (5)
	(1)	List, After Exclusions (2)	Receiving SNAP (3)	Not Receiving SNAP (4)	
	Observations (N)	229,584	143,923	84,038	
Panel A - Demographics					
Age (as of October 31, 2015)	72.91	70.45	69.77	71.42	68.83
Share Age 80+	0.27	0.18	0.15	0.23	0.16
Male	0.35	0.36	0.36	0.36	0.38
Share White ^a	0.71	0.79	0.79	0.79	0.75
Share Black ^a	0.17	0.10	0.11	0.07	0.08
Share Primary Language not English	0.04	0.03	0.03	0.03	0.04
Share Living in Philadelphia	0.18	0.00	0.00	0.00	0.00
Share Living in Pittsburgh	0.05	0.07	0.07	0.06	0.06
Share Last Medicaid Spell Starting before 2011	0.45	0.47	0.55	0.36	0.33
Share Enrolled in Medicaid for 2015 Full Year	0.83	0.84	0.89	0.77	0.73
Panel B - (Annual) Health Care Measures, 2015					
Total Health Care Spending (\$) ^b	18,347	7,683	6,036	9,995	11,838
Number of Hospital Days	5.41	1.51	1.24	1.88	2.16
Number of ER Visits	0.41	0.41	0.41	0.40	0.50
Number of Doctor Visits	6.25	5.87	5.97	5.74	7.11
Number of SNF Days	66.23	1.57	0.85	2.58	2.67
Number of Chronic Conditions	6.50	4.93	5.08	4.70	5.45

Notes: Observations correspond to a sample of Medicaid enrollees using data from Pennsylvania Dept. of Human Services (DHS). Column (1) shows the initial outreach list of individuals aged 60 and over enrolled in Medicaid as of October 31, 2015. In column (2) we make two exclusions from this list: we exclude all individuals enrolled in the Long-Term Care Medicaid program and individuals with an address in Philadelphia City. Columns 3 and 4 partition the resulting sample in column 2 into those in "households" enrolled in SNAP and those not, respectively, where a "household" is defined as individuals on the outreach list sharing the same last name and address; recall that SNAP is a household-level benefit. Column (5) shows the final study population, which is a subset of the individuals not enrolled in SNAP in column (4); we excluded all individuals in column (4) to whom BDT had previously sent outreach materials and randomly selected one individual from each "household". All data come from Medicaid administrative data; health care spending and utilization data come from the 2015 Medicaid claims files and all measures are annualized for individuals with less than a full year of Medicaid enrollment; see Appendix B for more details.

^aOmitted category is other or missing race.

^bTotal spending is truncated at twice 99.5th percentile of study population, which is 371,620 (99.5th percentile in study population is 185,810). Amounts greater than the threshold are set to missing.

Table 2: Behavioral Responses to “Information Only” and “Information Plus Assistance”

	Control	Information Only	Information Plus Assistance	P Value of Difference (Column 2 vs 3)
	(1)	(2)	(3)	(4)
SNAP Enrollees	0.058	0.105 [0.000]	0.176 [0.000]	[0.000]
SNAP Applicants	0.077	0.147 [0.000]	0.238 [0.000]	[0.000]
SNAP Rejections among Applicants	0.233	0.266 [0.119]	0.255 [0.202]	[0.557]
Callers	0.000	0.267 [0.000]	0.301 [0.000]	[0.000]
Adjusted Callers	0.000	0.289 [0.000]	0.301 [0.000]	[0.156]
SNAP Applicants among Non-Callers	0.077	0.086 [0.063]	0.081 [0.324]	[0.363]
SNAP Applicants among Callers	0.000	0.313 [0.000]	0.602 [0.000]	[0.000]
SNAP Enrollees among Non-Callers	0.058	0.061 [0.442]	0.059 [0.713]	[0.688]
SNAP Enrollees among Callers	0.000	0.226 [0.000]	0.450 [0.000]	[0.000]
Observations (N)	10,630	5,314	10,629	

Notes: Columns 1 through 3 shows means by intervention arm with the p-value (relative to the control arm) in [square brackets]. Column 1 shows the control. Column 2 shows the Information Only arm (for the two equally-sized pooled sub-treatments). Column 3 shows the Information Plus Assistance arms (weighted so that the two pooled sub-treatments received equal weight). Column 4 reports the p-value of the difference between the Information Plus Assistance and Information Only treatment arms. All outcomes are binary rates measured during the nine months from the initial mail date. All p-values are based on heteroskedasticity-robust standard errors. Callers are measured for the relevant call number and are therefore mechanically zero for the control; see text for a description of the adjusted caller rate.

Table 3: Behavioral Responses to “Information Only” Intervention with and without reminders

	Control	Information Only Standard	Information Only No-Postcard	P Value of Difference (Column 2 vs 3)
	(1)	(2)	(3)	(4)
SNAP Enrollees	0.058	0.112 [0.000]	0.092 [0.000]	[0.016]
SNAP Applicants	0.077	0.151 [0.000]	0.120 [0.000]	[0.001]
SNAP Rejections among Applicants	0.233	0.224 [0.751]	0.216 [0.536]	[0.777]
Callers	0.000	0.278 [0.000]	0.212 [0.000]	[0.000]
Adjusted Callers	0.000	0.300 [0.000]	0.234 [0.000]	[0.000]
SNAP Applicants among Non-Callers	0.077	0.089 [0.079]	0.074 [0.593]	[0.071]
SNAP Applicants among Callers	0.000	0.311 [0.000]	0.295 [0.000]	[0.524]
SNAP Enrollees among Non-Callers	0.058	0.064 [0.284]	0.054 [0.492]	[0.172]
SNAP Enrollees among Callers	0.000	0.237 [0.000]	0.234 [0.000]	[0.921]
Observations (N)	10,630	2,657	2,658	

Notes: Columns 1 through 3 shows means by intervention arm with the p-value (relative to the control arm) in [square brackets]. Column 1 shows the control. Column 2 shows the “standard” Information Only intervention (see Appendix Figure A2; this “standard” intervention is was half of the sample shown in Table 2 column (3) for the pooled Information Only analysis). Column 3 shows the results of the Information Only intervention without the reminder postcard; the outreach materials are otherwise identical to those in Appendix Figure A2. Column 4 reports the p-value of the difference between the standard Information Only intervention and the Information Only intervention without the reminder postcard. All outcomes are binary rates measured during the nine months from the initial mail date. All p-values are based on heteroskedasticity-robust standard errors. Callers are measured for the relevant call number and are therefore mechanically zero for the control; see text for a description of the adjusted caller rate.

Table 4: Enrollee Monthly Benefits and Predicted Benefits

	Control	Information Only	Information Plus Assistance	P Value of Difference (Column 2 vs 3)
	(1)	(2)	(3)	(4)
Benefit Amount	145.94	115.38 [0.000]	101.32 [0.000]	[0.013]
Share \$16 Benefit	0.178	0.299 [0.000]	0.357 [0.000]	[0.012]
Share \$194 Benefit	0.191	0.157 [0.131]	0.143 [0.009]	[0.421]
Share \$357 Benefit	0.055	0.050 [0.681]	0.039 [0.115]	[0.285]
Share Missing Benefit	0.073	0.043 [0.025]	0.028 [0.000]	[0.139]
Predicted Benefit for Enrollees w/ Actual Benefit	140.20	112.49 [0.000]	102.93 [0.000]	[0.086]
Predicted Benefit for All Enrollees	138.65	114.01 [0.000]	104.03 [0.000]	[0.068]
Share of Enrollees in Household Size of 1	0.657	0.714 [0.038]	0.760 [0.000]	[0.036]
Benefit Amount for Enrollees in Household Size of 1	116.97	93.35 [0.000]	85.82 [0.000]	[0.134]
Observations (N)	613	559	1,861	

Notes: Sample is individuals who enrolled in the 9 months after their initial mailing. Columns 1 through 3 shows means by intervention arm with the p-value (relative to the control arm) in [square brackets] for SNAP enrollees. Column 1 shows the control. Column 2 shows the Information Only arm (with the two equally-sized sub-treatments pooled). Column 3 shows the Information Plus Assistance arms (weighted so that the two pooled sub-treatments received equal weight). Column 4 reports the p-value of the difference between the Information Plus Assistance and Information Only treatment arms. See text for a description of the predicted benefits. All p-values are based on heteroskedasticity-robust standard errors. N reports the sample size of enrollees.

Table 5: Demographic and Health Characteristics: Applicants and Enrollees

	Applicants				Enrollees			
	Means			P Value Info Plus Assistance vs Info Only	Means			P Value Info Plus Assistance vs Info Only
	Control	Info Only	Info Plus Assistance		Control	Info Only	Info Plus Assistance	
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Panel A - Predicted Benefits								
Predicted Benefits	148.26	125.65 [0.000]	115.36 [0.000]	[0.037]	138.65	114.01 [0.000]	104.03 [0.000]	[0.068]
Panel B - (Annual) Health Care Measures, 2015								
Total Health Care Spending (\$) ^b	9,424	8,605 [0.517]	8,334 [0.300]	[0.781]	10,238	9,532 [0.661]	8,603 [0.208]	[0.459]
Total Number of Visits and Days	13.33	11.67 [0.331]	9.92 [0.018]	[0.166]	14.79	10.90 [0.058]	9.92 [0.008]	[0.467]
Weighted Total Number of Visits and Days	4,661	3,273 [0.128]	2,818 [0.022]	[0.442]	5,407	3,288 [0.064]	2,779 [0.011]	[0.461]
Number of Chronic Conditions	6.21	5.55 [0.094]	5.27 [0.006]	[0.383]	6.54	5.43 [0.019]	5.37 [0.005]	[0.875]
Panel C - Demographics								
Share Age 80+	0.06	0.11 [0.001]	0.14 [0.000]	[0.042]	0.07	0.12 [0.005]	0.14 [0.000]	[0.085]
Male	0.41	0.40 [0.983]	0.38 [0.232]	[0.250]	0.39	0.42 [0.446]	0.38 [0.444]	[0.104]
Share White ^a	0.67	0.73 [0.005]	0.74 [0.000]	[0.554]	0.71	0.78 [0.004]	0.78 [0.001]	[0.958]
Share Black ^a	0.10	0.08 [0.103]	0.11 [0.577]	[0.011]	0.11	0.07 [0.011]	0.10 [0.833]	[0.004]
Share Primary Language not English	0.08	0.06 [0.141]	0.04 [0.000]	[0.012]	0.06	0.05 [0.242]	0.03 [0.002]	[0.067]
Share Living in Pittsburgh	0.05	0.06 [0.385]	0.07 [0.066]	[0.459]	0.05	0.06 [0.374]	0.07 [0.028]	[0.310]
Share Last Medicaid Spell Starting before 2011	0.25	0.30 [0.022]	0.29 [0.017]	[0.704]	0.26	0.33 [0.009]	0.31 [0.026]	[0.348]
Observations (N)	817	781	2,519		613	559	1,861	

Notes: Columns 1 - 3 and 5 - 7 show means by intervention arm with the p-value (relative to the control arm) in [square brackets] for SNAP applicants who applied within 9 months of their initial mailing, and SNAP enrollees who enrolled within 9 months of their initial mailing, respectively. Column 1 and 5 show the control. Column 2 and 6 show the Information Only arms (with the two equally-sized sub-treatments pooled); columns 3 and 7 show the Information Plus Assistance arms (weighted so that the two pooled sub-treatments received equal weight). Columns 4 and 8 report the p-value of the difference between the Information Plus Assistance and Information Only treatment arms. All p-values are based on heteroskedasticity-robust standard errors.

^aOmitted category is other or missing race.

^bTotal spending is truncated at twice 99.5th percentile of study population, which is 371,620 (99.5th percentile in study population is 185,810). Amounts greater than the threshold are set to missing.