The Welfare Effects of Eligibility Expansions: Theory and Evidence from SNAP *

Jenna Anders†
Charlie Rafkin‡

November 2022

Abstract

We study the U.S. rollout of eligibility expansions in the Supplemental Nutrition Assistance Program. Using administrative data from the U.S. Department of Agriculture, we show that expanding eligibility raises enrollment among the inframarginal (always-eligible) population. Using an online experiment and an administrative survey, we find evidence that information frictions, rather than stigma, drive the new take-up. To interpret our findings, we develop a general model of the optimal eligibility threshold for welfare programs with incomplete take-up. Given our empirical results and certain modeling assumptions, the SNAP eligibility threshold is lower than optimal.

*We thank Hunt Allcott, Abhijit Banerjee, Judi Bartfeld, Leo Bursztyn, Clément de Chaisemartin, Raj Chetty, Jon Cohen, John Conlon, Zoë Cullen, Esther Duflo, Amy Finkelstein, Peter Ganong, Benny Goldman, Jon Gruber, Craig Gundersen, Basil Halperin, Emma Harrington, Nathan Hendren, Lisa Ho, Jeffrey Liebman, Stephen Morris, Whitney Newey, Ben Olken, Emily Oster, Amanda Pallais, Dev Patel, Jim Poterba, Indira Puri, Frank Schilbach, Amy Ellen Schwartz, Jesse Shapiro, Johannes Spinnewijn, Evan Soltas, Dmitry Taubinsky, and participants at workshops at Harvard and MIT for helpful discussions. We thank Rian Flynn and Amber Zheng for excellent research assistance. We thank Timothy Harris for sharing data on SNAP work requirement waivers. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 and Grant No. 1745303; by Harvard's Foundations of Human Behavior Initiative; and by the Harvard GSAS Professional Development Fund for PhD Students. The online experiment was pre-registered at the AEA RCT Registry under AEARCTR-0005566, and survey instruments are available at Rafkin's website. The experiment received exempt status from MIT's Committee on the Use of Humans as Experimental Subjects (#E1962) and Harvard's Institutional Review Board (#IRB20-0326).

†Department of Economics, Harvard University. janders@g.harvard.edu.
‡Department of Economics, MIT. crafkin@mit.edu.

1 Introduction

Social programs in the United States are characterized by incomplete take-up, and there is substantial heterogeneity in take-up across programs. Meanwhile, there is also heterogeneity in eligibility criteria across programs. In fact, in some social programs, such as the Supplemental Nutrition Assistance Program (SNAP, also known as food stamps), the eligibility threshold even varies across states. There is a suggestive positive correlation between U.S. welfare programs' eligibility thresholds and take-up: programs with less stringent income eligibility thresholds have higher take-up rates (Figure 1).

Regardless of the across-program correlation, the within-program relationship between take-up rates and eligibility is consequential. In the simplest model of how to set eligibility thresholds, policymakers trade off giving larger benefits to only the poorest people or spreading the benefit more thinly to a larger number of people. But if eligibility thresholds affect take-up within the eligible population, the policymaker no longer faces this basic trade-off alone. Targeting benefits only to the poorest households could decrease take-up for these groups. As a result, it is important to determine whether there is a causal relationship between the eligibility threshold and take-up of the already-eligible population.

Does the eligibility threshold affect take-up of social programs? If so, how does this phenomenon affect programs' optimal eligibility? In this paper, we provide novel evidence that the eligibility threshold affects take-up among low-income individuals who are always eligible for SNAP, regardless of the threshold. We explore the mechanisms underlying this take-up response using an online experiment and analysis of a government-commissioned survey on incomplete SNAP take-up. To interpret our findings, we propose a general model of welfare program participation that allows us to study optimal policy when the eligibility threshold endogenously affects take-up. The model makes precise how mechanisms — namely, stigma and incomplete information — affect welfare considerations, and we estimate the model empirically.

We focus on SNAP for several reasons. First, it is a large program (with an annual budget of about $70 billion) that forms an important part of the U.S. public assistance system. Second, SNAP eligibility rules are at the center of an ongoing public discussion.[1] Third, SNAP publishes anonymized public-use administrative data (the Department of Agriculture's Quality Control files), which we use to form our main outcome of log enrollment counts. The administrative data alleviate concerns that the results could reflect the mismeasurement of individuals' eligibility status or program participation reporting biases (Kreider et al., 2012; Meyer et al., 2015).

We begin by providing evidence that eligibility expansions in SNAP raised enrollment among the lowest-income individuals who are always SNAP-eligible. States can choose to expand SNAP eligibility standards beyond the federal minimum of 130% of the Federal Poverty Level (FPL). We focus on individuals at 50–115% of the FPL, a group eligible for SNAP in every state because they are poorer than the federal

[1] The Trump administration proposed eliminating state discretion in eligibility thresholds (Federal Register, 2019).

minimum eligibility. Leveraging an event-study design (using variation across states and years), we find that raising the eligibility threshold by 10 percentage points (pp) of the FPL (e.g., from 130% to 140%) boosts enrollment by over 1 percent among the inframarginal group that was always eligible for SNAP. Our setting also yields a clean placebo test: the policy change that permits states to change their SNAP eligibility threshold also gave other bureaucratic benefits to states, and we show that states which adopted the policy *without* expanding SNAP eligibility saw no increases in SNAP enrollment. As another way of benchmarking the magnitude, we find that for every person who joins SNAP because she becomes newly eligible, 0.9 inframarginal people join the program. We conduct a model-free cost-effectiveness exercise and find that the mechanical cost of raising the means test enough to increase inframarginal enrollment by 1 pp is $2.2 billion per year — about the same as the mechanical cost of increasing the SNAP benefit enough to achieve the same goal.

This take-up response among the inframarginal population is consistent with a small literature documenting a similar phenomenon for public health insurance programs, where it is called a "welcome-mat effect" or a "woodwork effect" (because already-eligible individuals appear "out of the woodwork" to take up the health program). We use the term "inframarginal effects" to avoid negative or positive connotations.

To further connect our findings to the literature on incomplete social program take-up (Currie, 2004; Bhargava and Manoli, 2015; Finkelstein and Notowidigdo, 2019), we next turn to uncovering the mechanisms underlying inframarginal effects. One hypothesis, motivated by models of social signaling (e.g., Bursztyn and Jensen, 2017), is that raising the income threshold could reduce SNAP stigma: with less stringent eligibility rules, taking up SNAP no longer conveys as much information about one's type. To test this hypothesis, we conduct an online experiment with a nationally representative sample of more than 2,000 participants. We provide truthful information about the eligibility threshold in *one* state to shock beliefs about the mean eligibility threshold *across* states. The experimental variation increases participants' beliefs about the share of individuals who are eligible for SNAP in the entire U.S. by 9 percentage points on average (standard error: 0.8 pp), and decreases an index of stigma by -0.050 standard deviations (SE: 0.027, $p = 0.061$). Effects on stigma are larger among people who are SNAP-eligible.

A second hypothesis is that relaxing the eligibility restrictions increases information about SNAP. To test this hypothesis, we analyze microdata from the Food Stamp Program Access Study (FSPAS), a nationally representative survey on SNAP awareness and stigma among both SNAP enrollees and non-enrollees conducted by the USDA (Bartlett et al., 2004). To our knowledge, this is the first academic analysis of this rich dataset on SNAP take-up mechanisms. Using the FSPAS, we identify demographic groups that are likely subject to SNAP awareness and stigma. We find that the demographic groups with the largest inframarginal effects are those with low levels of SNAP awareness and *not* the groups whose stigma is most sensitive to eligibility thresholds from the experiment. Thus, combining the FSPAS and the online experiment, we find that relaxing the eligibility threshold does reduce SNAP stigma, but information appears to play a larger role in the decisions of people who newly take-up.

To assess the quantitative importance of each mechanism, and determine the implications of inframarginal effects for social welfare, we propose a general economic framework for analyzing optimal eligibility in the presence of inframarginal effects. Individuals who are eligible for a welfare program take up the program benefit as long the benefit exceeds a private take-up cost (e.g., stigma) and they are aware of the program.[2] Both the cost and information (awareness) can depend on the eligibility threshold. The model emphasizes that the planner trades off (i) a standard redistributive motive in which she values giving a bigger benefit to people with higher welfare weights against (ii) a new motive, inframarginal effects, in which relaxing eligibility thresholds raises take-up.[3] We derive an optimality condition for the eligibility threshold in which our key empirical fact, the inframarginal effect, enters as an observable elasticity (Chetty, 2009; Kleven, 2021).

The model also explains the role of the two candidate mechanisms, stigma and information frictions. Similar to in a Baily (1978)-Chetty (2006) framework, the optimality condition features a fiscal externality of the inframarginal effects and recipients' willingness to pay (WTP) for a higher eligibility threshold. Recipients' WTP depends on why inframarginal effects exist. First, suppose that inframarginal effects are mostly driven by behavioral responses to changing costs (e.g., through stigma). Then, the new enrollees driving the inframarginal effects are just indifferent between taking up and not, so they do not value the eligibility expansion — a standard envelope condition. However, those who would have enrolled regardless now pay lower stigma costs to take up. The optimal eligibility threshold trades off the reduced stigma among inframarginals with the fiscal externality of new take-up. In this way, the model cleanly isolates two countervailing forces that govern the welfare effects of reducing stigma. On the other hand, suppose instead that inframarginal effects reflect improved awareness of the program. Then, the new take-up now confers first-order utility gains, since people who lacked awareness were not previously optimizing.

Our framework lets us conduct normative analysis about whether the planner should raise the eligibility threshold. We study when the "naïve" planner who ignores inframarginal effects but otherwise behaves optimally will set the eligibility threshold too low — or equivalently, the benefit size too high — relative to a "sophisticated" planner who is aware of inframarginal effects. We characterize a simple sufficient condition: the naïve planner will always set the threshold too low if information agents' take-up is weakly more elastic to a change in the threshold than stigma agents' take-up. Thus, the model yields a direct empirical test for whether the existence of inframarginal effects implies that the eligibility threshold should rise.

We proceed to implement this test using a model-based decomposition of the mechanisms. On the one hand, the experiment suggests that eligibility changes reduce stigma. On the other hand, the FSPAS

[2]While we focus on stigma costs, our framework permits any cost that depends on the eligibility threshold. Another possible cost embedded within our framework is uncertainty about eligibility, as in e.g. Kleven and Kopczuk (2011).

[3]In our benchmark model, we hold labor supply constant, but we show in the Appendix that similar intuitions apply in a more elaborate environment.

analysis suggests that stigma agents are not those who newly take-up in response to eligibility changes. We propose a decomposition that lets us empirically estimate the contributions of stigma and information in a regression framework. We conclude that that the types of people who are marginal to the eligibility increase are those who are misinformed, not those who are subject to stigma. Put another way, we find that the eligibility increase reduces take-up costs among people who always take up. Those who newly take-up do so because they were previously uninformed, so they capture the full utility gain of the program. When we implement our test, we reject that stigma agents are more elastic than information agents. The upshot of this test is that the eligibility threshold is set too low, if current policy naïvely ignores inframarginal effects.

Finally, we combine the model with our empirical estimates to conduct analysis of the optimal eligibility threshold. As noted, our propositions developed in the model deliver that the social planner will set the eligibility threshold too low if she ignores inframarginal effects and information is more important than stigma in driving inframarginal effects. But how large will the planner's mistake be? Traditionally, local analysis in the spirit of Baily (1978)-Chetty (2006) does not inform the analyst about whether the planner's mistake is large or small when the optimality condition does not hold exactly. In our first exercise, we propose a new method of estimating the magnitude of the planner's mistake, using only the local optimality condition. The core idea is to solve for how much the planner must misperceive the population's risk aversion in order that the optimality condition holds in the context of the naïve model. We find that the naïve planner would overestimate risk aversion by 30%, which corresponds to overvaluing the marginal utility of inframarginal types who always take up and hence over-transferring to them. In a second exercise, we impose more parametric structure to solve for the globally optimal eligibility threshold implied by our model and empirical estimates. We find that the optimal threshold will be 13% too low if the planner ignores inframarginal effects.

Contributions and related literature. Every social program makes some determination about program eligibility (even if the program is universal). Yet much of the vast literature on program design focuses on other policy instruments besides the eligibility threshold. To quantify this, we collected all 278 papers published in the *American Economic Review* between 2010–2018 and the *Quarterly Journal of Economics* between 2010–2019 that met one of 33 search terms about social welfare programs (see Figure 2 and Appendix A for details). Seventy-six of them were primarily about effects or design of social welfare programs, 49 of which involved the study of a specific policy instrument. Yet only 7 (14% of the 49) examined eligibility criteria as a policy instrument that the planner could manipulate to improve welfare. On the other hand, 25 of the 76 papers about welfare programs consider eligibility thresholds as a source of variation for estimating the program's treatment effect. In sum, while economists regularly exploit eligibility thresholds for causal inference, they are often neglected as an aspect of optimal program design. Our paper is among the first to combine empirical estimates of endogenous take-up from eligibility thresholds with a theoretical model that permits welfare analyses of current program rules.

Our work advances several literatures. First, we add to the large body of research in public economics that deals with the optimal design of social programs. Much of this work considers the optimal *benefit level* when take-up is distorted by moral hazard (Baily, 1978; Gruber, 1997; Chetty, 2006, 2008; Hendren et al., Forthcoming). Kroft (2008) introduced to this literature a new fiscal externality which is closer to ours — social spillovers, which are one potential microfoundation for inframarginal effects — and explored how this phenomenon affects optimal benefit size. Relative to Kroft (2008), we emphasize how the mechanism underlying peer effects drives different welfare effects, and we consider the implications for choosing the optimal eligibility threshold. Altogether, analyses of optimal eligibility are rare in this literature.[4,5]

As an additional contribution to the program-design literature, we propose a new strategy to help researchers assess the magnitude of the social planner's mistake when a given social optimality condition does not hold precisely, as is common when empirically testing Baily (1978)-Chetty (2006) conditions. Our approach uses only the local optimality condition and does not require extrapolation with additional parametric assumptions. This contribution therefore relates to other methods of conducting welfare analysis like the Marginal Value of Public Funds (Hendren and Sprung-Keyser, 2020).

Second, we contribute to the public-finance literature on barriers to social program take-up (Moffitt, 1983; Aizer, 2003; Currie, 2004; Heckman and Smith, 2004; Bhargava and Manoli, 2015; Friedrichsen et al., 2018)[6] and the role of social spillovers in program take-up (Bertrand et al., 2000; Dahl et al., 2014). Similar to Finkelstein and Notowidigdo (2019), we consider the welfare implications of these barriers to take-up, but doing so in the context of eligibility thresholds allows us to consider new trade-offs in the model. Our experiment provides clean evidence that aspects of program design may affect stigma costs. We emphasize that reducing stigma introduces two forces — a fiscal externality and a first-order gain to people who always enroll — and provide methods to analyze them empirically. Our discussion of restricted eligibility departs from the most common prior motivation for restricting eligibility, described in Nichols and Zeckhauser (1982), who suggest that limiting program participation can induce self-targeting.[7]

Third, we link research on optimal program design to the growing literature in behavioral public economics (Bernheim and Taubinsky, 2018). Our analysis suggests that individuals' utility depends on social norms, and government policy plays an important role in shaping these norms, like in Lindbeck et al. (1999). Economists have only begun to explore how policy may influence psychological forces like shame or guilt,

[4]Fetter and Lockwood (2018) is a recent example that studies optimal eligibility for old-age insurance. Other papers, e.g. Diamond and Sheshenski (1995), Low and Pistaferri (2015), and Golosov and Tsyvinski (2005) study optimal eligibility in the context of disability insurance.

[5]In many social programs, the eligibility threshold is defined by the benefit size and the slope of the benefit schedule. Thus studies of benefit levels may also contribute to our understanding of eligibility thresholds. Our empirical setting allows us to isolate the effect of eligibility thresholds alone, since the benefit remains constant. Our framework proposes a setting where the policymaker can set eligibility separately from the benefit.

[6]This literature includes several papers studying stigma surrounding SNAP take-up and the rollout of the Electronic Benefits Transfer (Daponte et al., 1999; Currie and Grogger, 2001; Atasoy, 2009; Klerman and Danielson, 2011; Manchester and Mumford, 2012; Eck, 2018).

[7]Two related papers, Kleven and Kopczuk (2011) and Hanna and Olken (2018), model the means test as an instrument for optimal program targeting in the presence of exclusion (Type I) and inclusion (Type II) errors. Our model differs from these papers by emphasizing how the eligibility threshold might directly affect stigma and information.

which may in turn may have important consequences for social welfare. For instance, we provide empirical support for the claim, promulgated in sociology and historical discussion of welfare programs, that programs like Social Security are not stigmatized precisely because they are not means tested (e.g., Katz, 1986). Our model-based decomposition of information shows a novel strategy for isolating information from stigma, which has proven to be difficult in many contexts (Chandrasekhar et al., 2019).

Fourth, we contribute to the study of the Supplemental Nutrition Assistance Program, the subject of a wide-ranging literature.[8] We draw on the data used in Ganong and Liebman (2018), who study how changes in the economic environment, coupled with changes in SNAP program design, affected SNAP enrollment through 2012. Relative to prior work, we highlight a previously unappreciated phenomenon in SNAP (inframarginal effects), show how inframarginal effects affect SNAP stigma and information, and consider their implications for optimal eligibility. As an auxiliary contribution, we also present an academic analysis of the USDA's FSPAS data.

Finally, we contribute to the small literature on inframarginal effects. These effects have received little attention in public economics. The health literature on Medicaid expansions finds evidence of inframarginal effects (Aizer and Grogger, 2003; Sommers and Epstein, 2011; Frean et al., 2017; Sacarny et al., 2022), but it has not considered their implications for optimal program design.[9,10]

Roadmap. Section 2 establishes evidence of an inframarginal effect in SNAP. Section 3 discusses mechanisms underlying the effect, and Section 4 proposes the model of optimal eligibility thresholds. Section 5 presents welfare analysis. Section 6 concludes.

2 Inframarginal Effects in SNAP

This section documents the empirical relationship that motivates this paper. States that have less stringent eligibility standards tend to have higher take-up in SNAP *among inframarginal people* — people whose incomes are low enough that they are eligible everywhere, regardless of the state's eligibility threshold. Appendix A provides more information about the dataset construction and policy variation.

[8]Currie (2003) provides a review of the U.S. food assistance programs and Bartfeld et al., eds (2016) gives extensive coverage to additional research on SNAP. Recent research studies how SNAP receipt affects household members' nutrition, health or other outcomes (Almond et al., 2011; Hoynes et al., 2016; Bronchetti et al., 2019; Bailey et al., 2020; Hastings et al., Forthcoming); whether the marginal propensity to consume food out of SNAP benefits differs from that out of cash (Hoynes and Schanzenbach, 2009; Hastings and Shapiro, 2018); and how SNAP affects recipients' labor supply (Hoynes and Schanzenbach, 2012; East, 2018; Harris, 2021). Ratcliffe et al. (2008) study the effect of *categorical eligibility* on SNAP take-up but do not examine the effect of eligibility thresholds. Homonoff and Somerville (2021) study the screening properties of the SNAP recertification process.

[9]There is little evidence that inframarginal effects would generalize outside the Medicaid setting. Much of the inframarginal effects documented in the health literature pertain to *within-household* take-up of the already-eligible population — for instance, new Medicaid take-up among children who are already eligible because children face less stringent Medicaid requirements than adults, as in Sacarny et al. (2022). By contrast, we show that entire households that were already eligible may sign up when eligibility requirements are relaxed.

[10]Outside of the health literature, Leos-Urbel et al. (2013) and Marcus and Yewell (2021) find that eligibility expansions boost take-up among inframarginal recipients of free school breakfast or lunch programs. These authors study reforms that granted universal eligibility; programs with universal eligibility may be very different than programs like SNAP where eligibility remains restricted. Moreover, program take-up among children may be subject to very different social dynamics and information frictions than among adults.

2.1 SNAP Data

We obtain the total number of people who participate in SNAP from the SNAP Quality Control (QC) files, which are administrative data from the U.S. Department of Agriculture on a random sample of SNAP participants (United States Department of Agriculture Food and Nutrition Service, 2019). The data record granular information about household characteristics, benefit size, and incomes of SNAP participants. The data are a repeated cross-section, so we cannot study households over time. Using these files, we construct the total counts of program participants and those below a given income threshold, for each state and year from 1996 until 2016, the last year for which systematic policy data are available.[11] The Quality Control files are administrative data, so they record people's incomes and household size accurately, thereby addressing concerns about measurement error from Meyer et al. (2015) and others. On the other hand, the dataset is relatively small at the state level. There are about 100,000 observations across 51 states and DC in each year from 2001 to 2016.[12]

Sample and outcomes. We begin with a sample of individuals with household income between 0%–130% of the FPL. In this section, we also focus on a sample including only individuals in households with income between 50%–115% of the FPL. We exclude individuals between 115%–130% of the FPL to address concerns about measurement error: we might consider an individual near the threshold as "inframarginal" when in fact she would be ineligible under a different eligibility regime because of additional restrictions such as asset tests. We focus on individuals above 50% of the FPL because take-up is very high among individuals below 50% of the FPL, regardless of the state's eligibility threshold. Thus there is little scope for increased take-up among this group.

Using this sample, our main outcome is a measure of take-up *counts*. In particular, we use log total enrollment within specific income groups — in our main regression, among people earning 50–115% of the FPL. Almost all individuals in this range are eligible for SNAP in every state. This allows us to study inframarginal recipients; we are not counting increased enrollment among people who are newly eligible. Compared to take-up rates, this outcome has the advantage of not involving imperfect measures of the share of people who are *eligible* for the program as the outcome variable. Instead, we rely on the assumption that the number of people who are eligible regardless of the eligibility threshold (e.g., the number of people in households earning 50–115% FPL) is not correlated with the eligibility threshold beyond the controls we include; we provide support for this assumption below.

We form take-up *rates* as a secondary outcome. Following Ganong and Liebman (2018), we divide the number enrolled (from the QC data) by the number of people within a given band of the income distribution in the state from the Current Population Survey's Annual Social and Economic Supplement (CPS ASEC)

[11]We construct our dataset by modifying the publicly available replication code for Ganong and Liebman (2018).

[12]Relative to comparable datasets, the QC data are best-suited for our analysis. The Survey of Income and Program Participation is not intended to be representative at the state level. The Panel Study of Income Dynamics and Current Population Survey both may be subject to measurement error about SNAP participation.

(Ruggles et al., 2020). For instance, in our main specification, the denominator is the number of people in the CPS who are between 50–115% of the FPL. Crucially, the denominator does not exclude people who are otherwise ineligible for SNAP due to work requirements or asset histories. Thus the take-up rates are likely underestimates. There also may be measurement error in reported incomes in the CPS. We show that measurement error in the CPS data cannot explain our results in Section 2.6.

Just as in other work estimating SNAP take-up, a possible limitation to the analysis is that we do not always observe household assets or work histories, which can affect SNAP eligibility. First, federal rules restrict households with sufficient assets from participating in SNAP. In practice, only a small fraction of households are ineligible for SNAP under these asset histories. Second, under the Personal Responsibility and Work Opportunity Act (PRWORA), single households must meet certain work requirements to participate. However, the changes in these requirements do not coincide with changes in the eligibility threshold, and we also show that the results are similar among households with dependents. We also control for the requirements in robustness checks.

Policy changes. Federal rules require that households below 130% of the FPL are eligible to participate in SNAP. Beginning in 2001, states had the option to expand eligibility to additional households up to 200% of the FPL under Broad-Based Categorical Eligibility (BBCE). The SNAP benefit schedule, which is set nationally, does not depend on a state's eligibility rules.

The eligibility thresholds relaxed under the BBCE correspond to gross income tests. Households must also pass a net income test: net of allowable deductions (e.g., an earnings deduction amounting to 20% of their earned income), their income must be below 100% of the Federal Poverty Level. This is true regardless of the gross income test set by the state. Moreover, because the SNAP benefit size falls in net income and is not changed by the BBCE, many people who become newly eligible from the BBCE receive a small SNAP benefit. Nevertheless, this section documents that raising the gross eligibility threshold led to persistent and large increases in inframarginal take-up.

Not every state that adopted the BBCE took the option to expand the eligibility threshold. In Section 2.6, we note that adopting the BBCE did entail additional changes to state welfare programs, but we reject that these changes can explain the inframarginal effects we document here. Ultimately, 30 states expanded SNAP eligibility through the BBCE through 2016, four of which adjusted eligibility twice during this period (Figure B.1A). Expansions occurred throughout the period, but they were especially likely to occur in 2001–2002 and 2010–2011. The states that do roll out an eligibility expansion are generally distributed across the country, although there are no states in the Great Plains region that implement an expansion (Figure B.1B).

2.2 Econometric Strategy

We estimate an event-study regression that leverages the variation in eligibility provided through the BBCE. We index each event by event-time τ, where $\tau = 0$ represents the first fully treated year. We set $\tau = -1$ in all years for untreated states. We define the "event eligibility rate" in each state s as the eligibility rate as a percent of the FPL after the BBCE expansion in treated states and the federal minimum (130%) in untreated states. We use a balanced panel: we limit the sample to the five years before and after treatment for treated states, and include all years in control states.[13] We normalize our coefficients relative to the year before the event and estimate:

$$y_{s,t,\tau} = \sum_{r \in \mathcal{R}} \eta^r \left(\mathbb{1}(\tau = r)_{s,t,\tau} \times \text{event eligibility rate}_s \right) + \delta_s + \gamma_t + X'_{s,t,\tau}\boldsymbol{\phi} + \varepsilon_{s,t,\tau} \tag{1}$$

where $\mathcal{R}$ is the set of event periods, s indexes states, t indexes years, event eligibility rate$_s$ measures the eligibility rate as a ratio of the FPL, δ_s is state fixed effects, and γ_t is year fixed effects.[14] We include X, a vector of additional linear controls for the state unemployment rate, the log of the number of people in a given income group in the state (measured in the CPS), SNAP outreach spending per person earning under 130% FPL in the states (transformed with $\sinh^{-1}$), and an index of other SNAP policies implemented around the same time (as in Ganong and Liebman (2018), henceforth the "Ganong-Liebman index").[15] In our primary estimates, we use $\ln(\text{enrollment}_{s,t,\tau})$ as the dependent variable ($y_{s,t,\tau}$). The coefficient of interest η^r represents the marginal effect of 1 pp increase in the eligibility rate (expressed in terms of the FPL) on enrollment in event-time r. This specification encodes a standard pre-trends test for whether $\eta^r = 0$ when $r < 0$. Our primary specifications are unweighted. We present standard errors clustered at the state level in this and all subsequent analyses that use state-year variation.

We also pool the data in this sample to estimate:

$$y_{s,t} = \eta \text{ eligibility rate}_{s,t} + \delta_s + \gamma_t + X'_{s,t}\boldsymbol{\phi} + \varepsilon_{s,t}. \tag{2}$$

The variable eligibility rate$_{s,t}$ represents the eligibility as a percent of the FPL in a given state-year, so η is the average effect on inframarginal people after an eligibility expansion.

Discussion of controls. Given that our state and year fixed effects remove fixed differences in outcomes across states and across years, the identifying assumption is that there are no time-varying within-state trends in enrollment (not absorbed by our time-varying state controls). One concern is that states that impose the eligibility increase have faster population *growth* in the inframarginal sample. To address this concern, we control for the log count of the people within the inframarginal income group from the CPS. The economic environment and the policy environment are also relevant for SNAP take-up (Mabli et al., 2014; Ganong and Liebman, 2018). We control for the state unemployment rate to address the concern that states with eligibility increases may simply have more financial distress, and we include the Ganong-Liebman index to address the concern that states that expand eligibility may also impose other policies

[13] We drop the four states with two events in the event-study analysis, as well as the two states that have events too recently to have sufficient post-period data. This leaves 45 states (including the District of Columbia).

[14] For instance, event eligibility rate$_s = 1.3$ represents that the state has the minimum threshold of 130% of the FPL.

[15] The Ganong-Liebman index is the average of several indicators for the presence of different policies that may influence SNAP take-up, such as whether households can apply to SNAP online. See Appendix A for details on the variables that enter the index.

relevant to enrollment. We present robustness to additional threats to identification later in this section. Altogether, these controls do not have a dramatic effect on our results. The most important control is for the count of people who are eligible, which we show eliminates a modest (and insignificant) pre-trend in our event study.

2.3 Results

Descriptive evidence. Before presenting the formal estimates, we begin by visualizing inframarginal effects in the raw data. In Figure 3A, we present total SNAP enrollment per 1,000 people (population-wide) in state-years with eligibility thresholds equal to 130% FPL versus above 130% FPL. We normalize the enrollment by the total population in all states with the relevant income rule to aggregate enrollment counts across states.

First, without the eligibility expansion, very few individuals with household income above 130% FPL take up the program, while with the eligibility expansion, mass appears above 130% FPL where individuals are newly eligible. This confirms that the QC data give sensible estimates of the enrollment counts, and that the eligibility changes relax a binding constraint for some individuals. Second, individuals *below* the threshold also enroll at higher rates with looser eligibility restrictions. These inframarginal effects — the increased enrollment below the threshold — are the subject of our attention.[16]

Figure 3B presents a binscatter of the cross-sectional relationship between SNAP take-up among these inframarginal individuals (i.e., earning 0-130% FPL) and the state's eligibility threshold at the state-year level. We observe five different eligibility thresholds chosen by states between 1996–2016. Mean take-up is roughly 10 pp lower in states with eligibility at 130% of the FPL, the most stringent eligibility standard permitted under federal law.

Event-study specifications. For confidence that the raw data reflect inframarginal effects and are not driven by confounds, we turn to our event study (Equation (1)). We plot log enrollment among our inframarginal sample by event period, relative to event period -1 (Figure 4A). We find no evidence of pre-trends leading up to the policy change. After the policy change, enrollment increases steadily. Figure 4B shows that the effect is concentrated among people in households earning over 50% FPL. We also exclude households over 115% to alleviate concerns (described above) about measurement error or unobserved assets. Our benchmark estimates suggest that increasing the eligibility level by 10 pp of the FPL boosts the number enrolled by 1-2 percent in the five years following the policy change.

To show the effect of controls on our empirical estimates, we present in Figure B.2 the event study with state and year fixed effects only (Panel A) and then add the control for the log of the total number of people between 50 to 115% of the FPL (Panel B). Overall the results are similar without controls. With no controls

[16] We note a slight excess mass around 75% of the FPL, which may be an artifact of the QC data; however, inframarginal effects appear throughout the income distribution.

at all (Panel A), we see some visual evidence of a pre-trend prior to treatment, although the trend is small in magnitude and vanishes three years before treatment. Once we control for the log of the CPS population totals (Panel B), any pre-trend vanishes, and the results in Panel B are very close to those in Figure 4. Note that we are running log take-up regressions: the moderate importance of controlling for CPS population simply confirms that the denominator of a take-up regression matters and is not on-face concerning.

The event-study figure suggests that results grow over time. The effects are larger in years 4–5 than years 1–3, suggesting that inframarginal effects persist or grow in the medium-term.[17] Such effects might grow even several years later if, for instance, information takes time to spread or cascades once others become eligible. Alternatively, stigma might respond only slowly to changes in the threshold.

Placebo. We conduct a placebo test that offers a useful validation of the above results. We observe nine states implement the BBCE *without* expanding eligibility beyond 130% of the FPL.[18] Most of these states adopted the BBCE around the same time as the states in the main event study (2009–2011). Thus, we study the effect of the BBCE in the states that did *not* expand eligibility but *did* implement the BBCE. To implement the placebo test, we show an event study (as in Equation (1)), where treatment represents states that implemented the BBCE but did not expand eligibility (Figure 4C).[19] We use log enrollment among the 0–130% of FPL sample as the dependent variable. This event study gives no effect; we find no evidence to support that the short- or long-term effects in placebo states are the same as the 5-year effect in states with an eligibility expansion.

The placebo test suggests that eligibility expansions, and not ancillary features of the BBCE, drive the results. We cannot completely rule out that the BBCE caused unobserved changes in outreach (not captured by our outreach control variable) or transaction costs (not captured by the vector of SNAP policy controls). But such forces would also be inconsistent with the placebo test, unless they only occurred in BBCE states that also raised the eligibility threshold.

Combined estimates. The event-study specification and placebo test confirm the existence of inframarginal effects. To obtain the pooled effect over all periods, and parsimoniously present robustness to different specifications, Table 1 estimates Equation (2). Our preferred specification (Column 1) uses the sample used in the event study and includes state and year fixed effects, and controls for the state unemployment rate, outreach spending, and the Ganong-Liebman controls.[20] The independent variable is the eligibility threshold as a ratio of the Federal Poverty Level, so that increasing it by 1 corresponds to increasing the threshold by 100% of the FPL. We find that $\eta = 0.107$ and reject $\eta = 0$ at $p < 0.05$. These estimates

[17] Tests of the null hypothesis that $\eta^4 = \eta^3$ and $\eta^5 = \eta^3$ both reject with $p < 0.01$. A joint test for both hypotheses also rejects the null with $p < 0.01$.

[18] States can implement the BBCE for bureaucratic reasons, as the policy can simplify program administration, or to relax the SNAP assets test. See Appendix A.

[19] We exclude states that did increase eligibility from this test, so the regression includes 19 states. A handful of states which adopted BBCE without changing their eligibility thresholds at that point did expand eligibility at a later date. Here, we exclude these states, but the results are similar when they are included and we add a control for the eligibility threshold.

[20] We control for the inverse hyperbolic sine of outreach spending to address state-years with zero outreach spending (Burbidge et al., 1988).

suggest that raising the eligibility rate by 10 pp of the FPL (e.g., from 130% to 140%) boosts take up by 1.07 percent. The modal eligibility increase in our sample is from 130% to 200% of the FPL, which delivers a 7.5 percent increase in take-up among this sample ($0.7 \times 0.107 \approx 0.75$). The results in Column 1 are consistent with the event study plot.

The rest of Table 1 shows that our estimate of inframarginal effects is robust to the particular choice of the specification. Column 2 separates the Ganong-Liebman index into separate indicators for each component variable. Column 3 reverts to the index form of these controls but adds new controls for lagged unemployment and the prevalence of waivers relaxing the SNAP work requirements for able-bodied adults without dependents (ABAWDs), beginning in 2010.[21,22] Column 4 excludes the years 2008–2011 (the Great Recession). Column 5 weights by state-year population. Column 6 computes the treatment effect as the difference between the average of the event study coefficients in the post period and the average of the coefficients in the pre period, weighting all post periods equally. Finally, Column 7 uses all years of data we have (a balanced panel of 50 states and D.C. from 1996–2016), instead of only the event study sample of a 5-year window around the eligibility increase. It also includes states that change eligibility several times or reduce eligibility. Throughout the table, the results are stable: estimates of η range from 0.10 to 0.12.

We also repeat the exercise for two different samples in Table B.1 and find similar results. Panel A shows enrollment responses in the 0-130% FPL sample. The estimates are consistent with the main results but generally lower. This attenuation reflects that our dependent variable (SNAP enrollment) has less scope to rise when almost all people from 0–50% of the FPL already take up SNAP. Panel B assesses enrollment among households with children, as these people are likely not subject to the additional ABAWD work requirements that were relaxed and reimposed during the sample period. Here, we see similar sized, though noisier, effects in this sample. Together, these results and Table 1 provide strong evidence of inframarginal effects from the BBCE, as the effect persists across specifications and samples.

Distributional effects. From which portion of the income distribution do inframarginal effects arise? We present treatment effect heterogeneity by household income (Figure B.5). We estimate a version of Equation (2), using take-up *rates* instead of log enrollment counts so that the values are more directly comparable across income groups with different base rates. Take-up rates increase most among those earning 130–160% FPL, who are barely ineligible before an expansion. The effect in this group is larger than the largest effect in the inframarginal population, among those earning 100–130% FPL. However, even after the expansion, take-up in the newly eligible group is still much lower than any other group. We also see that the treatment effect size is increasing with household income within the inframarginal sample; however, this may partially reflect that the base take-up rate is much lower among households with relatively more income.

[21] We use data on ABAWD waivers from data generously shared by Harris (2021).

[22] Figure B.3B also shows the event study where the sample includes only SNAP recipients in households with children.

Characterizing compliers. Who is most affected by eligibility expansions? To the extent that inframarginal effects are driven by reductions in barriers to take-up ("ordeals"), they may affect the targeting properties of the expansions (Nichols and Zeckhauser, 1982). If inframarginal effects influence SNAP's screening capacity, we expect the people who join the program after an eligibility expansion to look different on observables than the previously enrolled. On the contrary, we find little evidence that the eligibility threshold affects the characteristics of SNAP enrollees earning 50–115% FPL (Table 2). Of the characteristics we can analyze, we only find a significant positive effect on the average poverty level of enrollees. However, the magnitude of these effects is small: increasing the eligibility threshold from 130% FPL to 140% FPL, for example, would imply a 0.07% FPL increase in the average gross income of SNAP recipients. Together, these results suggest that whatever the ordeals behind inframarginal effects, they do not have substantial screening effects.[23]

2.4 Interpreting the Magnitude of the Results

We now provide three ways of interpreting the magnitude of the results.

Take-up elasticity. We estimate the elasticity of take-up with respect to the share of the population who is eligible. The elasticity will also play a critical role in the theoretical model.

We employ an instrumental variables approach to estimate this elasticity. The share of the population eligible for SNAP is affected by confounding conditions which also affect the number of people below a certain income level. The eligibility expansions provide plausibly exogenous shocks to the share eligible. Thus we instrument for the log share eligible for SNAP using the state-and-year-specific income cutoff as a ratio of the Federal Poverty Level. The exclusion restriction is that eligibility expansions are not associated with take-up of inframarginal people except through changes in the share eligible.

We return to Equation (2) from Section 2. We use a log-log specification, with $\ln(\text{take-up})$ and $\ln(\text{share eligible})$ as the dependent and independent variables, respectively. The estimating equation is:

$$\ln(\text{take-up})_{s,t} = \eta \ln(\text{share eligible})_{s,t} + X'_{s,t}\boldsymbol{\phi} + \delta_s + \gamma_t + \varepsilon_{s,t}, \tag{3}$$

where we instrument for $\ln(\text{share eligible})$ using the state eligibility threshold as a ratio of the FPL. Here η represents an elasticity rather than a level effect.

We present the IV estimates for the 0–130% sample using all the data (Table 4, Panel A) as well as the event-study sample (Panel B). We document a strong first stage: in the full sample, increasing the eligibility threshold by 10% of the FPL increases the share of a state population that is eligible by 7.28% (t-stat = 21.37), with similar results for the event-study sample. Our 2SLS estimate in the full sample is $\eta_m = 0.130$ (SE: 0.067); the estimate in the event-study sample is $\eta_m = 0.104$ (SE: 0.077). We also document that simple

[23] Table 2 also shows no evidence of an increase in the share of enrollees whose SNAP certification period is less than 6 months, suggesting that new enrollees also do not have more volatile income.

OLS regressions of log take-up on the log share eligible have the opposite sign, likely due to the omitted variables bias we described above. The full sample estimate is more precise, so we prefer it when used for welfare analysis.

Comparison to inframarginal effects in Medicaid. We now convert our inframarginal effect estimate to the same units as Sacarny et al. (2022) to compare magnitudes. Sacarny et al. (2022) find that about 0.1 previously-eligible children enter Medicaid for every adult who entered Medicaid from the Oregon Health Insurance experiment. To compare to this point estimate, we employ the magnitude of the inframarginal effect among the entire inframarginal population (Table 1A).[24] We find that .91 (standard error: 0.57) inframarginal people between 0–130% of the FPL are induced to take up the program for every newly eligible person who takes up the program.

We cannot reject that the treatment effects are equal to those in Sacarny et al. (2022). Even so, our point estimate is that inframarginal effects in this setting are nine times larger than in Sacarny et al. (2022), which warrants discussion. Altogether, we have no reason to expect that inframarginal effects will be of the same magnitude across programs and over time. In this setting, expanding the SNAP eligibility threshold for gross income does not loosen other eligibility criteria (e.g., the net income threshold). These criteria may bind for people with higher incomes. As a result, an eligibility expansion can lead to higher take-up among the inframarginal population without many newly eligible people joining the program.

Comparison to outreach spending. A final way of benchmarking our effects is to compare the take-up from inframarginal effects to the take-up from direct SNAP spending on information and outreach. The SNAP Policy Database contains information on states' outreach spending, but we do not have quasi-random variation in this spending. For an effect of outreach on take-up, we turn to the randomized control trial run by Finkelstein and Notowidigdo (2019), where the authors find that sending mailers to people who are likely eligible for SNAP but not enrolled boosts enrollment. They calculate that their intervention costs about \$20 per additional enrollee induced to join by the outreach intervention. At this rate, it would cost about \$66 million to increase enrollment in the inframarginal population by the same amount as raising the income eligibility threshold from 130% FPL to 200% FPL.[25]

On the one hand, \$66 million is a fraction of the total annual spending on SNAP (\$70 billion in 2016). On the other hand, it is more than three times what all states combined spent on outreach in 2016 (\$17.4 million). Finally, the mechanical cost of raising eligibility goes to program recipients who are newly eligible. But the mechanical cost of outreach does not go to program recipients. To summarize, outreach spending

[24] Let the point estimate for the entire inframarginal population from Table 1A, Column 1 be $\hat{\eta}_i$. We then estimate a version of Equation (2), using the log of the total number of people on the program as the dependent variable (and controlling for the log of the number of people below 130% of the FPL from the CPS). Let the point estimate from this regression be $\hat{\eta}_t$. We then present $\frac{\hat{\eta}_i}{\hat{\eta}_t - \hat{\eta}_i}$, where the denominator represents the increase in the marginal population and the numerator represents the increase in the inframarginal population.

[25] There were around 44 million SNAP enrollees in 2016. To derive the number of new enrollees from such an increase, we multiply 44 million by the increase in take-up (7.5%) implied by our estimates in Table 1 at the modal eligibility threshold increase (130% to 200% FPL). Finally, we multiply this by \$20 per additional enrollee to arrive at \$66 million.

may be an alternative instrument for increasing SNAP take-up, but it is not obviously a better one than increasing information by raising the eligibility threshold.

2.5 Cost-effectiveness

We conduct a back-of-the-envelope calculation to compare the mechanical costs of two natural interventions to raise take-up of the inframarginal population by 1 pp: raising the eligibility threshold and raising the benefit size. We find that the methods have similar mechanical costs (Table 3). Using the η_m estimated in Equation (3), we calculate that to increase take-up by 1 percentage point, an additional 4 pp of the US population would need to be eligible for SNAP. If take-up in the newly eligible population is similar to take-up among people who are just barely eligible (25%, Figure B.5), and the benefit size is similar to the benefit size in this group ($707 per person-month, calculated from the QC data), then this intervention costs an additional $2.2 billion per year. To compare raising the eligibility threshold to the cost of raising take-up by raising the benefit size, we assume that the elasticity of take-up with respect to the benefit size is 0.5 (see Section 5.3 for details). To get a 1 pp increase in take-up, the benefit size would need to increase by $56 per year for 44 million SNAP enrollees — costing $2.5 billion per year.

This cost-effectiveness point does not have direct implications for social welfare, since the cost and benefit of each policy instrument also depend on recipients' willingness to pay. However, it is a model-free way to compare the tools.

2.6 Robustness

Balance. The identifying assumption in our event study is that there are not other factors besides the eligibility threshold that contribute to inframarginal take-up and coincide with the means test policy change. A related concern is that the states which change their threshold are different from those that do not; note that with our event study framework, internal validity does not require that control and treatment states are similar.

We first test whether states that implement the BBCE bundle the change with other adjustments to SNAP policy. We note that, following Ganong and Liebman (2018), all our regressions control linearly for an index of eight other SNAP policies that occur during the same period (measured by the SNAP Policy Database). As Table 1 shows, including this index makes little difference, which gives additional confidence that unobserved policies do not affect the results. Moreover, when the index is separated into its component parts (Column 2), the magnitude of the effect is not diminished. One might nevertheless worry that the eligibility expansions were bundled with informal policies (e.g., flyer campaigns) that the SNAP Policy Database does not measure. To allay this concern, we present a placebo event study, with the SNAP policy index as the dependent variable (Figure B.4A). We find no evidence that the SNAP index increases after the eligibility expansions. Overall, the test is inconsistent with economically material bundling of SNAP policies. Finally, because we control for this index, this objection requires unobserved policies to affect the outcomes even after residualizing by the index.

We next examine whether economic conditions change leading up to the changes in the eligibility threshold. We estimate Equation (1) with the log of the CPS counts of the people at 50–115% of the FPL as the dependent variable (Figure B.4B). We find a slight pre-trend in the CPS populations three years before the event, but the effects are modest. We discuss whether including this control affects the results above; it helps alleviate a moderate but insignificant pre-trend in the main event study. Similarly, we estimate Equation (1) with the unemployment rate as the dependent variable. Although the unemployment rate appears to grow in advance of the policy, the trends are insignificant (Figure B.4C). Moreover, the time series pattern of the changes in the unemployment rate do not align with our main results: the unemployment rate returns to 0 after 5 years, whereas our main effects persist. That is why when we control for the unemployment rate, this control does not materially affect our results (Figure B.2B versus Figure 4B). We conduct two additional tests to address the concern about the unemployment rate changing in advance of the policy. We include a further control for lagged unemployment (Column 3, Table 1). We also exclude the Great Recession (Column 4), when unemployment rates had the greatest fluctuation. Our results remain robust.

Our final balance exercise is a standard one: we compare states which did and did not ever change their eligibility threshold (Table B.2). Because our main results use an event study, imbalance in levels in the pre-period is not itself concerning; however, it can still be helpful to understand whether treated states were different from untreated states. In Panel A, we see that states which ever changed their eligibility threshold have significantly higher average family incomes in the pre-period (measured in 2000, the last year before any state changed its eligibility threshold), and marginally significantly higher measures of SNAP access-related policy (the Ganong-Liebman index and SNAP outreach spending). However, Panel B shows that these measures are not strongly associated with the size of the means-test change, which provides suggestive evidence that the policy decision is not driven by these measures.

Measurement error. To mitigate concerns about measurement error, our key empirical fact (inframarginal effects) uses the QC numerators as the dependent variable, as in Table 1. Even so, we control for the size of the eligible population, which may be measured imperfectly in the CPS.[26] First, we note the above point that the share eligible does not change with treatment. This pushes against concerns that differential measurement error in the pre- and post-periods drives our results. Appendix B.3 presents a simulation that shows that only an implausible amount of measurement error, exactly coinciding with the event and only in treated states, could explain our results.

There may also be measurement error in the timing of the policy implementation.[27] We use data at the

[26]We show the main event study with take-up rates on the left-hand side in Figures B.3C and B.3D.

[27]We follow the date of the policy implementation in the SNAP Policy Database. However, the precise implementation date may vary across sources, and the legal implementation date may not coincide with the date that the program actually began accepting

annual level in our main specification because we measure the number of people who are eligible from the March CPS, which is only available annually. Moreover, the QC data contain relatively few people at the month-state-income group level. However, BBCE policies can be implemented mid-year. In Figure B.3A, we show our event study using monthly data to estimate Equation (2). It looks broadly similar, although the inframarginal response is slightly slower to appear. This reflects the fact that in our main specification, we index policy implementation to the beginning of the first fully treated year.

Other effects of the BBCE. A related concern is that some states grant extra eligibility through the BBCE together with explicit referrals or brochures to SNAP. As a part of the BBCE, states sometimes use the budget from the Temporary Assistance for Needy Families to fund referrals to state services, including SNAP. As Figure 4C shows, states which adopted the BBCE but did not expand eligibility did not see similar effects on SNAP enrollment. This placebo test thus constitutes strong evidence that only the eligibility threshold, and not ancillary BBCE-related policies, are responsible for the take-up effect.

The BBCE also waives some rules on the maximum assets that block families in states without BBCE from obtaining SNAP. First, the above placebo study also rejects this concern, since BBCE states that maintain eligibility at 130% of the FPL but do change their asset limits do not exhibit a take-up increase. Second, in practice, these asset rules affect a small number of families. Ganong and Liebman (2018) find asset waivers were responsible for only a small share of increased take-up in recent years. Eslami (2015) finds that 4 percent of inframarginal people who participate in SNAP are eligible only due to state asset eligibility rules.[28] There are a host of such asset waivers, including many not linked to the BBCE. But even assuming all these households were only eligible due to the BBCE, the asset waivers could not explain even half of the inframarginal effects we find.

Two-way fixed effects and negative weights. Concerns about negative weights (Callaway and Sant'Anna, 2020; Sun and Abraham, 2021) are unlikely to apply in our setting, since: (i) there is a large pool of never-treated units, and (ii) we do not have always-treated units. As a check we implement the heterogeneity-robust stacked estimator from Cengiz et al. (2019) and the Sun and Abraham (2021) estimator. We obtain similar results (Appendix Figure B.6); the Sun and Abraham (2021) estimator delivers somewhat larger results in years 4 and 5 (but the confidence interval safely contains the original point estimates).

Policy salience. An additional concern is that the inframarginal effects arise in our setting because the expansions are *salient* to people, but they are not steady-state responses. First, we show that eligibility expansions boost take-up up to five years after the expansion, so they at least have effects in the medium-term. Second, the event study plots also show that the jump in take-up does not coincide with the expansion but grows over time.

people with incomes larger than 130% of the FPL (e.g., if program social workers need to be trained on the new procedures). In practice, measurement error along these lines would merely add noise to the event study.

[28]See computation in Ratcliffe et al. (2016).

3 Mechanisms: Information and Stigma

Why does the eligibility threshold affect inframarginal take-up? The question relates to a long-standing literature on incomplete take-up of social programs that categorizes barriers to take-up into incomplete information, stigma, and other enrollment costs. Furthermore, the model in the following section will also make clear that the mechanisms matter for welfare analysis.

One hypothesis is that the eligibility change affects stigma around SNAP take-up. For example, it is possible that when SNAP becomes available for relatively wealthier people, SNAP no longer conveys as much of a negative signal. We test this hypothesis using an online experiment in which we exogenously change participants' beliefs about the SNAP means test. A second hypothesis is that changing the eligibility threshold increases the information about the program. For example, because more people are eligible, people can more easily obtain information about how to apply from friends or family. We test this hypothesis by making novel use of USDA survey data on SNAP stigma and information.

We do not emphasize non-stigma enrollment costs as a potential mechanism, because Section 2 provides evidence that the eligibility changes did not meaningfully change enrollment costs. For example, we find no differential effect on people with different recertification periods. However, our theoretical model will permit changes in these costs.

3.1 Online Experiment: Evidence of Stigma

Here, we present evidence from an online experiment that the eligibility threshold may affect perceived stigma around SNAP take-up.[29]

3.1.1 Experiment Design

The objective of the experiment is to induce variation in participants' beliefs about the share of people who are income-eligible for SNAP. In particular, we study how raising people's beliefs about the share eligible affects self-reported stigma.[30] Figure C.1 summarizes the experiment design.

Main experiment. Our main experiment was embedded in a question asking respondents to report what share of Americans they thought were income-eligible for SNAP in 2016.[31] On this page of the survey, all respondents were given a truthful hint: *"In 2016, in one of the U.S. states, roughly [X] of the population had low enough income that they could qualify for SNAP."*

[29] We used the survey provider Lucid; other papers using Lucid include Wood and Porter (2019) and Bursztyn et al. (2020). We ran the experiment in March 2020. The onset of the coronavirus pandemic should not complicate the treatment-control differences via our randomized information provision.

[30] The complete survey instruments are available from Rafkin's website.

[31] Reports were incentivized as follows: participants were told at the beginning of the survey that a lottery would be conducted among respondents who answered a factual question correctly, and the winner would have $50 donated to her choice of charity.

X was randomly either 15% or 38%, which were the highest and lowest state-level eligibility shares we see in the administrative SNAP data from 2016. We refer to those participants who saw the 38% hint as those in the "high-share" treatment.

Belief elicitation. After implementing the treatment, we conduct a manipulation check by eliciting people's beliefs about the share of people eligible for SNAP. We asked: *"In 2016, how many out of every 100 people (in all U.S. states) do you think have low enough income that they could qualify...?"*

Auxiliary experiment. Following the belief elicitation, we included an auxiliary randomization: we informed a random subset of participants about the correct share (27%, as per our calculations combining the CPS and the SNAP Policy Database). Depending on their prior beliefs, this treatment (which we call the "belief-correction" treatment) is intended to cause participants to update up or down about the share of people who are eligible for treatment. Unlike the main treatment, the auxiliary belief correction treatment does not have a tight connection to changes in an eligibility threshold.[32] As a result, we relegate discussion of the belief-correction treatment to the Appendix.

Stigma elicitation. We asked respondents to rate their agreement, on a scale from 1 to 9, to a series of eight statements about SNAP: *(1) I would prefer not to use food stamps because I would rather be self-reliant and not accept help from the government; (2) I believe that people should do what they can to avoid being on food stamps; it is better to make it on your own; (3) Most people believe that someone who uses food stamps is just as hard-working as the average citizen; (4) If I used food stamps, I would be concerned that people would treat me disrespectfully at stores; (5) Most people believe that someone who uses food stamps does so because of circumstances outside their control; (6) Most people think less of a person who uses food stamps; (7) Most people who use food stamps would go out of their way to prevent others knowing about their food stamp receipt; (8) If I used food stamps, I would avoid telling other people about it.*

We aggregate the statements into two indices: (i) "first-order stigma," which ask respondents about their own attitudes (statements 1, 2, 4, and 8 above), and (ii) "second-order stigma," which asks respondents about others' attitudes (statements 3, 5, 6, and 7).[33] We standardize these outcomes using the mean and standard deviation of the control group and then average the standardized values as in Kling et al. (2007). We also show the effects on an aggregated index.

Either first- or second-order stigma could play a role in inframarginal effects, depending on the model. If people care about social image and take-up is partly observable, the extent to which others condone or sanction SNAP may affect take-up costs. With first-order stigma, people may have a hedonic aversion to SNAP that does not depend on others' views. Such aversion could easily influence take-up if modeled as a direct take-up cost.

Sample construction and balance. We drop participants who fail either of two pre-registered attention

[32] We originally included the auxiliary experiment because recent papers, e.g. Bursztyn et al. (Forthcoming), use similar belief corrections to manipulate people's prior beliefs.

[33] We reverse the scale for questions 3 and 5 so that positive numbers always indicate more stigma.

checks, as well as those who did not provide a prior or respond to all stigma questions. Our final sample has 2,131 participants (79% of the original sample). Table C.3 summarizes these sample limitations and confirms that attrition, inattention, and non-response were balanced between treatment and control. Appendix A describes the data cleaning in more detail.

The sample is balanced across the high-share treatment (joint p-value: 0.94) and has a relatively similar composition as the U.S. on average (Table C.1). In some tests, we restrict the sample only to the 512 people below 130% of the FPL, because this subgroup — the inframarginal SNAP sample — is of particular interest for inframarginal effects. Among this subgroup only, a joint F-test suggests experimental imbalance (p-value: 0.02).[34] The experiment was randomized but not stratified, and any imbalance in this subgroup occurred by chance. To address the lack of balance when studying treatment effects in this subgroup, we present robustness tests that control for available demographics. We stress that the experimental treatment is balanced in the full sample, and we emphasize results from the full sample as a result.

Econometric strategy. In our primary specification, we simply compare the difference in means across treatments:

$$y_i = \beta \mathbb{1}(\text{high})_i + \gamma \mathbb{1}(\text{truth})_i + \varepsilon_i, \tag{4}$$

for individual i, where β represents the coefficient of interest. In robustness exercises, we estimate a version of Equation (4) with additional demographic controls. We conduct inference using robust standard errors.

3.1.2 Experiment Results

Beliefs about eligibility. The high-share treatment successfully moved beliefs about eligibility (Figure C.2). Both groups report beliefs that are slightly overestimated but reasonable; the mean for the low group is that about 30% of people are eligible, and the mean for the high group is that about 39% are eligible. The raw difference in means is 9.21 pp (SE: 0.80, p-value < 0.001). The standard deviation of beliefs in the control group is 19.8 pp, so the treatment raised the beliefs by a sizable 0.47 standard deviations. Moreover, while the low- and high-share treatments anchored a large fraction of people toward the numbers we provided them (15% and 38%), it also moved beliefs for others throughout the distribution.

Stigma. First, we note that responses to the eight stigma statements are somewhat but not overwhelming correlated (Figure C.3), so each question may contain independent information about the participants' views. A concern is that participants simply anchor to their responses on the first question since the question order was not randomized. In fact, while we find that responses to the second question are relatively correlated with the first question (correlation $\approx$ 0.65), other questions do not display a large correlation with the first question.

Next, we turn to investigating the treatment effects. Increasing individuals' beliefs about the share of

[34] The most imbalanced covariate is that the high-share treatment is less concentrated in the Northeast region than the low-share treatment (p-value of difference: 0.02).

Americans eligible for SNAP decreases their self-reported second-order stigma (Figure 5A). Aggregating the results into indices, the high-state treatment reduced second-order stigma by -0.050 standard deviations (SE: 0.027, $p = 0.061$). Effects are larger in magnitude among the 512 participants below 130% FPL (point estimate: -0.109, SE: 0.058, $p = 0.061$).

The treatment effects for second-order beliefs are similar across questions that form the second-order index. In the full sample, the high-share treatment reduces stigma the most in the question about whether most people believe recipients "go out of their way to prevent others knowing about their food stamp receipt." We find larger effects among people who have ever taken up SNAP, men, and Democrats, although treatment effect heterogeneity is not generally significant (Figure C.4). On the other hand, we find positive but statistically insignificant results on first-order stigma (Panel B).

We summarize these results in Table C.4, and we find very similar results when we include demographic controls (Table C.5). Moreover, when we aggregate the second- and first-order stigma results into a combined index, we find no statistical evidence of an average effect on stigma, although the point estimate is negative. The null result is mechanically driven by the null or slightly positive effect on first-order stigma.

3.1.3 Experiment Conclusions and Caveats

This experiment provides new empirical evidence on one possible mechanism underlying inframarginal effects. It serves as a useful contribution in its own right. The health literature on inframarginal effects has not provided clean evidence that either information frictions or stigma costs contribute to inframarginal effects. Additionally, evidence about stigma in social welfare programs remains elusive (Currie, 2004; Bhargava and Manoli, 2015). Our experiment suggests key aspects of program design, e.g. the eligibility threshold, indeed have the potential to affect program stigma.

While the experiment suggests that stigma could, in principle, drive inframarginal effects, the evidence we provide is not dispositive. We note several caveats. First, we find no effects on first-order stigma. First-order beliefs about, say, whether one should accept help from the government may represent deep-seated aspects of one's identity. It is therefore not surprising that people's first-order beliefs may be hard to move in a light-touch survey experiment. Second, an important caveat about our design is that we presented the high- and low-share treatments before a belief correction exercise. The belief-correction exercise itself does not provide evidence that the means test affects stigma (see results in Appendix C). Third, as with any online experiment, one may worry about external validity. We cannot experimentally manipulate the actual SNAP eligibility threshold, only people's perceptions of it.

Finally, we do not have a measure of whether the intervention affects SNAP take-up; this motivates our next empirical analysis.

3.2 Stigma, Information, and Take-Up

In the previous section, we found evidence that the means test affects perceived stigma around SNAP take-up. In this section, we show that the subgroups whose stigma decreases the most in the online experiment do *not* have the largest changes in take-up in the administrative data used in Section 2. Instead, those subgroups who appear to have the lowest stigma about SNAP, and those that are least likely to have information about SNAP, are those that see the largest changes in take-up. Together, this suggests that the means-test affects take-up largely by increasing information availability.

Data. For this exercise, we include data from an additional source: the USDA's Food Stamp Program Access Study (FSPAS) (Bartlett et al., 2004).[35] The USDA's FSPAS involved phone and in-person interviews conducted in 2001 with a reference month of June 2000. Since the analysis of inframarginal effects uses QC data from 1996–2016, the FSPAS data occur toward the beginning of the sample period. We use data from two subsurveys: one of a random sample of approved SNAP applicants, and another of a nationally representative sample of likely eligible nonparticipants.[36]

In both surveys, respondents are asked a series of four questions about their perceived stigma around SNAP; they are also asked a number of questions about the information they have about SNAP. We consider respondents who reported any feelings of stigma to be affected by stigma. We consider any nonparticipants who reported a lack of information about any of three information questions to be affected by information frictions.[37]

The data include demographic information, including gender, age, race, marital status, and number of children. Because we also have these variables (as well as household income) in our online experiment and in the administrative SNAP data, we can compare statistics at the demographic cell level between datasets. Each cell is defined by the gender and age (binned into 18-30 year-olds, 31-65, and 65-100) of the household head; whether or not the household head is a non-Hispanic white; the household composition (married adult with children, unmarried adult with children, or adult(s) without children); and, where available, the income decile of the household when compared to the distribution of incomes in the US Current Population Survey.

Descriptives. Figure B.7 shows the stigma and information statements presented to FSPAS respondents and the share of respondents who agreed with each statement. About 40% of the sample agreed with at least one of the stigma statements, leading us to categorize them as being affected by stigma. Of those who agreed with any stigma statements, almost half (47%) agreed with only one, and another 29% agreed with

[35] To our knowledge, this is the first academic study of the FSPAS, which the USDA generously shared with us.

[36] Among nonparticipants deemed to be eligible from an initial screener, 96.3% completed the survey. Among applicants randomly sampled from lists provided by SNAP offices, 56.7% of were reached and completed the survey. We analyze a sample of 1,585 respondents who either answered questions about stigma or answered questions about information (and have non-missing weights assigned by the USDA).

[37] These asked whether participants had heard of SNAP; whether they thought they were eligible for SNAP; and whether they knew where to go to get SNAP benefits.

two. Meanwhile, about 60% of the nonparticipant sample disagreed with any of the information statements, leading us to categorize them as being affected by information. Finally, we show descriptive statistics by whether we consider the respondent to be affected by stigma (Table B.3). Those who report any stigma are more likely to be white, are on average younger, and are more likely to have children in their household. Notably, those who report any stigma are *more* likely to be enrolled in SNAP.

Results. Next, we study whether demographic cells with many stigma or information types have larger inframarginal effects (the binned scatterplots in Figures 6).[38] First, we find that cells with many stigma types have smaller inframarginal effects (panel A), and cells with many information types have larger inframarginal effects (panel B). While the relationships are noisy, we can statistically reject that the slopes are equal to zero at $p < 0.05$. The fact that the cells with many stigma types are statistically *less* likely to have large inframarginal effects is particularly suggestive that inframarginal effects are not driven by stigma. To complete the story, Figure 6C presents the correlation between the treatment effect from the online experiment (i.e., the effect of the perceived means test on reported stigma) and the treatment effect from the main analysis (the effect of the means test on take-up). Subgroups with the largest reductions in stigma when the means test increases do *not* have the largest inframarginal effects.

Discussion. Taken together, we find no evidence that stigma contributes to inframarginal effects. We find some suggestive evidence that information is responsible. Nevertheless, the experiment shows that increases in the means test decrease stigma costs. How should we interpret these facts? We use a model to conduct welfare analysis.

4 Model

In this section, we develop a model for analyzing optimal eligibility in the presence of inframarginal effects. Our model takes as given that SNAP — or, more generally, any lump-sum, means-tested transfer program — exists. We do not model the optimality of SNAP above and beyond redistribution via an income tax. Instead, we use the model to consider how to determine the share of the population which should be eligible for a redistributive program that has incomplete take-up. We use the model to emphasize the relevance of distinguishing between different mechanisms for the effects in Section 2. Our main argument is that whether take-up barriers are consistent with agent optimization affects both the incidence and the size of the welfare gains.

[38] Appendix D gives details about forming these measures. Because these binned scatterplots plot cell-level coefficients estimated with error, we conduct our tests weighting by the inverse of the product of the variances of the coefficients, also discussed in the Appendix.

4.1 Benchmark

We begin by analyzing a benchmark model where take-up responds endogenously to the eligibility threshold, but all consumers optimize. We add optimization failures due to imperfect information in Section 4.2.

We start by assuming take-up costs are normatively relevant (i.e., consumers are perfectly rational optimizers with respect to the take-up decision). In the following discussion, we often refer to these costs as "stigma costs," since we are especially interested in the case in which raising the eligibility threshold can reduce stigma and therefore boost take-up. However, the costs refer to any cost that inhibits take-up, e.g. hassle costs. Other possible mechanisms that might be cast as changing costs include transaction costs, for instance if more stores accept SNAP once more people become eligible. If the threshold reduces *uncertainty* about eligibility, that might be either an increase in the expected net benefit or an increase in awareness, depending on the model.

There is a continuum of individual types $\theta \sim F$, which correspond to ability or higher consumption. Types are perfectly observable, but we consider an environment in which the government cannot give a type-specific transfer (e.g., due to political economy or implementation constraints). The government offers a social program with a lump-sum consumption benefit B. The government provides B only to types $\theta < m$, where m is the eligibility threshold (or means test/income cutoff) also chosen by the government. We normalize the distribution of types to be *quantiles* of the distribution used to determine program eligibility (for example, the income distribution), i.e. $F := U[0,1]$.[39]

Denote the welfare weight on type θ by λ_θ, which refers to the welfare weight of quantile θ in the type distribution. For example, λ_0 refers to the weight that the planner places on the lowest-quantile person. We assume that the welfare weights are weakly decreasing in θ.

Assume all people have the same twice continuously differentiable and concave utility function from taking up the benefit, denoted by $u(B)$. Normalize individuals' outside income to be 0 and outside utility to be $u(0) = 0$. We already permit differences in realized consumption utility for each type to enter the planner's problem through λ_θ. We can simply redefine a type's welfare weights to capture the different consumption utility that the type experiences.

Individuals choose whether to take up the benefit. We incorporate inframarginal effects by allowing the take-up probability to depend on the eligibility threshold m. In particular, every individual faces a take-up utility cost c, drawn from a continuously differentiable distribution H (which we additionally assume has a finite first moment). We suppose H depends on m, so $H(\cdot|m)$ and $h(\cdot|m)$ are the CDF and PDF of c.

We assume separability between the consumption benefits and take-up cost. Write realized utility as

[39] Note that this normalization is innocuous: it amounts to letting type m simply refer to the m-th quantile of the type distribution. The planner chooses what fraction of people are eligible, rather than the threshold type who is eligible.

$U(B,c) = u(B) - c$. Then individuals participate in the program if $u(B) - u(0) > c$, i.e. $u(B) > c$.[40] Because $H(u(B)|m)$ is the take-up probability, define $p(B,m) := H(u(B)|m)$. We sometimes suppress arguments and write $p(B,m)$ as p, so that the probability an individual of type θ takes up the program is p_θ. We also assume that each type takes a cost draw from the same distribution, so that $p_\theta = p$.

Labor supply. We assume households' labor supply is fixed: there are no labor supply responses to the threshold. We relax this assumption in Appendix G and show how a general problem with endogenous labor supply nests the key insights in this framework. Assuming fixed labor supply simplifies the framework considerably and permits us to focus on our novel mechanism (inframarginal effects).

Planner's problem. The planner faces a budget constraint T. In our setting — as in, e.g., Finkelstein and Notowidigdo (2019) — the social planner cannot simply set the optimal nonlinear income tax. For instance, the planner in our benchmark model might correspond to a state-level administrator tasked with choosing the parameters of a fixed program budget allocated by Congress. Indeed, such a setting is especially natural with SNAP, where state administrators choose the eligibility threshold but face an exogenous federal income tax.[41]

The planner solves:

$$\max_{B,m} p(B,m) \left(\int_0^m \lambda_\theta u(B) d\theta - \int_0^m \int_{c \le u(B)} \lambda_\theta c h(c|c < u(B), m) dc d\theta \right) \tag{5}$$

subject to $p(B,m) \int_0^m B d\theta \le T$ and $m \in [0,1]$.

Let $\eta_m := \frac{\partial p}{\partial m} \frac{m}{p(B,m)}$ be the take-up elasticity with respect to the eligibility threshold. The parameter η_m is the inframarginal effect, represented as an elasticity. We assume throughout that increases in m reduce costs, so $\frac{\partial p(B,m)}{\partial m} > 0$ for all B. For instance, raising the eligibility threshold might decrease stigma costs if stigma directly depends on the share of people who are eligible or take-up, as in Lindbeck et al. (1999).

Define η_B as the elasticity of take-up with respect to the benefit size, B: $\eta_B := \frac{\partial p(B,m)}{\partial B} \frac{B}{p(B,m)}$. Let $\gamma(B,m) := \frac{E[c|c<u(B),m]}{u(B)}$, noting $\gamma(B,m) < 1$. The parameter γ is the expected cost-benefit ratio conditional on take-up. It represents the share of the welfare gain from the benefit dissipated by the cost of taking up the benefit. For instance, if $\gamma = 0.5$, then costs represent half the utility gain (at $u(B)$).

Let $\lambda_{\text{avg}}(m)$ be the average welfare weight up to type m:

$$\lambda_{\text{avg}}(m) := \frac{\int_0^m \lambda_\theta d\theta}{\int_0^m d\theta} = \frac{\int_0^m \lambda_\theta d\theta}{m}.$$

Then the first-order conditions yield the following benchmark:

[40] Of course, utility is also a function of other consumption. The model takes this consumption as exogenous and normalizes $u(0) = 0$. In Appendix G we show that the model can accommodate different consumption across types, at the cost of notational complexity.

[41] We close the planner's budget constraint by trading off eligibility threshold increases with per-person benefit size decreases. While this is a natural tradeoff to consider theoretically, in practice in SNAP, the benefit schedule and the eligibility threshold are chosen by different decision-makers (federal and state, respectively).

Proposition 1. *At an interior optimum, m and B satisfy:*

$$\overbrace{\frac{\frac{\lambda_m}{\lambda_{\text{avg}}}(1-\gamma)u(B)}{u'(B)}}^{\text{Welfare-weighted WTP of newly eligible}} + \overbrace{\frac{\frac{m}{p(B,m)}\int_0^{u(B)}\frac{\partial H(c|m)}{\partial m}dc}{u'(B)}}^{\text{WTP for lower } c} = \overbrace{\frac{B(1+\eta_m)}{(1+\eta_B)}}^{\text{Fiscal externality}}. \tag{6}$$

All proofs are in Appendix F.[42] Proposition 1 has familiar Baily (1978)-Chetty (2006) logic. At an optimum, the social planner equates the willingness to pay for a higher means test (the left-hand side) to its fiscal externality (the right-hand side). The willingness to pay for a higher means test combines: (i) the (welfare-weighted) utility gains of people who are newly eligible, and (ii) the utility gains from lower costs to previously eligible types who would have enrolled irrespective of the means test. The fiscal cost incorporates two fiscal externalities, one positive and one negative: raising the means test causes higher take-up from reduced costs, but it also causes lower take-up from the lower benefit amount given to each enrollee.

Notably, our model embeds Baily (1978)-Chetty (2006) logic in the context of a *redistributive* program, rather than as an analysis of social insurance against risk. Similar intuitions appear regardless because the curvature of the utility function gives the planner a motive to smooth consumption across individuals.

Simple case: the means test does not affect stigma. Equation (6) nests the case where there are no inframarginal effects and $\eta_m = 0$. In that case, the planner seeks to equalize welfare-weighted marginal utility across people. Due to the concavity of the utility function, she does not give the entire budget to the lowest type so the solution is interior. On the other hand, as the welfare weight schedule is decreasing, the planner values the marginal utility of the lower types more than that of higher types. The solution will thus depend on the utility function's curvature as well as the schedule of welfare weights.

Stigma costs. As is standard, our optimality condition is governed by an envelope argument: people who take up the program due to a reduction in costs are just indifferent. They impose a fiscal externality because they take up the program, thus reducing how much the planner can transfer to others, but they experience no first-order utility gain. In this setting, the planner has an additional way to raise the utilities of people who always take up the program. She can reduce stigma by raising the eligibility threshold. Since these people are not indifferent, they do experience first-order utility gains. A change in the eligibility threshold itself also has first-order implications for social welfare, as those who are newly eligible enjoy the benefit of the program.

In this way, our model embeds a key trade-off in policies that reduce stigma either as an end goal or incidentally. On the one hand, reducing stigma can give a fiscal externality by raising take-up for people who do not value the program. But people who would take-up anyway will enjoy a first-order gain.

[42]This statement refers to a necessary but possibly not sufficient condition for an interior optimum. We describe the statement in more detail in Appendix F.

4.2 Incorporating Information Frictions

In this section, we present our main optimality condition. We now permit some share of consumers not to optimize. Assume share $s \in [0,1]$ of consumers are "stigma(-only)" agents who behave as in the previous section. We introduce share $(1-s)$ of consumers who suffer from optimization frictions: raising the eligibility threshold for these consumers raises take-up because it increases information. We call these consumers "information(-only)" agents. We assume that the probability of being a stigma-only agent is independent of m.

Let the take-up probability for stigma agents be p^s and for information agents be p^i. For information agents, costs are distributed:

$$c = \begin{cases} \infty, \text{ with probability } 1 - p^i(m) \\ 0, \text{ with probability } p^i(m) \end{cases}, \tag{7}$$

for continuously differentiable $p^i(m)$. Put another way, information agents always participate if they know about the program. If they know about the program, the cost they face is 0 and they take it up. An agent's awareness does not depend on her type.

Let η^i_m and η^s_m represent the take-up elasticities with respect to the eligibility threshold for information and stigma agents, respectively. In Appendix F, we set up the planner's problem and obtain the following optimality condition:

Proposition 2. *At an interior optimum, m and B satisfy:*

$$\frac{\overbrace{\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)u(B)}^{\textit{Welfare-weighted WTP of newly eligible}} + \overbrace{\frac{1-s}{p^s}\eta^i_m u(B)}^{\textit{WTP of info agents who now take-up}} + \overbrace{\frac{s}{p^i}\frac{m}{p^s}\int_0^{u(B)}\frac{\partial H(c|m)}{\partial m}dc}^{\textit{WTP for lower c}}}{u'(B)}$$
$$= \frac{\overbrace{B\left[\frac{(1-s)}{p^s}\left(\eta^i_m + 1\right) + \frac{s}{p^i}\left(\eta^s_m + 1\right)\right]}^{\textit{Fiscal externality}}}{\frac{p^s}{sp^s + (1-s)p^i}s\eta_B + 1}. \tag{8}$$

The key difference between Equation (8) and Equation (6) from the benchmark case is the appearance of the term $\frac{1-s}{p^s}u(B)$ on the LHS. This new expression gives the utility gains of the information-only agents who were previously eligible but learn about and join the program when the eligibility threshold increases. These information agents are *not* subject to an envelope condition like the stigma agents because they do not initially optimize. As a result, they would be willing to pay for the full benefit of the program, now that they know about it. Note that because we assume the distribution of stigma agents is independent of the welfare schedule, there is no welfare weight adjustment to the new inframarginal term. Other differences between the equilibrium conditions constitute simple rescaling factors to adjust for the share of the population that is affected by stigma and information.

If $s = 1$, Equation (8) nests Equation (6). Moreover, if $s = 0$ so all people are information agents, the distribution of stigma costs captured by γ and $\frac{\partial H}{\partial m}$ no longer enter the expression; in this case, since B does not affect the take-up rate for information agents, η_B no longer enters the planner's optimality conditions. Then, the information-only case has the especially parsimonious expression:

$$\frac{u(B)\left(\frac{\lambda_m}{\lambda_{\text{avg}}} + \eta_m\right)}{u'(B)} = B\left(1+\eta_m\right). \tag{9}$$

The LHS encodes the welfare-weighted WTP for the newly eligible types and the WTP for the inframarginal types who now take-up. The RHS captures the fiscal externality from take-up.

Empirical implementation. In Appendix F, we show that a second-order Taylor expansion as in Gruber (1997) gives:

$$1+\frac{1}{2}\rho \approx \frac{\frac{(1-s)}{p^s}\left(\eta_m^i+1\right)+\frac{s}{p^i}\left(\eta_m^s+1\right)}{\left(\frac{p^s}{sp^s+(1-s)p^i}s\eta_B+1\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s+\frac{1-s}{p^s}\eta_m^i-\frac{sm}{p^i}\frac{\partial\gamma}{\partial m}+\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i}+\frac{1-s}{p^s}\right)\right)}, \tag{10}$$

for coefficient of relative risk aversion $\rho := -\frac{u''(B)}{u'(B)}B$.[43]

Equation (10) is the main condition that we examine empirically. Relative to Equation (8), this expression substitutes out the utility function $u(B)$ and derivative of the distribution of stigma costs with respect to the eligibility threshold $\frac{\partial H(\cdot|m)}{\partial m}$, which reduces the number of parametric assumptions we need to make. In their place, we add the risk aversion parameter ρ, which is more familiar to calibrate. The upshot is that we can take Equation (10) to the data by estimating η_m^i and η_m^s for a given social program. While estimating separate elasticities by type may seem daunting, Section 5 shows how the combination of our empirical approaches yields estimates of these parameters.

4.3 Policy Implications

We next derive sufficient conditions for when inframarginal effects unambiguously serve as a force to increase the eligibility threshold. Along the way, we derive an empirical test for whether the eligibility threshold is set suboptimally low. We proceed informally, to emphasize intuition, but present a formal treatment in Appendix G.

We study how the *naïve* planner's choice of (B, m) will differ from the *sophisticated* planner's choice. We define the naïve social planner as one who sets policy according to Equation (8) but: (i) erroneously believes that inframarginal effects arising from either agent are zero ($\eta_m^i = \eta_m^s = 0$), and (ii) does not realize that the eligibility threshold affects stigma. On the other hand, the sophisticated planner sets policy optimally according to Equation (8) and knows the true values of inframarginal effects.

[43]We use the coefficient of relative risk aversion in B, evaluated at the sub-utility $u(B)$, since c is just an additive shifter and does not affect curvature.

In this section, we hold fixed the parameters $\{p^i, p^s, \lambda_\theta, s, \gamma, \eta_B, u(\cdot)\}$. We state some basic assumptions in Appendix G that rule out edge cases. In the following comparative statics, we assume that the coefficient of relative risk aversion $\rho \geq 1$. We also employ the following non-trivial assumption that merits discussion:

Assumption 1. $\frac{\partial \gamma}{\partial m} = \frac{\partial (E[c|c<u(B),m]/u(B))}{\partial m} \leq 0.$

This assumption imposes that the average cost-benefit ratio, conditional on taking up the program, does not rise with a looser eligibility threshold. As m rises, stigma costs fall, which tends to reduce γ. On the other hand, new people may take up the program. Since they are nearly indifferent, they have relatively high draws of c, which raises γ. The assumption is true as long as the mass of just-indifferent people who newly sign up for the program as a result of reduced stigma costs do not raise the cost-benefit ratio more than the reduction in inframarginal stigma costs. For instance, in the case where costs are distributed uniformly, $\frac{\partial \gamma}{\partial m} = 0$.

Assumption 1 is sufficient but not necessary. In Appendix G we give a substantially weaker but less concise necessary condition. We also prove that the assumption always holds for costs that are distributed normally or exponentially.

We then arrive at the following proposition.

Proposition 3. *Inframarginal effects raise the sophisticated planner's eligibility threshold relative to the naïve planner's eligibility threshold, if stigma types' take-up is less elastic to the threshold than information types' ($\eta^s_m \leq \eta^i_m$).*

Due to the planner's budget constraint, this proposition equivalently implies that, if the same hypotheses hold, the sophisticated benefit size B is smaller than the naïve benefit. The condition $\eta^s_m \leq \eta^i_m$ in Proposition 3 is sufficient but not necessary. There exist cases with stigma types who are more elastic than information types where the sophisticated eligibility threshold is larger than the naïve threshold.

Proposition 3 implies an empirical test for whether the eligibility threshold is unambiguously too low. The threshold value determining whether the statement is sharp is $\eta^s_m = \eta^i_m$: for any $\eta^s_m \leq \eta^i_m$, the naïve planner unambiguously sets the eligibility threshold too low. Accordingly, testing $H_0 : \eta^s_m > \eta^i_m$ permits the analyst to determine whether the eligibility threshold should optimally rise. If the test fails to reject that $\eta^s_m > \eta^i_m$ then the threshold may still be set too low. But if the test does reject, then the normative conclusions are unambiguous if one accepts the assumptions in the model. We conduct this test in the following section.

We discuss both possible cases to aid intuition.

Example 1 ($\eta^s_m \leq \eta^i_m$). We first consider the case where inframarginal effects are driven by reduced information frictions, i.e. the hypotheses to Proposition 3 hold. Then, for a small change in the eligibility threshold, more people who take up capture the full benefit than people who take up and are just indifferent. The naïve planner thus employs a version of Equation (8) that unambiguously underestimates the welfare gains of a small increase in the threshold.

Example 2 ($\eta^i_m < \eta^s_m$). The policy implications in the second case are not sharp. There exist parameterizations in which the naïve planner sets the eligibility threshold too high or too low. To understand why, note that a high η^s_m introduces two forces. On the one hand, if just-indifferent stigma types are very sensitive to the eligibility threshold and many newly take-up, a small increase in the threshold introduces a large fiscal externality. The naïve planner will ignore this cost, a force pushing the sophisticated planner to lower the threshold. On the other hand, the fact that some stigma types are very sensitive implies that stigma types who always take up will enjoy large reductions in stigma from a small change in the threshold. Put another way, because many people react strongly to the threshold, that implies the threshold has a large effect on stigma. This logic of course requires a connection between the stigma gains from those who always enroll and the stigma gains from those who are just indifferent. That is precisely the role that Assumption 1 plays: it gives a sufficient value for how similar the stigma responses between those who always enroll and the indifferent types need to be.

To summarize, in this case, the naïve planner neglects two forces that accrue from reducing stigma. The net contribution of these forces (as well as the gains to information types) is unsigned.

We note that these normative conclusions are *not* sensitive to the share of stigma or information agents s. In fact, as $s \to 0$, Appendix G shows that we can substantially relax Assumption 1. In the limit case where $s = 0$, neither of Assumption 1 or the condition $\eta^s_m \leq \eta^i_m$ is required at all (as is intuitive, since $\frac{\partial \gamma}{\partial m}$ vanishes from the optimality condition). Thus, the case with $s > 0$ is *conservative* for the model's normative conclusions. If all people are information types, then Proposition 3 holds under weaker conditions.[44]

4.4 Discussion of Model Assumptions

Our framework yields a tractable benchmark for welfare analysis that we can take to the data. Even so, it involves several stark assumptions.

Lump-sum benefits. Many social programs, including SNAP, have non-linear benefits schemes that vary based on income and household size. If the planner could give non-linear benefits, she might extend a small benefit to a larger share of people, to reduce stigma costs and boost take-up without incurring as large a fiscal externality. Our model abstracts from this choice, but we view B as representing the (appropriately weighted) average benefit given to inframarginal types. Relatedly, we assume that people are perfectly informed about the benefit to which they are entitled.[45]

Campaigns to inform or destigmatize. Our model does not feature an instrument by which the planner

[44] The statement that welfare analysis is conservative if $s = 1$ does *not* mean that the planner should increase the eligibility threshold by a greater amount as $s \to 0$. It means that Proposition 3 holds without Assumption 1. Proposition 3 deals with infinitesimal changes in the eligibility threshold. Analyzing non-marginal changes requires more structure, which we develop in Section 5. Moreover, if $s = 1$ and the reduction in always-takers' stigma costs are large, then that serves as another motive to increase the eligibility threshold.

[45] An alternative model, as in Finkelstein and Notowidigdo (2019), casts information frictions as a noisy (mis)perception of benefits. Even if misperceptions are symmetric, correcting them can still increase take-up in our model, since benefits enter a concave utility function. The welfare implications of this model are different: the utility gain to the newly enrolled inframarginals is bounded above in relation to the size of the misperception, while the previously enrolled do not gain.

can spread information about SNAP or reduce the stigma of SNAP directly. Even if the planner has other means of spreading information or reducing stigma (i.e., the eligibility expansion is not the most effective way to do so), the model highlights that eligibility expansions could nevertheless affect information and stigma. The planner must choose *some* eligibility threshold for her means-tested program. She must contend with the trade-offs inherent in setting the threshold.

Identical take-up probabilities. If in fact p varies with θ, it is possible to undo some of our normative conclusions. For example, suppose most of the increase in take-up from inframarginal effects is concentrated in types for whom λ_θ is small. Then inframarginal effects can yield a smaller transfer to the types for whom λ_θ is large. A fruitful extension of the model could consider different take-up probabilities.

5 Welfare Analysis

5.1 Set-Up

In this section, we combine the model and empirics to show that under reasonable assumptions, inframarginal effects meaningfully increase the optimal eligibility threshold. To make this point, we compare the optimal means test under the naïve social planner, who sets policy optimally but erroneously believes $\eta_m = 0$, to that under a sophisticated social planner who understands that $\eta_m > 0$. The model does not capture every relevant economic force, so we view this welfare analysis as illustrative.

First, we implement the empirical test we proposed in Section 4. This test signs whether inframarginal effects should cause the eligibility threshold to rise. We then extend our analysis to quantify how much the inframarginal effects we measure should affect optimal policy. We implement a novel method to quantify the planner's mistake using only local policy analysis. Finally, we impose more structure to make global claims about the optimal eligibility threshold.

5.2 Decomposition

A takeaway from our model is that welfare effects of the eligibility threshold depend on the mechanism underlying inframarginal effects. Figure 6 suggests that, although increasing the eligibility threshold appears to reduce stigma, this effect does not drive the results in Section 2. We use the model and our empirical estimates to decompose inframarginal effects between information and stigma. This decomposition quantifies the mechanisms underlying inframarginal effects and is therefore useful in its own right. Moreover, it gives estimates of η_m^i and η_m^s, the inframarginal take-up elasticities for information and stigma types. With these in hand, we can directly implement our empirical test, proposed in Section 4, of whether inframarginal effects should rise.

The key piece of model structure that we leverage is that all agents are either stigma or information types. In that case, it is an identity that:

$$\frac{\partial p}{\partial m} = s \times \frac{\partial p_s}{\partial m} + (1-s) \times \frac{\partial p_i}{\partial m}, \tag{11}$$

and manipulations give:

$$\frac{\partial \ln \text{(Number Enrolled)}}{\partial m} = s\frac{1}{p}\left(\frac{\partial p_c}{\partial c}\frac{\partial c}{\partial m}\right) + \frac{1}{p}(1-s) \times \frac{\partial p_i}{\partial m} + \frac{\partial \ln \text{(Number Eligible)}}{\partial m}. \tag{12}$$

That is, the increase in take-up after a change in the means test can be decomposed into the increase in take-up among stigma agents (mediated by their change in stigma costs) and the increase in take-up among info agents. We are able to estimate a demographic-cell level version of Equation (12) by combining our various datasets. Specifically, we estimate:

$$\frac{\partial \ln\text{(Number Enrolled)}}{\partial m}_d = \frac{1}{p}\left(\frac{\partial c}{\partial m}_d s_d\right) \times \beta^s + \frac{1}{p}(1-s)_d \times \beta^i + \epsilon_d, \tag{13}$$

which replaces unobserved terms in Equation (12) with coefficients to be estimated, β^i and β^s.[46] For each demographic cell d, we estimate $\frac{\partial \ln\text{(Number Enrolled)}}{\partial m}_d$ using the QC data and instrumental variables regressions analogous to Equation (3). We estimate the cell-level effect of changing eligibility on cost $\frac{\partial c}{\partial m}$ using the second-order stigma results from the experiment. We extract cell-level values of s using the FSPAS.

Noting that $\beta^s = \frac{\partial p^s}{\partial c}$ and $\beta^i = \frac{\partial p^i}{\partial m}$, Equation (13) permits us to recover estimates of elasticities $\hat{\eta}^s_m$ and $\hat{\eta}^i_m$. We provide more estimation details (including details of bootstrapping the estimates on both the right- and the left-hand sides of the equation and jointly weighting by cell sizes across datasets) in Appendix D.

This exercise yields that the inframarginal effect principally arises from information frictions, rather than stigma (Table 5). We are unable to reject the null that $\beta^s = 0$ (Row 1) but robustly reject that $\beta^i = 0$ (Row 2). Combining these parameters with our other empirical estimates, we have $\eta^s_m \approx 0$ (Row 3), depending on the specification, and $\eta^i_m > 0$ (Row 4).

The upshot of conducting the formal decomposition is that it gives the machinery to test $H_0 : \eta^s_m > \eta^i_m$ empirically. We conduct a (conservative) two-sided test of the null that $\eta^s_m = \eta^i_m$, and we reject this null at the 5% significance level in all specifications (Row 6). This implies that $\eta^s_m \leq \eta^i_m$. Therefore, Proposition 3 holds: in our setting inframarginal effects imply that the eligibility threshold is too low, assuming the social planner does not presently account for them when setting SNAP's eligibility threshold.

5.3 Calibration

Next we calibrate the parameters necessary to quantify the implications of inframarginal effects. Informed by the evidence in the previous section, we henceforth assume that information frictions are the dominant

[46]Equation 13 also uses the assumption that $\frac{\partial \ln \text{Number Eligible}}{\partial m} = 0$, if our identifying variation is valid; put another way, changing m should not change the number of *inframarginals* who are eligible.

mechanism underlying inframarginal effects.[47] Additionally, for mathematical simplicity, we also assume that at the optimum, $p^i = p^s$. We state the assumptions as follows:

Assumption 2. At the planner's solution, $p^i = p^s$ and $\frac{\partial p^s}{\partial m} = 0$.

From Assumption 2, we proceed using $\eta_m^i = \eta_m(1-s)^{-1}$. We use the full-sample estimate of the effect of eligibility expansions on the 0–130% take-up rate from Table 4.

Now, we discuss calibration of other parameters, summarized in Table 6. As is common in welfare analysis, some of the economic primitives have a high degree of uncertainty. As a result, we show robustness to the particular choice of parameter.

Cost-benefit ratio of taking up the program (γ). Although our evidence suggests that information frictions are the dominant mechanism underlying inframarginal effects, reductions in stigma costs among eligible people who would have taken up SNAP regardless can still confer welfare gains when the means test increases. Thus, we need some assumptions on stigma costs among stigma agents. We choose a conservative assumption which simplifies the analysis: $c \sim U[0, \overline{A}(m)]$ for $\overline{A} > u(B)$.[48] We then have that $\gamma = \frac{1}{2}$, as $E[c|c < u(B), m] = \frac{1}{2}u(B)$.

With this assumption, $\frac{\partial \gamma}{\partial m} = 0$: no welfare gains accrue to inframarginal stigma agents. We see this assumption as therefore being conservative: if inframarginal effects also deliver utility to inframarginal types who already take up the program, then that only increases the planner's motive to raise the eligibility threshold. Intuitively, this assumption lets us avoid needing to compute enrollees' willingness to pay for reduced stigma — an interesting and policy-relevant parameter that future work should explore.

Take-up elasticity with respect to benefit size (η_B). The SNAP benefits schedule B is set nationally. As a result, we cannot use an event-study design to estimate η_B. Instead, we collect estimates of the typical elasticity of take-up with respect to benefit size for related programs. Krueger and Meyer (2002) review papers estimating η_B for UI and worker's compensation and conclude that, for these programs, η_B ranges from 0.3 to 0.6. We choose $\eta_B = 0.5$ as a sensible midpoint and show robustness to other values.[49]

Other parameters. We use $\rho = 3$ as a benchmark.[50] For the share of stigma types s in the population — informed by our analysis of the FSPAS — we use $s = 0.4$ as a benchmark. For the means test m, we take the population-weighted average of states' share eligible across years (accounting for varying eligibility thresholds) to obtain $m^* = 0.27$ in 2016. For the take-up probability p^*, our data from 2016 suggest the take-up probability is $p^* = 0.53$.[51]

[47] Note that this is different from saying that all agents are information agents; we continue to allow some share of the population $s \in [0, 1]$ to face stigma costs, but we assume that these agents' take-up decision does not respond to changes in the means test.

[48] Economically, this assumption posits that: (i) changing m does not change the shape of the cost distribution, and (ii) there exist people for whom the take-up cost exceeds the utility gain.

[49] Kroft (2008) also cites the Krueger and Meyer (2002) review and uses $\eta_B = 0.5$. Auray and Fuller (2020) is an example of a recent paper that finds a similar η_B in later years. In their data from 2002–2015, $\eta_B = 0.63$ (SE: 0.23), where η_B is the elasticity of UI take-up with respect to the replacement rate.

[50] Chetty and Finkelstein (2013) note that this parameter is notoriously difficult to calibrate, but review other papers that test values of $\rho \in [1, 4]$ (e.g., Gruber, 1997).

[51] This number is below the number the USDA reports because our denominator includes some people who are not eligible for

5.4 Local Policy Analysis

A standard problem with conducting empirical analysis of social optimality conditions is that if one rejects that the optimality condition exactly holds, it is difficult to estimate the magnitude of the planner's mistake. For concreteness, imagine one has data to statistically reject that the LHS and RHS of a standard Baily (1978)-Chetty (2006) condition exactly coincide in the analysis of a given social insurance program. Is the planner's mistake large or small? A typical approach is to impose structure so that the researcher can extrapolate agents' behavior away from equilibrium; we will take this approach in the next section. First, we propose a new method for estimating the size of the planner's mistake that uses only the local optimality conditions. The advantage of this approach is that it does not require extra parametric structure. The disadvantage is that it does not permit making policy recommendations like what the optimal eligibility threshold should be.

In short, we estimate the magnitude of the naïve planner's mistake by studying the implied value of ρ that would be required to make the current means test optimal (when inframarginal effects are present). We establish that, if the planner assumes $\eta_m = 0$ but otherwise optimizes according to the theory, she will treat people as if they are much less risk averse than they really are.[52]

Our approach is as follows. We assume some ground-truth value of ρ, say $\rho = 3$. Consider the naïve planner who chooses m and B to solve Equation (10) assuming $\eta_m = 0$ and that $p^i = p^s$. Given $\eta_m \neq 0$, what is the implied $\tilde{\rho}$ that keeps the optimality condition equated? We use the following algorithm:

1. Obtain inverse-optimum weights: Assuming $\eta_m = 0$, solve for $\lambda_m/\lambda_{\text{avg}}$ that satisfies Equation (10).
2. Obtain implied $\tilde{\rho}$: Given the inverse-optimum weights $\lambda_m/\lambda_{\text{avg}}$ and *true* value of η_m, solve for the $\tilde{\rho}$ that satisfies Equation (10).

Intuitively, because the planner ignores η_m, she treats people "as-if" they have risk aversion $\tilde{\rho}$, when they really have risk aversion ρ. Put another way, there is some value of $\tilde{\rho}$ that satisfies the optimality condition even under the (incorrect) assumption that $\eta_m = 0$. We focus on the value of $100 \times \frac{\tilde{\rho}-\rho}{\rho}$, which is a measure of the *bias* in the coefficient of relative risk aversion.

Results. We show that the magnitude of the planner's mistake can be substantial for the range of η_m that we estimate (Figure 7). The x-axis plots values of η_m. On the y-axis, we plot the bias in the "as-if" risk aversion parameter relative to the true risk aversion parameter, assuming $\rho = 3$. If $\eta_m = 0$, there is no bias: the naïve and sophisticated solutions coincide by construction. As η_m grows, the bias rises; for $\eta_m = 0.05$,

SNAP due to work requirements, asset thresholds, or other tests; moreover, it is not clear that the USDA number includes people with incomes above 130 if they live in states with an eligibility threshold beyond 130. We use the number for illustrative purposes in this exercise, but the results are not sensitive to adjusting the equilibrium p^*.

[52] We use the standard interpretation of ρ as risk-aversion. However, it also corresponds to the planner's unweighted valuation of transferring B to someone who is ineligible from someone who takes up (has a benefit of B). To see this, note that $\rho = -B\frac{u''(B)}{u'(B)} \approx \frac{u'(0)-u'(B)}{u'(B)}$.

the bias is about 10%. For our primary estimate of $\eta_m = 0.13$, we find that the bias can be quite large: the naïve planner's solution will treat people as if they are about 30% less risk averse than they really are. We show robustness to parameterizations in Appendix E.

Intuitively, the planner who ignores inframarginal effects transfers too much to inframarginal types who already take up the program. She overvalues inframarginal types' marginal utility and undervalues the gain in utility from those who would take up the program if she raised the eligibility threshold. As a result, she optimizes as if the coefficient of relative risk aversion were smaller than it really is.

MVPF. Another approach to gauging the size of the naïve planner's mistake with limited parametric assumptions is the Marginal Value of Public Funds (MVPF) (Hendren and Sprung-Keyser, 2020). We study the MVPF of an eligibility expansion in Appendix E. We document the potential for the naïve planner to have substantial bias in her estimate of the MVPF.

5.5 Global Policy Analysis

In this section, we impose structural assumptions to extrapolate take-up probabilities and welfare weights away from what we observe in equilibrium. Appendix E provides details on our parameterizations. With these assumptions, we numerically invert Equation (10) to solve for the optimal m^{opt} and B^{opt} as a function of η_m.

5.5.1 Results

We first ask how much larger is m^{opt} relative to today's $m^* = 0.27$. Because we use the inverse optimum approach to calibrate the welfare weights, Proposition 3 guarantees that the optimal m^{opt} exceeds today's m^*: $m^{\text{opt}} - m^* > 0$.

We present the percent increase in m^{opt} relative to m^*, i.e. percent increase $:= 100 \times \frac{m^{\text{opt}} - m^*}{m^*}$. This value represents the percent increase in the optimal eligibility threshold relative to today's threshold. Because the threshold is measured in terms of the share eligible, it equivalently represents the percent increase in the share of people who should be eligible relative to today.

We present our estimates of the percent increase in m as a function of η_m (Figure 8A) for both $s = 0.4$ (black line) and $s = 0.8$ and $s = 0$ (gray dashed lines). By construction, if $\eta_m = 0$, we find the optimal eligibility threshold coincides with today's threshold. As η_m rises, the optimal m rises too. At our preferred value of $\eta_m = 0.13$, about 13% more people should be eligible than are eligible today.

We also show the optimal take-up rate (Figure 8B). Because of our social planner's fixed budget, the optimal take-up rate is not monotonic: increases in m require decreases in B, which, through η_B, decrease take-up. At some point, however, take-up falls enough that those on the program are granted larger B, and take-up begins to rise again. This dynamic does not exist for $s = 0$ since η_B has no effect for information agents.

While Panel A and Proposition 3 show that the sophisticated planner will expand eligibility beyond today's m, Panel B highlights that we cannot conclude that take-up today is suboptimally low.[53] The naïve planner erroneously believes take-up will fall *more than it actually would* for a small increase in m because she does not account for η_m and only accounts for η_B. She therefore sets m too low. We show similar conclusions for $\rho \in \{1, 2, 4\}$ (Figure E.3A) and $\eta_B = 0.3$ (Figure E.3B). Notably, the magnitude of ρ does not have a large effect on the percent change in m. In this setting, η_m is much more important than ρ.

The planner has a fixed budget, so the increase in the optimal eligibility threshold and small change in optimal take-up rates imply that the optimal benefit is decreasing in the inframarginal effect, even for non-local changes (Figure E.4). For various s, at our preferred estimate of η_m, the optimal benefit is 5–10% lower than the current optimum.

A weakness of our numerical approach is that we assume that $\rho = 3$ both in today's equilibrium and also at an optimum, but that third-order utility terms vanish (in order that Equation (10) holds). We conduct a second exercise where we assume a quadratic utility function that imposes that $\rho = 3$ at today's B (Figure E.5).[54] This exercise gives similar results, although the magnitudes of the increase in eligibility are attenuated because risk aversion changes rapidly for quadratic utility.

6 Conclusion

This paper documents the existence of inframarginal effects in SNAP. We find that the inframarginal effects arise from increased information after states relax eligibility thresholds, but our online experiment also finds that relaxing eligibility thresholds can reduce stigma. We develop a general model for incomplete take-up of social welfare programs when the planner can control program eligibility. We apply our model to SNAP and assess the implications for the optimal eligibility threshold given the inframarginal effects. Because the information mechanism dominates, inframarginal effects unambiguously increase optimal SNAP eligibility.

All social programs, even universal ones, make some determination about eligibility. This threshold is often chosen by the planner and thus is not an exogenous feature of the policy environment. As a result, our normative insights have applications in many areas in public economics. When inframarginal effects are present, our theoretical framework highlights that they may serve as a motive to raise the eligibility threshold. Future work could enrich the model to include a larger set of policy instruments and more heterogeneity in individual responses.

References (including for Appendix)

Aizer, Anna, "Low Take-Up in Medicaid: Does Outreach Matter and for Whom?," *American Economic Re-*

[53] Note that take-up depends on the benefit size as well as the eligibility threshold, so the higher threshold does not necessarily imply higher take-up on net.

[54] Together with $u(0) = 0$, this assumption yields that utility is: $u(B) = -(B-k)^2 + k^2$ for $k := \frac{\rho+1}{\rho} B^*$.

view, 2003, *93* (2).

— **and Jeffrey Grogger**, "Parental Medicaid Expansions and Health Insurance Coverage," Working Paper 9907, National Bureau of Economic Research 2003.

Al-Athari, Faris Muslim, "Estimation of the Mean of Truncated Exponential Distribution," *Journal of Mathematics and Statistics*, 2008, *4* (4), 284–288.

Almond, Douglas, Hillary W. Hoynes, and Diane Whitmore Schanzenbach, "Inside the War on Poverty: The Impact of Food Stamps on Birth Outcomes," *Review of Economics and Statistics*, 2011, *93* (2), 387–403.

Atasoy, Sibel, "The End of the Paper Era in the Food Stamp Program: The Impact of Electronic Benefits on Program Participation," 2009.

Auray, Stéphane and David L. Fuller, "Eligibility, Experience Rating, and Unemployment Insurance Take-up," *Quantitative Economics*, 2020, *11* (3), 1059–1107.

Bailey, Martha J., Hillary W. Hoynes, Maya Rossin-Slater, and Reed Walker, "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program," National Bureau of Economic Research Working Paper 26942, Cambridge, MA April 2020.

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

Bartfeld, Judith, Craig Gundersen, Timothy M. Smeeding, and James P. Ziliak, eds, *SNAP Matters: How Food Stamps Affect Health and Well-Being*, Stanford University Press, 2016.

Bartlett, Susan, Nancy Burstein, and William Hamilton, "Food Stamp Program Access Study: Final Report," Technical Report, USDA Economic Research Service, https://www.ers.usda.gov/webdocs/publications/43390/30283_efan03013-3_002.pdf?v=0 2004.

Bernheim, B. Douglas and Dmitry Taubinsky, "Behavioral Public Economics," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 381–516.

Bertrand, Marianne, Erzo F. P. Luttmer, and Sendhil Mullainathan, "Network Effects and Welfare Cultures," *Quarterly Journal of Economics*, August 2000.

Bhargava, Saurabh and Dayanand Manoli, "Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment," *American Economic Review*, November 2015, *105* (11), 1–42.

Bronchetti, Erin T., Garret Christensen, and Hilary W. Hoynes, "Local Food Prices, SNAP Purchasing Power, and Child Health," *Journal of Health Economics*, 2019, *68.*

Brooks, Tricia, Sean Miskell, Samantha Artiga, Elizabeth Cornachione, and Alexandra Gates, "Medicaid and CHIP Eligibility, Enrollment, Renewal, and Cost-Sharing Policies as of January 2016: Findings from a 50-State Survey," Technical Report, The Henry J. Kaiser Family Foundation 2016.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb, "Alternative Transformations to Handle Extreme Values of the Dependent Variable," *Journal of the American Statistical Association*, 1988, *83* (401), 123–127.

Bursztyn, Leonardo, Alessandra González, and David Yanagizawa-Drott, "Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia," *American Economic Review*, Forthcoming.

— **and Robert Jensen**, "Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure," *Annual Review of Economics*, 2017, *9* (1), 131–53.

— **, Ingar Haaland, Aakaash Rao, and Christopher Roth**, "I Have Nothing Against Them, But. . .," Technical Report May 2020.

Callaway, Brantley and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," 2020.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *Quarterly Journal of Economics*, August 2019, *134* (3), 1405–1454.

Chandrasekhar, Arun G., Benjamin Golub, and He Yang, "Signaling, Shame, and Silence," Technical Report 25169, National Bureau of Economic Research 2019.

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

_ , "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, April 2008, *116* (2), 173–234.

_ , "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 2009, *1* (1), 451–487.

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

Child Trends, "Key facts about Head Start enrollment," 2018.

Congressional Research Service, "The Supplemental Nutrition Assistance Program (SNAP): Categorical Eligibility," Technical Report R42054 October 2019.

Cunnyngham, Karen, "Reaching Those in Need: Estimates of State Supplemental Nutrition Assistance Program Participation Rates in 2016q," Technical Report, United States Department of Agriculture 2019.

Currie, Janet, "U.S. Food and Nutrition Programs," in Robert A. Moffitt, ed., *Means-Tested Transfer Programs in the United States*, Chicago: University of Chicago Press, 2003.

_ , "The Take Up of Social Benefits," Technical Report 10488, National Bureau of Economic Research, Cambridge, MA May 2004.

Currie, Janet M. and Jeff Grogger, "Explaining Recent Declines in Food Stamp Program Participation," *Brookings-Wharton Papers on Urban Affairs*, 2001, *2001* (1), 203–244.

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

Daponte, Beth Osborne, Seth Sanders, and Lowell Taylor, "Why Do Low-Income Households Not Use Food Stamps? Evidence from an Experiment," *Journal of Human Resources*, 1999, *34* (3), 612–618.

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

East, Chloe N., "Immigrants' Labor Supply Response to Food Stamp Access," *Labour Economics*, 2018, *51* (202-226).

Eck, Chase S., "The Effect of Electronic Benefit Transfer on the Marginal Propensity to Consume Food out of SNAP," 2018, p. 43.

Eslami, Esa, "Trends in Supplemental Nutrition Assistance Program Participation Rates: Fiscal Year 2010 to Fiscal Year 2013," Technical Report, United States Department of Agriculture, Washington, D.C. August 2015.

Federal Register, "Revision of Categorical Eligibility in the Supplemental Nutrition Assistance Program (SNAP)," July 2019, *84* (142), 35570–35581.

Fetter, Daniel K. and Lee M. Lockwood, "Government Old-Age Support and Labor Supply: Evidence from the Old Age Assistance Program," *American Economic Review*, August 2018, *108* (8), 2174–2211.

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

Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers, "Premium Subsidies, the Mandate, and Medicaid Expansion: Coverage Effects of the Affordable Care Act," *Journal of Health Economics*, 2017, *53*, 72–86.

Friedrichsen, Jana, Tobias König, and Renke Schmacker, "Social Image Concerns and Welfare Take-Up," *Journal of Public Economics*, December 2018, *168*, 174–192.

Ganong, Peter and Jeffrey B. Liebman, "The Decline, Rebound, and Further Rise in SNAP Enrollment: Disentangling Business Cycle Fluctuations and Policy Changes," *American Economic Journal: Economic Policy*, November 2018, *10* (4), 153–176.

Giannarelli, Linda, "What Was the TANF Participation Rate in 2016?," Technical Report, Urban Institute 2019.

_ , Christine Heffernan, Sarah Minton, Megan Thompson, and Kathryn Stevens, "Welfare Rules Databook: State TANF Policies as of July 2016," Technical Report, Urban Institute 2017.

Golosov, Mikhail and Aleh Tsyvinski, "Designing Optimal Disability Insurance: A Case for Asset Testing," *Journal of Political Economy*, 2005, *114* (2), 257–279.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 1997, *87* (1), 192–205.

Haley, Jennifer M., Genevieve M. Kenney, Robin Wang, Victoria Lynch, and Matthew Buettgens, "Medicaid/CHIP Participation Reached 83.7 Percent Among Eligible Children in 2016," *Health Affairs*, 2018, *37* (8), 1194–1199.

Hanna, Rema and Benjamin A. Olken, "Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries," *Journal of Economic Perspectives*, November 2018, *32* (4), 201–226.

Harris, Timothy F., "Do SNAP Work Requirements Work?," *Economic Inquiry*, 2021, *59*, 72–94.

Hastings, Justine and Jesse M. Shapiro, "How are SNAP Benefits Spent? Evidence from a Retail Panel," *American Economic Review*, 2018, *108* (12), 3493–3540.

_ , Ryan Kessler, and Jesse M. Shapiro, "The Effect of SNAP on the Composition of Purchased Foods: Evidence and Implications," *American Economic Journal: Economic Policy*, Forthcoming.

Heckman, James J. and Jeffrey A. Smith, "The Determinants of Participation in a Social Program: Evidence from a Prototypical Job Training Program," *Journal of Labor Economics*, 2004, *22* (2), 243–98.

Hendren, Nathaniel, "The Policy Elasticity," *Tax Policy and the Economy*, 2016, *30.*

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

_ , Camille Landais, and Johannes Spinnewijn, "Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design," *Annual Review of Economics*, Forthcoming.

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

Hoynes, Hilary and Diane Whitmore Schanzenbach, "Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program," *American Economic Journal: Applied Economics*, 2009, *1* (4), 109–139.

_ and _ , "Work incentives and the Food Stamp Program," *Journal of Public Economics*, 2012, *96*, 151–162.

_ , _ , and Douglas Almond, "Long-Run Impacts of Childhood Access to the Safety Net," *American Economic Review*, 2016, *106* (4), 903–34.

Internal Revenue Service, “About EITC,” 2020.

Katz, Michael B., *In the Shadow of the Poorhouse: A Social History of Welfare in America*, Basic Books, Inc., 1986.

Klerman, Jacob Alex and Caroline Danielson, “The Transformation of the Supplemental Nutrition Assistance Program,” *Journal of Policy Analysis and Management*, September 2011, *30* (4), 863–888.

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

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

Kling, Jeffrey R, Jeffrey B Leibman, and Lawrence F Katz, “Experimental Analysis of Neighborhood Effects,” *Econometrica*, 2007, *75* (1), 83–119.

Kreider, Brent, John V. Pepper, Craig Gundersen, and Dean Jolliffe, “Identifying the Effects of SNAP (Food Stamps) on Child Health Outcomes When Participation Is Endogenous and Misreported,” *Journal of the American Statistical Association*, 2012, *107* (499), 958–975.

Kroft, Kory, “Takeup, Social Multipliers and Optimal Social Insurance,” *Journal of Public Economics*, April 2008, *92* (3-4), 722–737.

Krueger, Alan B and Bruce D Meyer, “Labor Supply Effects of Social Insurance,” in Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*, Vol. 4, Elsevier, 2002, pp. 2327–2392.

Leos-Urbel, Jacob, Amy Ellen Schwartz, Meryle Weinstein, and Sean Corcoran, “Not Just for Poor Kids: The Impact of Universal Free School Breakfast on Meal Participation and Student Outcomes,” *Economics of Education Review*, October 2013, *36*, 88–107.

Lindbeck, Assar, Sten Nyberg, and Jorgen Weibull, “Social Norms and Economic Incentives in the Welfare State,” *The Quarterly Journal of Economics*, February 1999, *114* (1), 1–35.

Low, Hamish and Luigi Pistaferri, “Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off,” *American Economic Review*, October 2015, *105* (10), 2986–3029.

Mabli, James, Thomas Godfrey, Nancy Wemmerus, Joshua Leftin, and Stephen Tordella, “Determinants of Supplemental Nutrition Assistance Program Participation from 2008 to 2012,” Technical Report, United States Department of Agriculture Food and Nutrition Service 2014.

Manchester, Colleen Flaherty and Kevin J. Mumford, “How Costly Is Welfare Stigma? Separating Psychological Costs from Time Costs in Food Assistance Programs,” 2012, p. 44.

Marcus, Michelle and Katherine G. Yewell, “The Effect of Free School Meals on Household Food Purchases: Evidence from the Community Eligibility Provision,” Technical Report 2021.

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

Moffitt, Robert, “An Economic Model of Welfare Stigma,” *American Economic Review*, December 1983, *73* (5), 1023–1035.

National Center for Education Statistics, “Digest of Education Statistics,” Technical Report, U.S. Department of Education 2017.

Nichols, Albert L. and Richard J. Zeckhauser, “Targeting Transfers through Restrictions on Recipients,” *American Economic Review Papers and Proceedings*, 1982, *72* (2), 372–377.

Ratcliffe, Caroline, Signe-Mary McKernan, and Kenneth Finegold, “Effects of Food Stamp and TANF Policies on Food Stamp Receipt,” *Social Service Review*, June 2008, *82* (2), 291–334.

—, —, **Laura Wheaton, Emma Kalish, Catherine Ruggles, Sara Armstrong, and Christina Oberlin**, "Asset Limits, SNAP Participation, and Financial Stability," Technical Report, Urban Institute, Washington, D.C. June 2016.

Ruggles, Stephen, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, "IPUMS USA: Version 10.0," www.ipums.org 2020.

Sacarny, Adam, Katherine Baicker, and Amy Finkelstein, "Out of the Woodwork: Enrollment Spillovers in the Oregon Health Insurance Experiment," *American Economic Journal: Economic Policy*, August 2022, *14* (3), 273–295.

Sampford, M. R., "Some Inequalities on Mill's Ratio and Related Functions," *The Annals of Mathematical Statistics*, March 1953, *24* (1), 130–132.

Sommers, Benjamin D. and Arnold M. Epstein, "Why States Are So Miffed about Medicaid — Economics, Politics, and the "Woodwork Effect"," *New England Journal of Medicine*, July 2011, *365* (2), 100–102.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

United States Department of Agriculture Food and Nutrition Service, "SNAP Quality Control Data," https://snapqcdata.net/ 2019.

— , "National School Lunch Program: Participation and Lunches Served," 2020.

— , "WIC 2017 Eligibility and Coverage Rates," 2020.

U.S. Department of Agriculture Economic Research Service, "SNAP Policy Data Sets," https://www.ers.usda.gov/data-products/snap-policy-data-sets/ August 2019.

U.S. Department of Health and Human Services, "Head Start Impact Study. Final Report.," Technical Report, Administration for Children and Families 2010.

Wood, Thomas and Ethan Porter, "The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence," *Political Behavior*, March 2019, *41* (1), 135–163.

7 Figures

Figure 1: Eligibility Thresholds and Program Take-Up

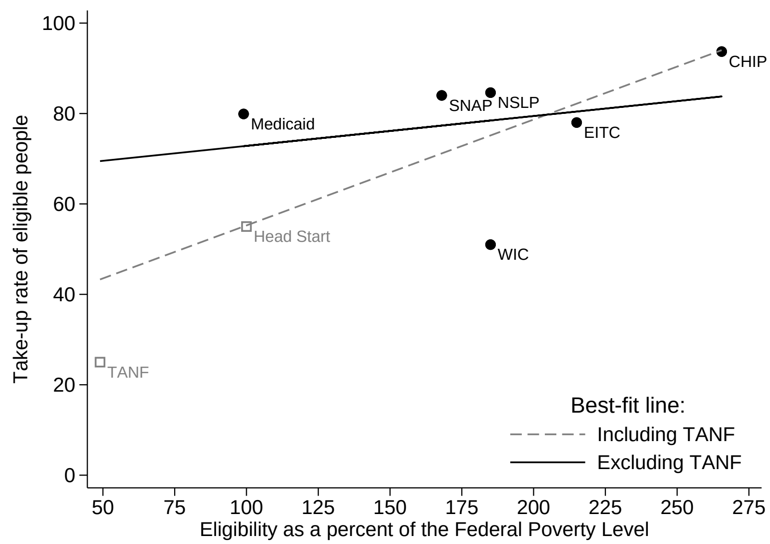

The figure shows income eligibility thresholds as a percent of the Federal Poverty Level (FPL) for the largest U.S. means-tested social programs against estimates of their national take-up rates, compiled from different sources. We plot TANF and Head Start in a separate series because eligibility and take-up rates for these programs are particularly difficult to estimate; see Appendix A for information on constructing these data. Take-up rates are estimated out of the eligible population for each program. In programs with different eligibility thresholds per state, the level plotted is the population-weighted average of those thresholds. The SNAP take-up rate displayed here is higher than that used in our paper because the USDA uses a more involved and restrictive method for assessing eligibility than we do; our empirical results are not affected by a denominator that is too large. Where the eligibility threshold is defined in dollars (e.g., EITC, TANF), the figure shows the threshold as in terms of percent of the FPL for a family of three. Some programs (e.g., WIC, TANF) are restricted to certain subgroups in addition to imposing income thresholds — for example, families only — or have additional requirements. Given Head Start's capacity constraints, additional assumptions were made to estimate a take-up rate. These are also documented in Appendix A.

Figure 2: Literature Review: *AER* and *QJE* papers about Eligibility Criteria

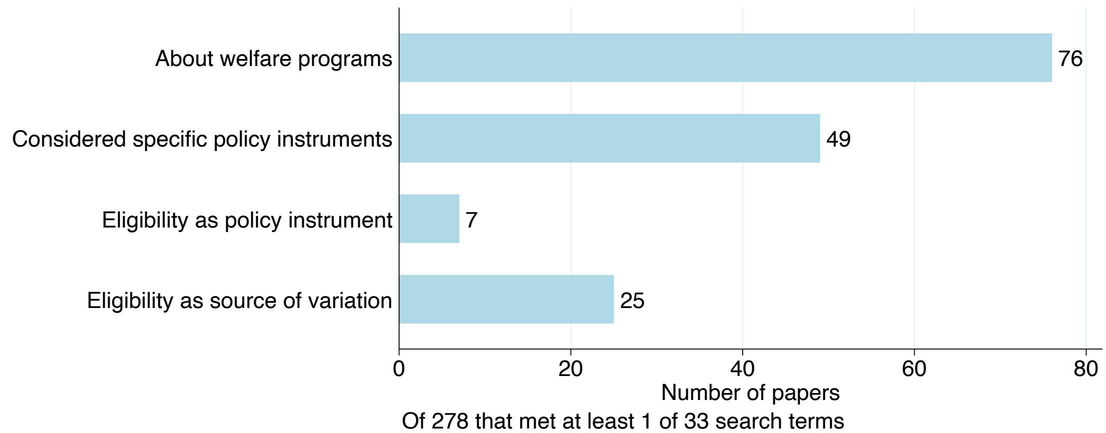

The figure presents the results from our literature review of papers in the *Quarterly Journal of Economics* (2010–2019) and the *American Economic Review* (2010–2018). Appendix A provides details about the sampling frame. The first row shows the total number of papers that we concluded were about welfare programs, after reading the abstract and introduction. The second row shows the number of papers that considered instruments with which the planner could enact optimal policy, e.g. the benefit size or duration. The third row shows the number of papers that considered the eligibility threshold as an instrument with which the planner could enact optimal policy. The fourth row shows the number of papers that use the eligibility threshold as a source of variation with which the authors estimated a treatment effect for the program.

Figure 3: Descriptive Evidence of Higher Inframarginal Enrollment with Expanded Eligibility

(A) Enrollment by Household Income and Eligibility Threshold

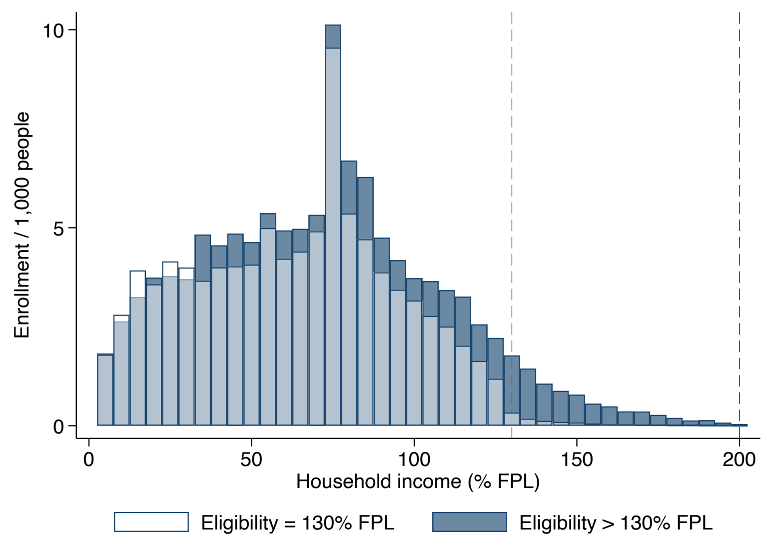

(B) Take-up by Eligibility Threshold

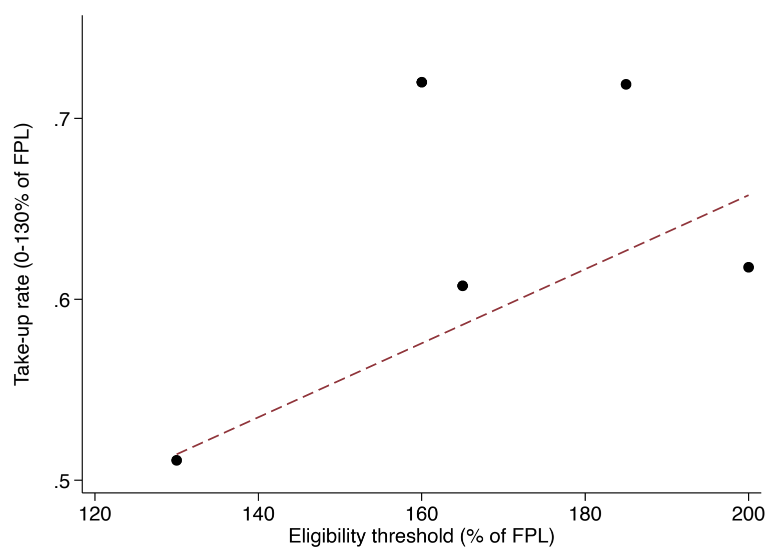

This figure presents the relationship between the eligibility threshold and SNAP take-up and enrollment. Panel A shows SNAP enrollment per 1,000 people in states and years where the eligibility threshold is 130% of the Federal Poverty Level (FPL) versus above 130% (and up to 200%). Each bar takes the number of people in the USDA Quality Control data whose household income is in each income bin, divided by the total population (i.e., all people, with any household income) in all state-years with the indicated eligibility regime. The data are limited to the sample we use in the main event study, and household income is top-coded at 200% FPL. Panel B shows average take-up among those earning 0–130% of the FPL in states with each eligibility threshold observed in the data. The USDA Quality Control data provide estimates of the numerator for the outcome (take-up counts, by state-year), and the Current Population Survey data provide estimates of the denominator (total counts of individuals within this sample).

Figure 4: Event Study of Changes to Eligibility Threshold

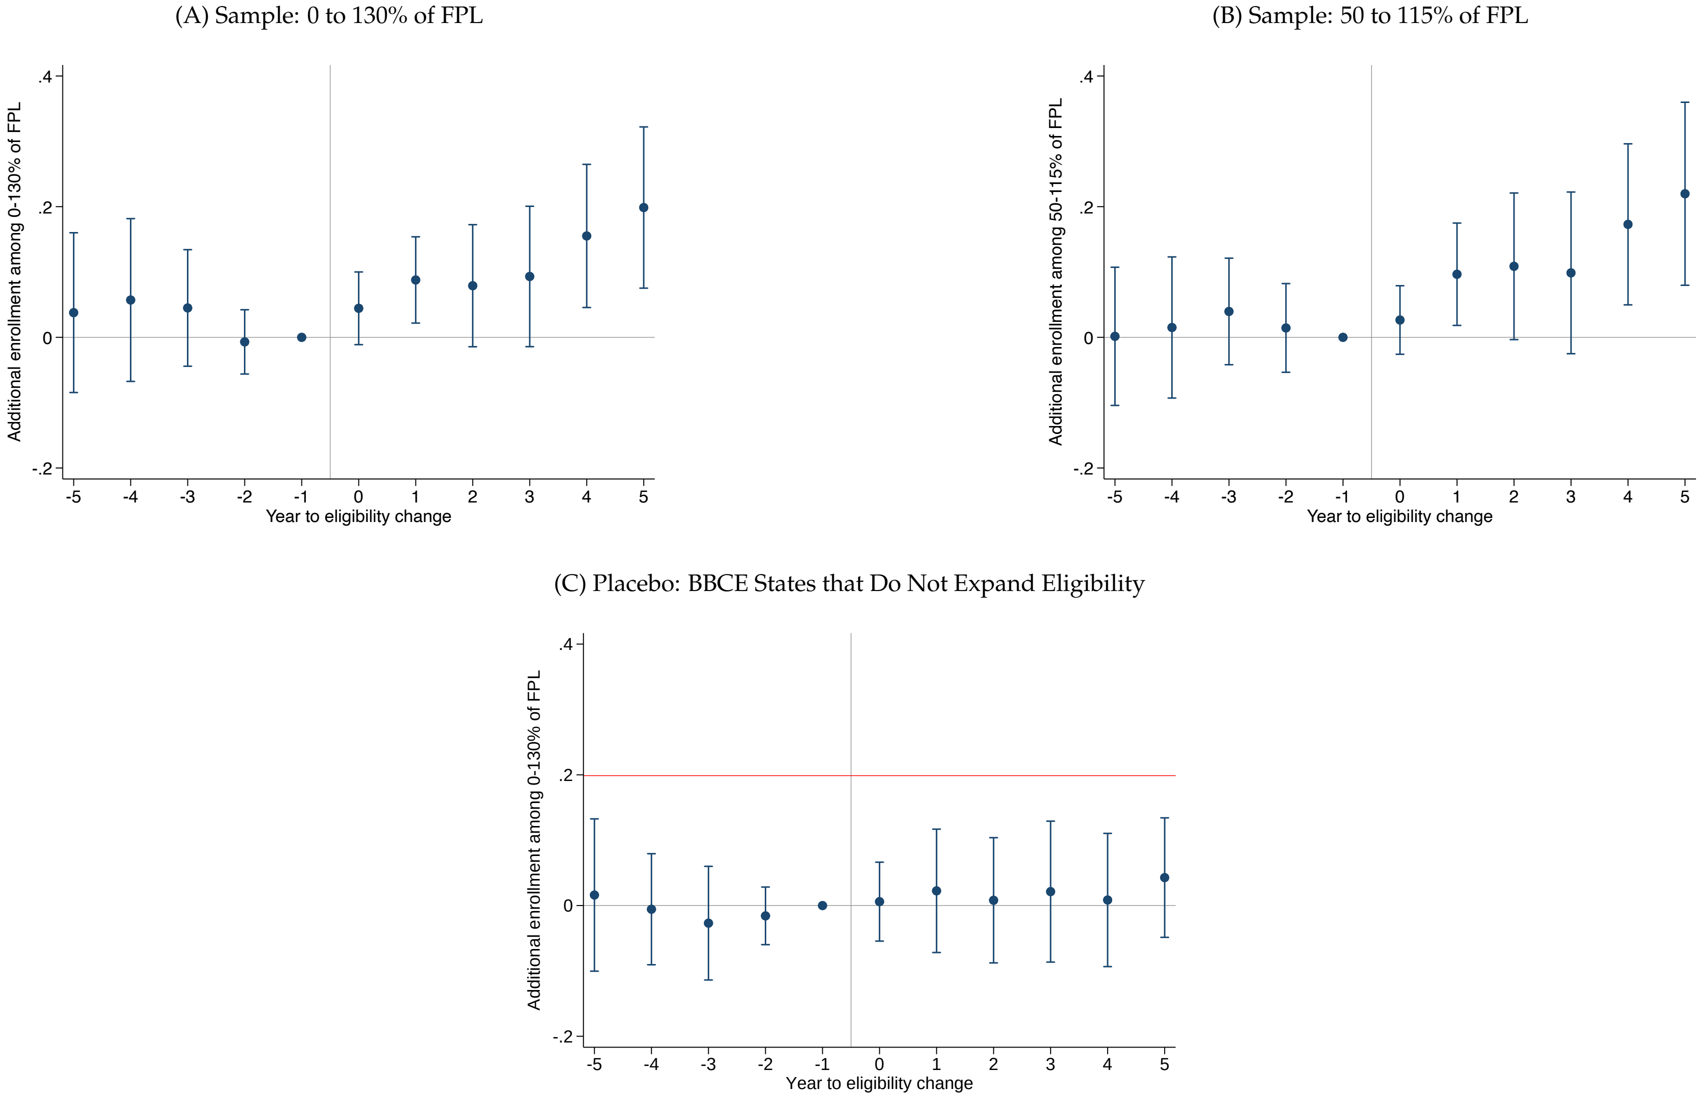

This figure presents the event-study estimate of η (Equation (1)), the effect of the eligibility rate on inframarginal take-up. Panel A presents results for the sample of individuals from 0–130% of the Federal Poverty Level (FPL); Panel B presents results for 50–115% of the FPL. Panel C presents a placebo event study, using the nine states that adopt the Broad Based Categorical Eligibility policy but do not expand eligibility (see Section 2). The red line in Panel C plots the 5-year point estimate from Panel A. The minimum eligibility in all states is 130% of the FPL. Standard errors are robust to heteroskedasticity and clustered by state.

Figure 5: Effect of High-Share Treatment on Stigma

(A) Second-order beliefs

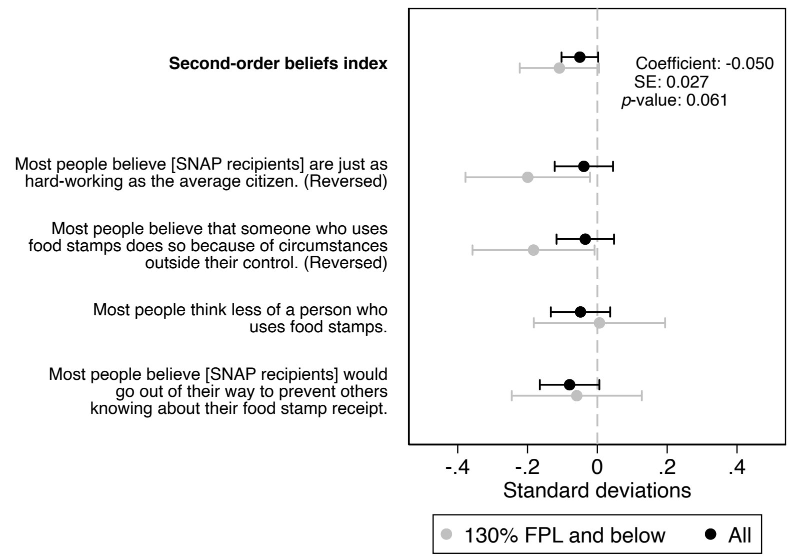

(B) First-order beliefs

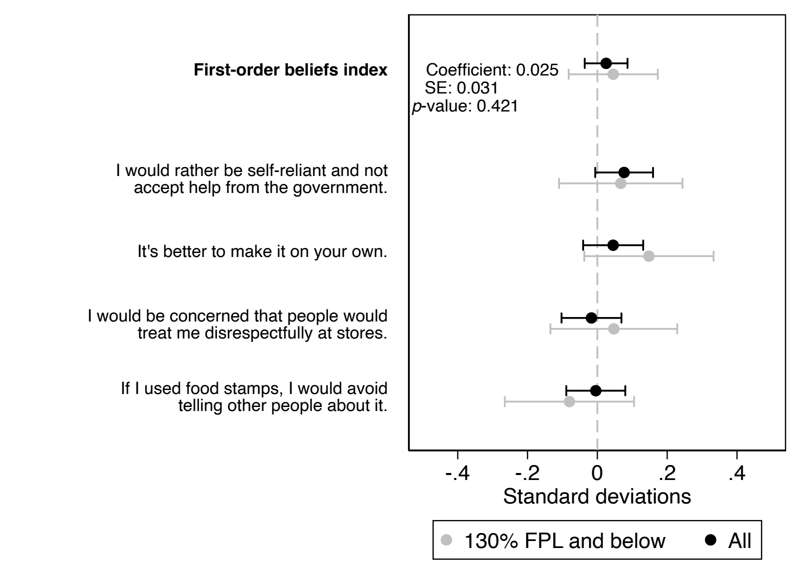

This figure presents results from the online experiment; it shows the effect of the "high-share" treatment (where respondents were randomly given a hint that *increased* their reported beliefs about the share of Americans who are eligible for SNAP) on agreement with each statement in the stigma instrument (Equation (4)). Outcomes marked with "(Reversed)" were reverse-coded so that for all items, a higher score indicates more stigma. The coefficients correspond to a reduced-form (intent-to-treat) estimate and do not account for the amount by which the treatment moved people's beliefs about the share of Americans who are eligible for SNAP. Each outcome is in units of standard deviations, and the indices average the set of outcomes displayed in each panel. Bars plot 95% confidence intervals.

Figure 6: Inframarginal Effects Heterogeneity by Demographic Cell

(A) Share Stigma Type vs. Inframarginal Effects

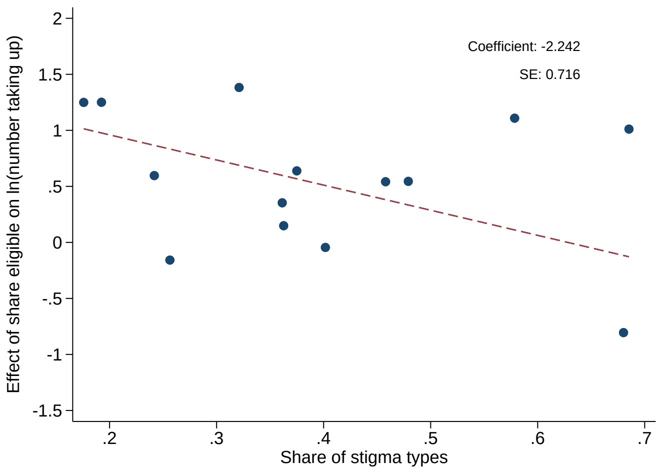

(B) Share Information Type vs. Inframarginal Effects

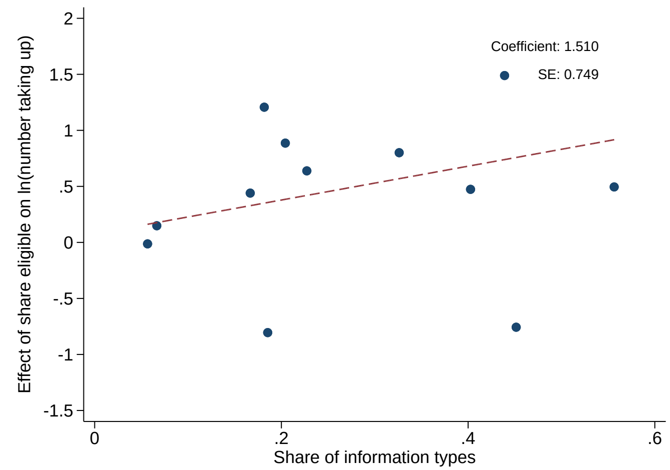

(C) Experiment Treatment Effect vs. Inframarginal Effects

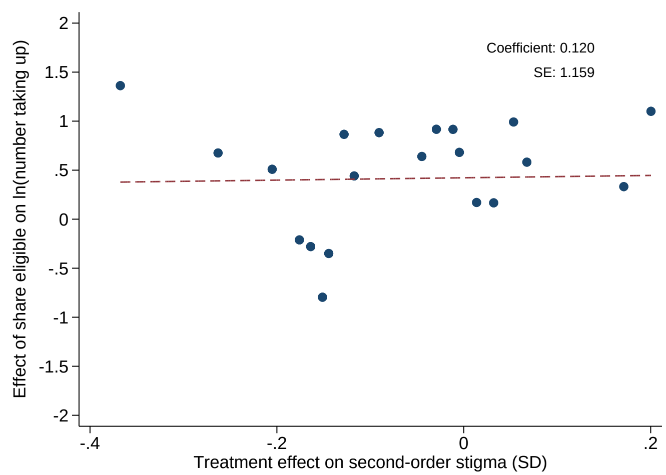

Panels A and B show the correlation between the subgroup-specific inframarginal effects with the share of respondents in the USDA FSPAS survey who reported (A) any stigma and (B) less than complete information. Panel C shows the correlation between the subgroup-specific inframarginal effects and the subgroup-specific treatment effect in the online experiment. Subgroups are defined by household head age bin, gender, and race/ethnicity (non-Hispanic white vs. other), as well as by their household composition and income decile in the national distribution. Estimates are weighted by the inverse of the product of the variances of the cell-level coefficients; see Appendix D for details.

Figure 7: Naïve Planner's Biased Risk Aversion

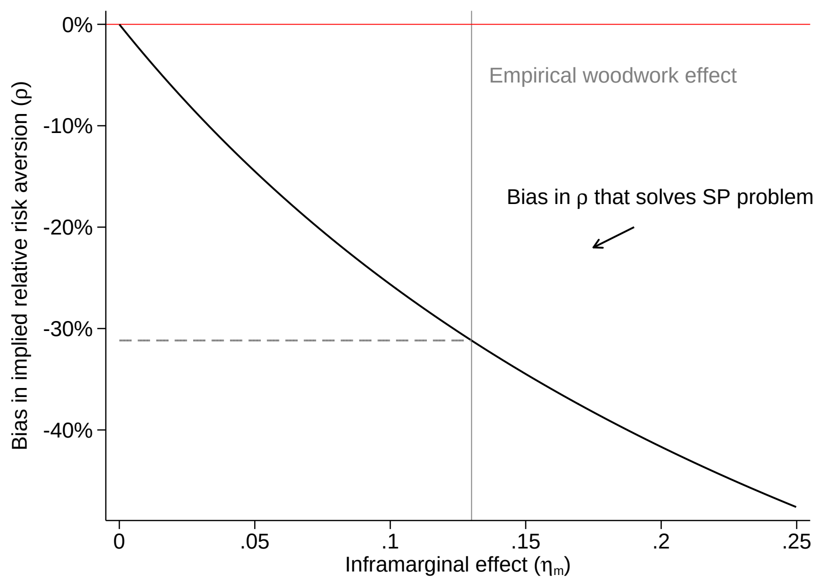

This figure shows the percent bias between the planner's "as-if" risk aversion ($\tilde{\rho}$) and the ground-truth risk aversion (ρ) (black line). Negative numbers indicate that the planner is behaving as if people are less risk averse than they really are. Panel A plots the bias as a function of the inframarginal effect; the vertical gray line plots the empirical inframarginal effect presented in Table 6. Panel B fixes η_m at the empirical inframarginal effect from Table 6 and varies s, the share of stigma agents.

Figure 8: Numerical Simulations: Optimal Eligibility Threshold and Take-Up

(A) Optimal Eligibility Threshold vs. Inframarginal Effects

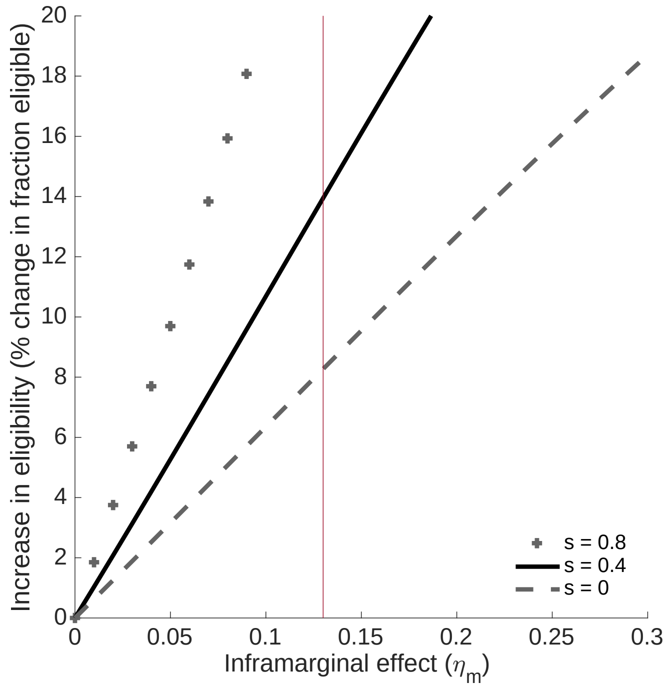

(B) Optimal Take-Up Rate vs. Inframarginal Effects

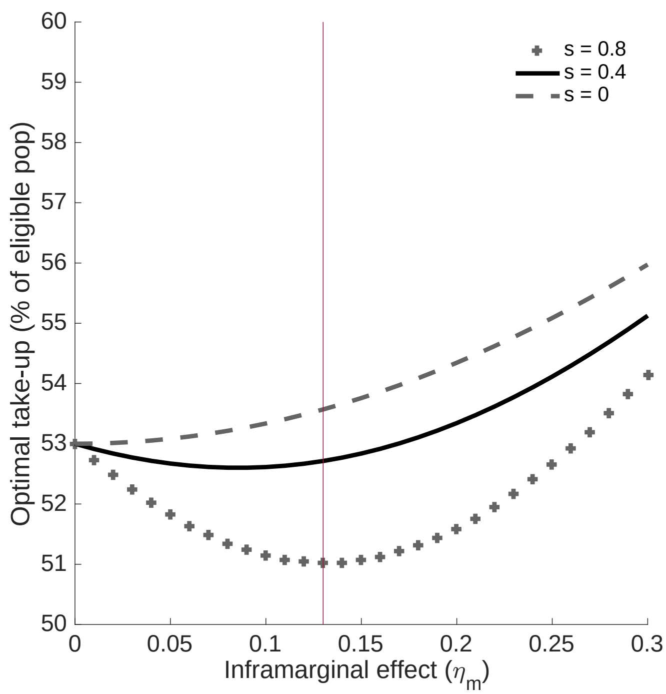

This figure shows the results from our numerical simulation exercise: it presents the change in the percent of people who are eligible relative to current policy if the planner were to acknowledge inframarginal effects (Panel A) and the optimal take-up rate (Panel B), as a function of the inframarginal effect η_m, using our preferred optimality condition (Equation (10)). Auxiliary parameters are set according to the values in Table 6.

8 Tables

Table 1: Estimates of the Inframarginal Effect

	(1) Main estimate	(2) Extra controls	(3) Waivers, lag unemp.	(4) Excludes recession	(5) Weighted	(6) Avg of coefficients	(7) All data
Panel A. 0–130% FPL							
Income limit (% FPL) / 100	0.085	0.087	0.074	0.086	0.082	0.076	0.091*
	(0.056)	(0.055)	(0.055)	(0.059)	(0.072)	(0.064)	(0.048)
Panel B. 50–115% FPL							
Income limit (% FPL) / 100	0.107**	0.112**	0.097*	0.114**	0.116*	0.103*	0.121**
	(0.051)	(0.054)	(0.050)	(0.053)	(0.064)	(0.056)	(0.047)
Observations	705	705	680	628	705	705	1071
N states	45	45	45	45	45	45	51

This table shows the effect of the eligibility threshold on log enrollment among the inframarginal population (0–130% FPL in Panel A and 50–115% FPL in Panel B). Column 1 estimates Equation (2), and the following columns present various extensions to show robustness. Column 2 separates the Ganong-Liebman policy index into separate indicators. Column 3 includes a control for the previous year's unemployment rate in each state and a control for the population-weighted average number of months a state had ABAWD work requirement waivers in effect. Column 4 excludes years 2008–2011, during the Great Recession. Column 5 weights observations by the state-year population. Column 6 presents the difference between the average pre- and post-period event study coefficients. Finally, Column 7 estimates Equation (2) using all the data available instead of only the event study sample. Standard errors are robust to heteroskedasticity and clustered by state.*, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table 2: Effects on Demographic Composition (50–115% FPL)

	(1) Female	(2) Black	(3) Age	(4) Has child	(5) Avg net income	(6) % FPL	(7) Certification $\leq$ 6 mo.
Income limit (% FPL) / 100	-0.001	0.059	0.391	-0.002	-28.557	0.732**	0.013
	(0.004)	(0.064)	(0.420)	(0.010)	(20.033)	(0.299)	(0.105)
Baseline mean	0.59	0.22	28.94	0.71	817.41	79.62	0.40
Observations	705	705	705	705	705	705	705
R^2	0.70	0.81	0.85	0.84	0.89	0.70	0.67

This table presents results from estimating the effect of the SNAP eligibility threshold on the composition of enrollees earning 50–115% FPL. The columns present estimates of Equation (2) with the indicated outcome variable. The independent variable is the eligibility threshold as a ratio of the Federal Poverty Level, so that increasing by 1 corresponds to increasing the eligibility threshold from, e.g., 130% FPL to 230% FPL. In each column, we use the specification described in Equation (2), where the outcome is indicated by the column header: Column 1 shows the effect of the eligibility threshold on the fraction of the 50–115% FPL enrollee sample who are female, and so on. "Baseline mean" refers to the average of the outcome indicated by the column in state-years where the eligibility threshold is 130% FPL. Outcomes are calculated using the USDA's Quality Control (QC) data, limiting the data to households earning 50-115% FPL. Standard errors are robust to heteroskedasticity and clustered by state. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table 3: Cost-Effectiveness Calculation

	Eligibility threshold	**Benefit size**
Change required for 1 pp take-up increase	3.9 pp	\$56 per person-year
1. Number of people affected	12 million (newly eligible)	44 million (enrolled)
2. Take-up among affected people	25%	100%
3. Cost per person-year	\$707	\$56
4. **Mechanical cost of intervention** (= Row 1 × Row 2 × Row 3)	\$2.2 billion	\$2.5 billion

This table shows the cost-effectiveness of increasing take-up by raising the means test versus by increasing the benefit size. Let m be the share of the U.S. population eligible for SNAP, B be the benefit size per person, and p be the take-up probability. The top row shows the required change in the instrument (m or B) to achieve a one percentage-point increase in take-up. We calculate this row by noting that $\eta_m = \frac{dp}{dm}\frac{m}{p}$ and rearranging to solve for dm when $dp = 0.01$ (and likewise for B). The remaining rows show the mechanical cost to the program (without including the costs incurred by inframarginal effects) of changing these instruments. When using the means test m to increase take-up, 12 million more people become eligible, but we estimate only 25% of those would take-up. When using the benefit size B, benefits are increased for all program participants. The cost per person uses averages from the QC data. The final row of the table shows the total mechanical cost for each policy tool, which multiplies rows 1-3.

Table 4: Estimates of the Take-up Elasticity with Respect to Eligibility Cutoff (η_m)

	OLS	IV		
		First Stage	Reduced Form	2SLS
Panel A. All data				
ln(Share eligible)	-0.105*			0.130*
	(0.060)			(0.067)
Income limit (% FPL) / 100		0.728***	0.094*	
		(0.034)	(0.048)	
Observations	1071	1071	1071	1071
Panel B. Event study sample				
ln(Share eligible)	-0.153**			0.104
	(0.069)			(0.077)
Income limit (% FPL) / 100		0.756***	0.079	
		(0.038)	(0.057)	
Observations	705	705	705	705

This table presents estimation results for η_m, the elasticity of take-up with respect to the share of the population who are eligible, controlling for the covariates included in Equation (3). We estimate this elasticity using the eligibility threshold as an instrument for the share of residents in a state who are eligible for SNAP. The first column shows results from a naïve OLS regression of ln(take-up) on ln(share eligible). The second column presents the first stage — the coefficient from a regression of ln(share eligible) on the eligibility threshold as a % of the Federal Poverty Level (FPL). The third column, the reduced form, gives the relationship between the eligibility threshold and ln(take-up). The final column gives the 2SLS estimate, our final estimate for η_m. Standard errors are robust to heteroskedasticity and clustered by state. ** and *** indicate $p < 0.05$ and 0.01, respectively.

Table 5: Decomposition: Stigma vs. Information

	(1)	(2)	(3)	(4)	(5)
1. Estimate of β_s	-1.211	6.686	0.580	-1.740	20.256
	(2.060)	(2.047)	(2.916)	(2.137)	(3.745)
2. Estimate of β_i	0.750	1.284	0.921	0.988	1.515
	(0.241)	(0.228)	(0.337)	(0.218)	(0.483)
3. Estimate of η_m^s	0.077	-0.434	-0.040	0.109	-1.316
	(0.151)	(0.150)	(0.213)	(0.156)	(0.274)
4. Estimate of η_m^i	0.593	0.952	0.677	0.780	1.038
	(0.090)	(0.085)	(0.126)	(0.082)	(0.181)
5. N cells	80	80	80	80	80
6. p-value for $H_0 : \eta_m^s = \eta_m^i$	0.006	< 0.001	0.008	< 0.001	< 0.001
Weights:	QC, Exp.	QC, FSPAS	Exp.	QC	FSPAS

This table shows the result of the formal decomposition exercise described in Section 5.2. Standard errors and p-values are formed from 99 bootstraps. Appendix D describes the estimation and weighting procedures.

Table 6: Summary of Parameters for Welfare Analysis

Parameter	Description	Primary Value	Range of Reasonable Values	Source
$\boldsymbol{\eta_m}$	**Take-up elasticity with respect to eligibility threshold (inframarginal effect)**	**0.13**	$\mathbf{[0.02, 0.24]}$	**Table 4 (and 90% CI)**
η_B	Take-up elasticity with respect to benefit size	0.5	$[0.3, 0.6]$	Krueger and Meyer (2002)
ρ	Coefficient of relative risk aversion	3	$[1, 4]$	Chetty and Finkelstein (2013)
s	Share of stigma-only types	0.4	$[0, 1]$	Food Stamp Program Access Study
γ	Cost-benefit ratio, conditional on take-up	0.5		Uniform costs assumption
$\frac{\partial \gamma}{\partial m}$	Change in cost-benefit ratio, conditional on take-up	0		Uniform costs assumption
m^*	Eligibility threshold (share eligible)	0.27		QC and CPS data
p^*	Take-up rate (all eligible)	0.53		QC and CPS data
$\lambda_m / \lambda_{\text{avg}}$	Ratio of marginal to inframarginal welfare weights	0.427		Inverse-optimum approach

This table summarizes the parameters used in the welfare analysis. We note the preferred value and source, but also show robustness to the range of values. The uniform costs assumption implies that γ and $\frac{\partial \gamma}{\partial m}$ are precisely 0.5 and 0, respectively.

Online Appendix: Not Intended for Publication

A Data and Institutional Context

A.1 SNAP Sample Construction

We build off the sample in Ganong and Liebman (2018), and adapt their public-use code and data associated with the published paper. We extend the sample to 2016. Our main outcome (the number of people enrolled in SNAP, for different income groups) uses the USDA's Quality Control (QC) data from 1996–2016. The QC data provides information on the household's income as a fraction of the FPL. We use the QC data (together with its household weights) to obtain counts of the number of people in a given state-year that enroll in SNAP who are within some income band (as a fraction of the FPL).

In our welfare exercise and in some supplemental analyses, we are interested in SNAP take-up *rates*. For these, we treat the QC data as the numerator in the take-up rate, and form the denominator from the CPS, which contains the count of people within a household income band in each state and year.

Our data on state-level SNAP policies, including the income eligibility threshold and other policies (e.g., outreach spending), come from the USDA's SNAP Policy Database (2019).

The QC data include individuals in the household who are not in the SNAP unit. As in Ganong and Liebman (2018), we include these individuals as taking up SNAP. Many of these individuals are relatives of the individuals in the SNAP unit and may, in practice, have their consumption subsidized by SNAP. Results are very similar if we limit only to individuals in the SNAP unit.

A.2 Broad Based Categorical Eligibility

We provide more information about the BBCE provision that permits states to expand SNAP eligibility.

Broad Based Categorical Eligibility permits states to expand eligibility using Temporary Assistance for Needy Families (TANF) or State Maintenance of Effort (MOE) budgets. States cannot expand eligibility beyond 200% of the FPL.

There are two concerns about other effects of the BBCE that could affect our analysis of inframarginal effects. In practice, states are legally required to fund small auxiliary services (e.g., telephone hotlines) using TANF/MOE funds in order to grant eligibility to more people in SNAP. Congressional Research Service (2019) writes:

> "As of July 2019, 42 jurisdictions have implemented what the U.S. Department of Agriculture (USDA) has called "broad-based" categorical eligibility. These jurisdictions generally make all households with incomes below a state-determined income threshold eligible for SNAP. States do this by providing households with a low-cost TANF-funded benefit or service such as a brochure or referral to a telephone hotline. There are varying income eligibility thresholds within states that convey "broad-based" categorical eligibility, though no state may have a gross income limit above 200% of the federal poverty guidelines."

The first concern, which we address in Section 2.6, is that this policy requires that SNAP administrators must notify households that they are eligible. In practice, the policy discussion around BBCE centers around the eligibility expansion, and the notification of receipt may not be much different than typical state efforts to notify recipients, especially for households below 115% of the FPL. The core of our robustness tests uses states that are treated with BBCE but do not expand eligibility. We find no evidence take-up increases in these states.

A secondary concern is that BBCE expansions sometimes waive asset rules. We also address this concern in Section 2.6.

A.3 Components of SNAP Policy Index

We use the SNAP policy index defined in Ganong and Liebman (2018), but without the BBCE. It is the average of dummies for each of seven policies. Six policies are directly from the SNAP Policy Database (2019). These are defined to be 1 if at least some parts of the state use the policy:

- At least one household vehicle is exempted from the asset test.

- Households with at least one recipient of Supplemental Security Income can use a simplified application for SNAP.
- Households can recertify with a telephone interview instead of a face-to-face interview.
- Households can apply to SNAP online.
- The state has fewer requirements for reporting changes in household earnings.
- There are call centers in the state for households to ask questions about SNAP, and in some places, recertify.

The final policy is a dummy if fewer than 20% of households have a certification period of 3 months or less, indicating that only a low share of SNAP households in the state must recertify at frequent intervals.

The index averages all seven policies except for when information about vehicle exemptions is unavailable; in this case, we average the remaining six.

In cases in which the index varies throughout the year, we use the minimum of the index in that year.

A.4 Experiment Sample Construction

We document several data cleaning decisions.

- A small number of participants had missing information about their household size or composition. We assume people with missing information were single, non-married, with no children (so had a household size of 1).
- A small number of participants had missing income. We assume they were in the bottom income bin and therefore had an income of $7,500.
- We top-coded household size at 6 because the most number of children that participants could report was 4.
- Incomes were top-coded at $250,000. We assume these participants had incomes of $300,000.
- Fewer than five participants took the experiment multiple times, and we drop them.
- **Attention checks.** The attention checks are the following. First, before treatment, we tell people: *"In this survey, we will ask you about your beliefs and attitudes about the Supplemental Nutrition Assistance Program (SNAP), also known as food stamps."* After eliciting the preferred charity (the incentive), we ask: *"What does SNAP stand for?"*. There are four multiple choice responses: *"Sufficiently Noisy Animal Parties"*; *"Supplementary Names Artful Program"*; *"Supplemental Nutrition Assistance Program"*; *"Salty Noodles And Pasta."* We drop the 106 participants who answer the acronym question incorrectly. Second, we drop the 145 participants who report that either 0 or 100% of people in the U.S. are eligible for SNAP.
- **Below 130% FPL Sample.** To form the "inframarginal" sample of experiment respondents, we predicted the relevant 2020 poverty threshold for each respondent using (1) the midpoint of their household income bin and (2) their household size, constructed via their marital status and number of kids. Anyone who reported a household income bin with a midpoint below 1.3 $\times$ the result is included in the sample of respondents under 130% FPL. This may have excluded some respondents from the inframarginal sample if they were also living with or supporting parents or elders.

A.5 Figure 1 Details

We collected income eligibility rules and take-up rates from various sources for a subset of U.S. social programs. To the extent possible, all values are from 2016. The set of programs was determined by the following process: We began by limiting to programs with FY 2016 budgets over $5 billion. We eliminated tax credits. Then we eliminated the following programs for specific reasons. We eliminated Section 8

Housing because the notion of participation is difficult to define where there are long wait lists and barriers to take-up are very high (often requiring moving). We eliminated Old Age Assistance and Social Security because income-based means tests are not meaningful for a population that often does not work and lives in households with other earners. Finally, we eliminated Pell Grants because eligibility is not based on a specific income threshold.

- CHIP
  - Eligibility data are from Brooks et al. (2016), Table 1, which gives income thresholds for children's eligibility to receive Medicaid or CHIP benefits, assuming a family of 3. In some states, the income threshold varies for different subgroups. The figure uses a population-weighted average of all the states' highest income thresholds.
  - The take-up rate is from Appendix Exhibit 1 of Haley et al. (2018), also as referenced by The Kaiser Family Foundation (KFF).
- EITC
  - Eligibility is calculated using the IRS.gov EITC maximum allowable AGI for a family of three.
  - The take-up rate is from the IRS.gov "About EITC" webpage (Internal Revenue Service, 2020), estimated by the Census Bureau using the CPS.
- Head Start
  - Eligibility is generally 100% of the FPL (HHS).
  - The take-up rate was calculated as follows:
    1. Participation rates are 35% (Child Trends, 2018), calculated using the total number of children enrolled in Head Start divided by the total number of children in poverty (ages 3-5).
    2. However, Head Start is oversubscribed. We use details from the Head Start Impact Study (U.S. Department of Health and Human Services, 2010): this study found that 85% of Head Start centers were oversubscribed. Within oversubscribed Head Start centers, the study randomized 60% of applicants into acceptance, while the remaining 40% were wait listed. In some centers, not all applicants were included in the randomization; in others, there were not enough applicants to attain this ratio in the randomization. We assume that take-up is 35*(1)*(15%) + 35*(10/6)*(85%). That is, the take-up rate is 35% among the 15% of centers which were not oversubscribed and 35*(10/6) in the oversubscribed centers, on average.
- Medicaid (parents only)
  - Eligibility data are from Brooks et al. (2016), Table 5, which gives income thresholds for parents' eligibility to receive Medicaid or CHIP benefits, assuming a family of 3. In some states, the income threshold varies for different subgroups. The figure uses a population-weighted average of all the states' highest income thresholds for parents.
  - The take-up rate is from Appendix Exhibit 2 of Haley et al. (2018), as referenced by KFF.
- NSLP (National School Lunch Program)
  - Eligibility for free lunch is 130% FPL in most districts; eligibility for reduced-price lunch is 185% FPL in most districts.
  - The take-up rate is calculated as follows:
    1. First, we take the total number of students eligible for free or reduced-price lunch in the 2015-2016 school year, according to Table 204.10 in National Center for Education Statistics (2017). This is around 26 million.
    2. We take the average number of free and reduced-price meals served daily in 2016, provided by the USDA Food and Nutrition Service: around 22 million (United States Department of Agriculture Food and Nutrition Service, 2020a).

3. The take-up rate is 22 / 26

- SNAP
  - Eligibility data use a population-weighted average of states' eligibility thresholds.
  - The take-up rate is from Cunnyngham (2019), which gives estimates of 2016 take-up rates.
- TANF (Temporary Assistance for Needy Families)
  - Eligibility data are from Giannarelli et al. (2017), which provides, for all states, the income cutoff in dollars for TANF initial eligibility for a family of three. These cutoffs were converted to percent of the 2016 Federal Poverty Level for a family of three. The final eligibility level is the population-weighted average of these.
  - The take-up rate estimate comes from Giannarelli (2019).
- WIC (The Special Supplemental Nutrition Program for Women, Infants, and Children)
  - Eligibility is capped at 185% of the FPL.
  - The take-up rate is an estimate from the USDA FNS (United States Department of Agriculture Food and Nutrition Service, 2020b).

A.6 Figure 2 Details

Using JSTOR and EBSCO, a research assistant collected all *AER* and *QJE* papers that met one of 33 search terms according to the search engine.[55] The search terms were: "welfare program," "social insurance," "social program," "social assistance," "social welfare," "social benefit," "income threshold," "participation threshold," "means-testing threshold," "means-tested program," "means-tested welfare," "means-tested benefit," "means-tested subsidy," "income means testing," "eligibility rule," "eligibility threshold," "eligibility criteria," "eligibility criterion," "eligibility requirement," "woodwork effect," "program eligibility," "program benefit," "program subsidy," "program duration," "optimal program," "optimal provision," "benefit schedule," "program schedule," "public insurance," "program take-up," "incomplete take-up," "welfare take-up," "benefit take-up."

We limit the sampling frame to the 2010–2018 *AER* and 2010–2019 *QJE*. Appendix A.6 provides the search terms. On the authors' websites, we also provide a spreadsheet of all the papers, their inclusion criteria, and how we classified them. We also provide a list of judgment calls involved in this exercise and our rationale for our decision. We exclude the papers and proceedings but include comments. We exclude the 2019 AER because it was not available on JSTOR or EBSCO. We then read the abstract and/or introduction of each of the 278 papers that met at least one of the 33 search terms. We determine whether a paper was about a social welfare program.

We impose the following additional criteria when categorizing papers.

- We exclude papers that are principally about optimal income or capital taxation.
- We exclude transfers that are not intended to alleviate poverty (e.g., the effects of giving people computers).
- We exclude papers about credit market restrictions only, such as papers about mortgage deductions. We do include papers about consumer bankruptcy.
- We exclude papers about search and matching in labor markets if they do not have a substantial social insurance angle (e.g., UI).

[55]The research assistant also searched the downloaded PDFs to see which search terms were most often met. Two of the papers that the search engines specified met the search terms did not actually include the search terms in the downloaded PDF, perhaps due to a bug in the search engine. Neither paper was deemed to be about social welfare programs so this issue does not substantively affect the conclusions.

- Because of the important theoretical connection between optimal social insurance and welfare design, we include papers that are about private insurance markets (including health insurance), as long as they have a significant angle about optimal policy.
- We define "program eligibility" as rules that determine whether a person has access to a social program. We do not consider eligibility to include access to different plan choices within a health program; our decision to exclude these papers is conservative, since they would only estimate a treatment effect using eligibility but not use optimal eligibility as an instrument.

B Empirics Appendix

B.1 Additional Figures

Figure B.1: BBCE implementation background

(A) Rollout of Eligibility Changes Per Year

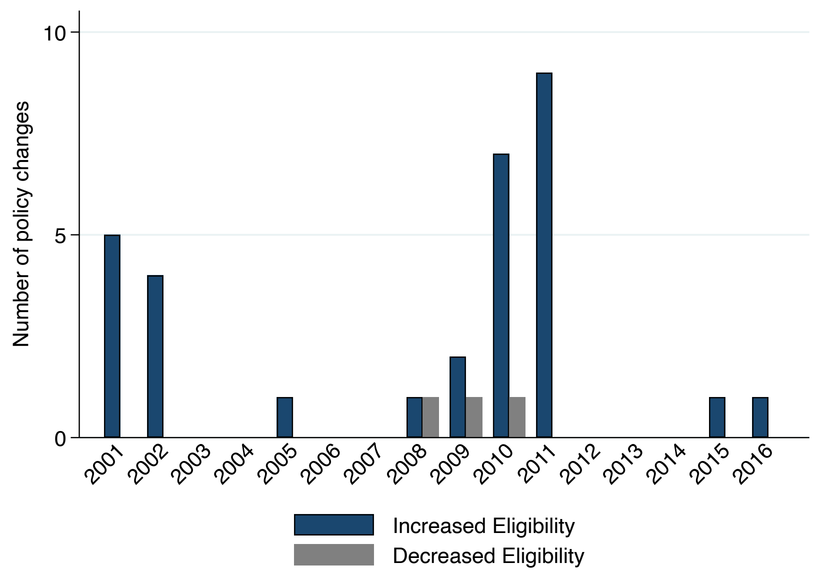

(B) Map of States that Implement Eligibility Expansions

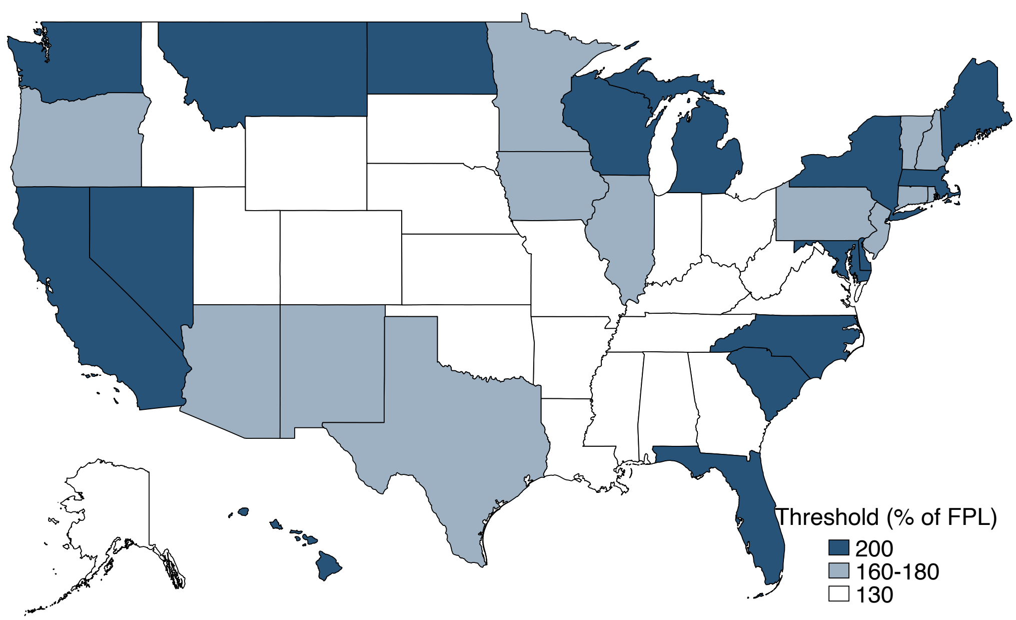

Panel A presents the number of states in each year that increased (blue bars) or decreased (gray bars) eligibility to the Supplemental Nutrition Assistance Program. Four states are counted twice, because they exhibit multiple changes. Panel B presents the maximum gross income eligibility threshold in a state from 1996–2016. The color coding refers to the maximum gross income eligibility threshold as a percent of the FPL; e.g., states colored in dark blue have maximum eligibility threshold of 200%. In two states that increase and then reduce the eligibility threshold, we present the largest eligibility threshold in the data. Source: SNAP Policy Database.

Figure B.2: Event Study of Changes to Eligibility Threshold: Without Controls

(A) Sample: 50 to 115% of FPL, No Controls

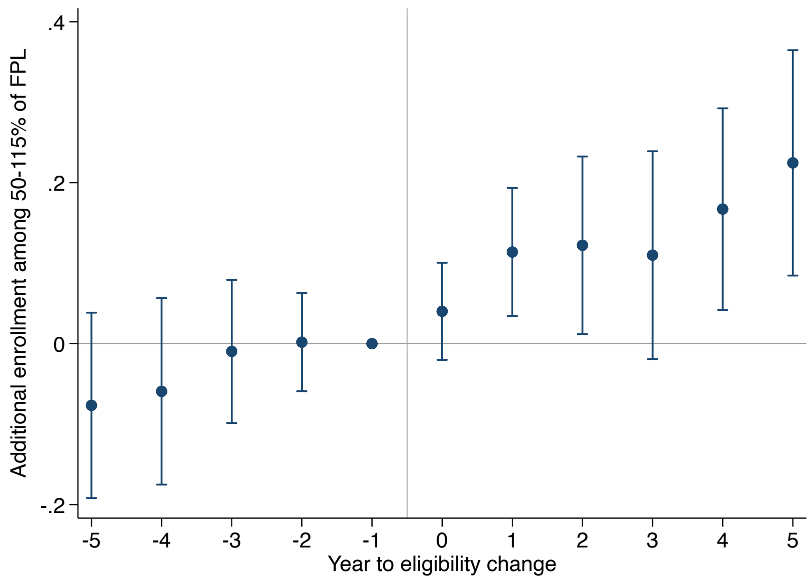

(B) Sample: 50 to 115% of FPL, Only Denominator Control

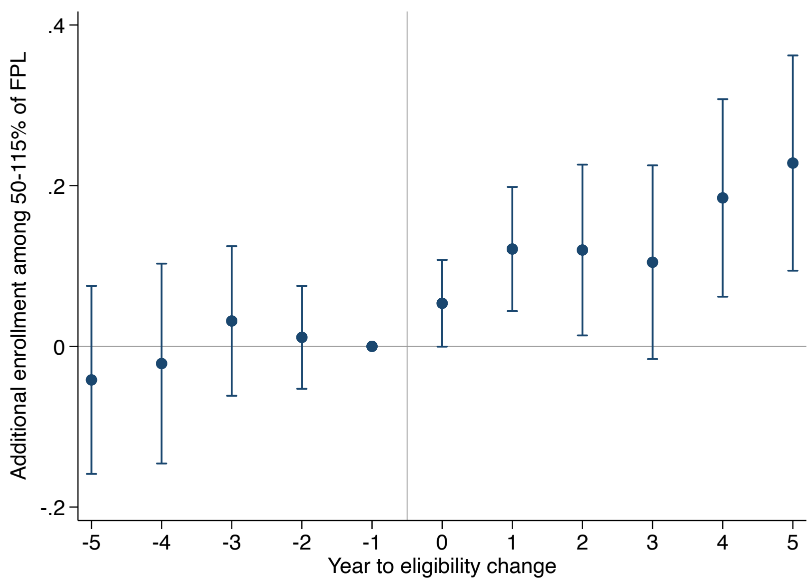

This figure is similar to Figure 4, but Panel A presents the specification with no controls beyond state and year fixed effects. Panel B presents the specification with state and year fixed effects, only controlling for the log of the total number of people between 50 and 115% of the FPL (from the CPS).

Figure B.3: Extra robustness checks

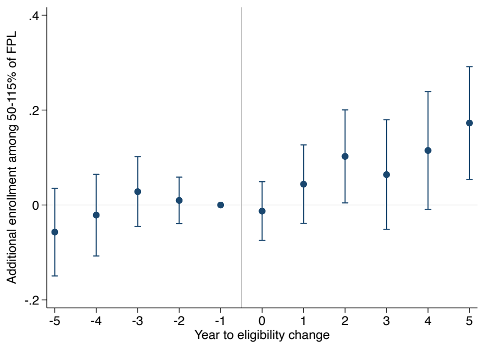

(A) Monthly data

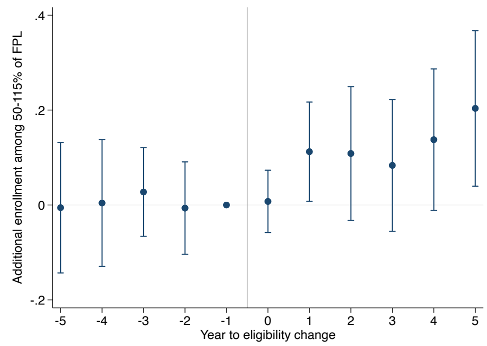

(B) Only Households with Dependents

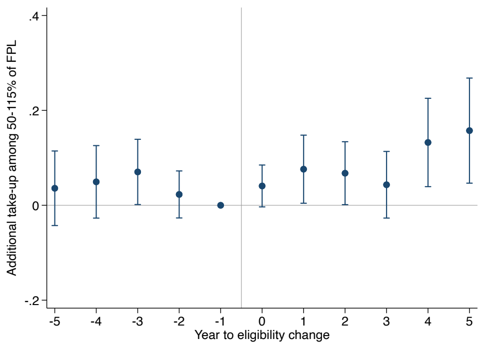

(C) Take-up

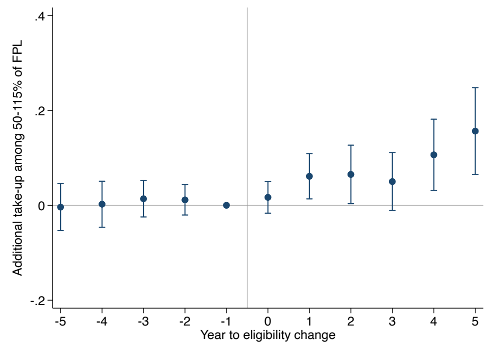

(D) Take-up, controlling for Ln(CPS count in 50-115%)

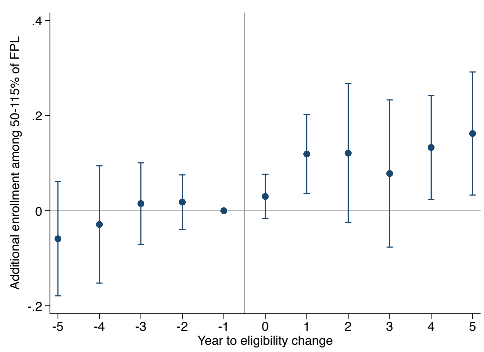

(E) Enrollment, weighted by state-year population

Panel A presents the results of estimating Equation (2) with monthly data instead of annual data. Panel B includes only SNAP recipients with any dependents—households that will not be affected by ABAWD work waivers. Panels C and D use the take-up share instead of the log of enrollment as the regressand, where the numerator in the take-up share comes from the USDA Quality Control data and the denominator uses the CPS. Panel C has no controls for state-year CPS population, while Panel D controls for the log of count of individuals in the CPS with household income in 50-115% FPL. Panel E uses the main specification and weights by population size in each state-year.

Figure B.4: Balance Tests

(A) SNAP Policy Index

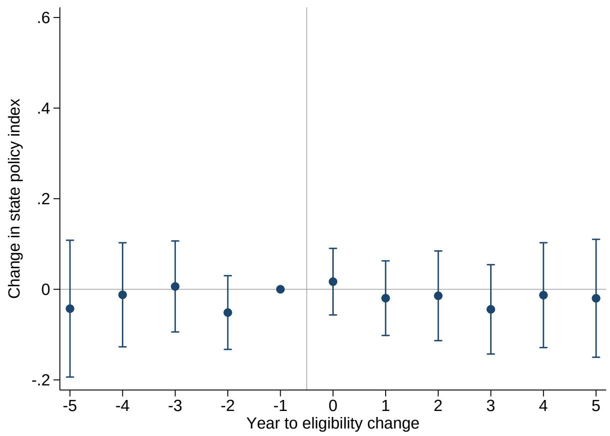

(B) CPS Population Count: 0 to 130% of FPL

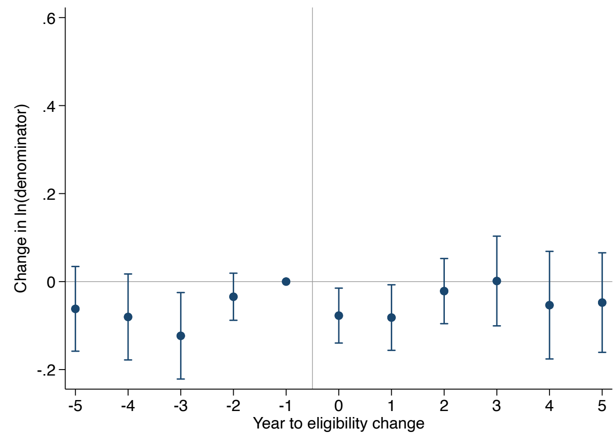

(C) Unemployment Rate

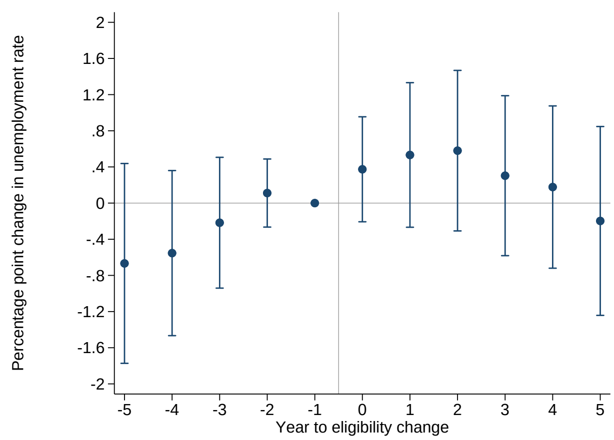

This figure presents placebo event studies with the main specification from Equation (1) but replacing the outcome with the main control variables. The event time is indexed around changes to state eligibility thresholds. Panel A uses the "Ganong-Liebman" index of SNAP policies, which are found in the USDA's SNAP Policy Database, as the outcome. Panel B uses the (ln of) the number of people in a state earning below 130% FPL (from the CPS) as the outcome. Panel C uses the state unemployment rate as the outcome. Standard errors are robust to heteroskedasticity and clustered by state.

Figure B.5: Effect on Take-Up Rates by Income Group

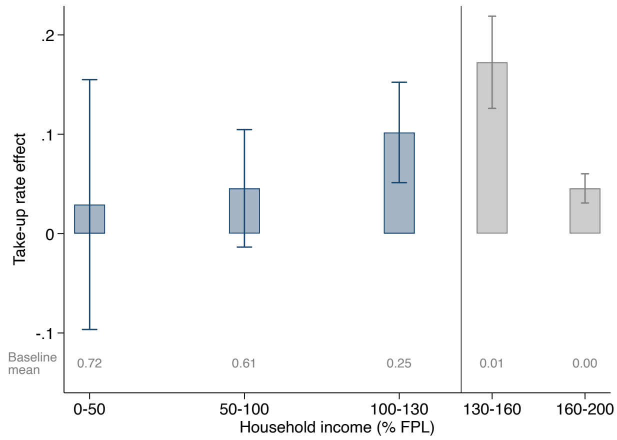

This figure presents estimates of Equation (2) using take-up rates as the outcome variable. The bars show the effect of the eligibility threshold on SNAP take-up by income group, and the whiskers show the 95% confidence intervals. While the regression specification is the same for all bars (with only the reference group changing), they are colored blue and gray to distinguish the effects on the inframarginal population versus the effects on the newly eligible population. Take-up rates are calculated using the enrollment counts from the USDA Quality Control (QC) data in the numerator and total counts of individuals within the income group from the Current Population Survey (CPS) in the denominator. Standard errors are robust to heteroskedasticity and clustered by state.

Figure B.6: Two-Way Fixed Effects Robustness

(A) Stacked Estimator (Cengiz et al., 2019)

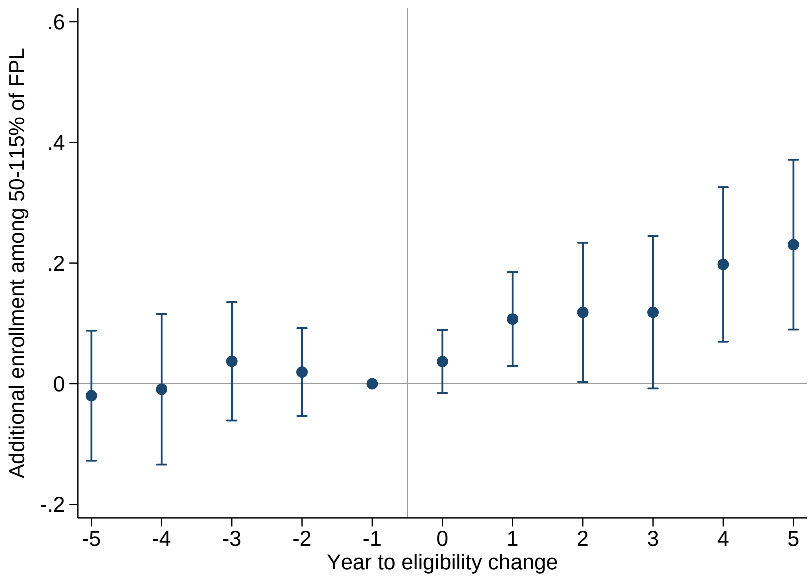

(B) Sun and Abraham (2021) Estimator

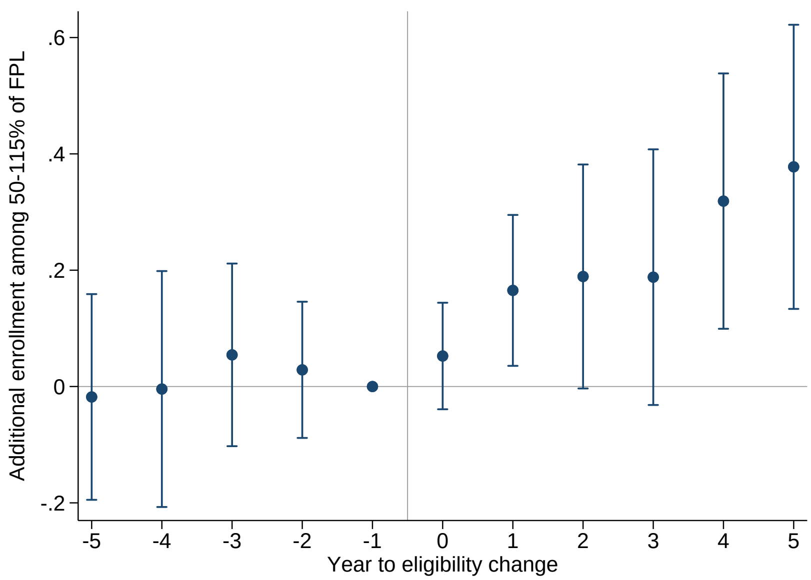

This figure presents heterogeneity-robust event study estimates using the 50–115% sample. Panel A presents the "stacked estimator" developed in Cengiz et al. (2019). For each treated state, we form a dataset keeping just one treated state and all never-treated states. We then stack all datasets and estimate a version of Equation 1, controlling for dataset-state fixed effects. We employ two-way clustering by dataset and state. Panel B presents the results from the estimator in Sun and Abraham (2021), using never-treated states as a comparison group.

Figure B.7: FSPAS Descriptives

(A) Stigma

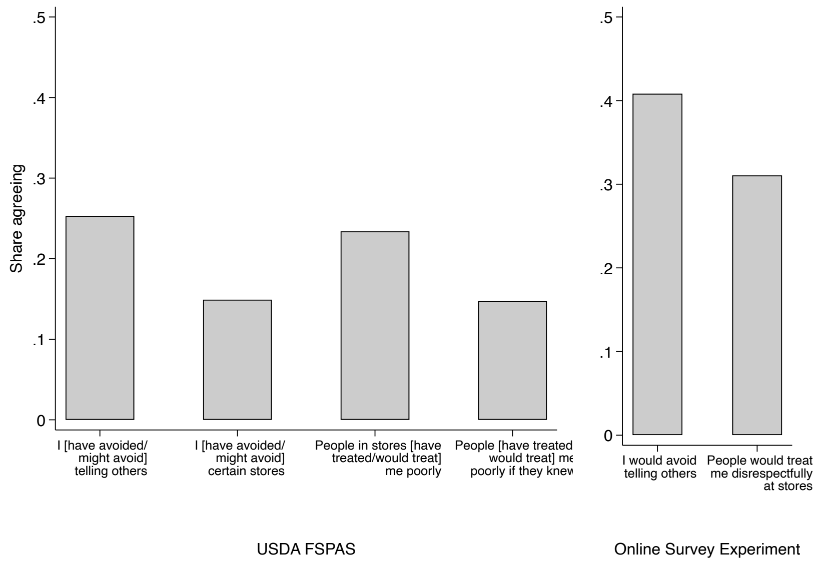

(B) Information

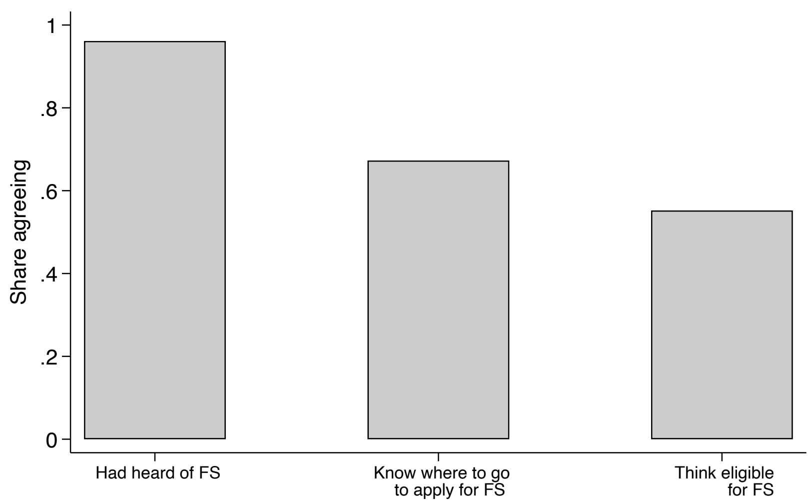

This figure shows the share of respondents (among approved applicants and eligible nonparticipants) agreeing with different statements presented in the USDA Food Stamp Program Access Study about the stigma around SNAP (in Panel A) and their access to information about SNAP (Panel B). In Panel A, we compare results to those from the online experiment, limited to respondents earning under 130% FPL.

B.2 Additional Tables

Table B.1: Estimates of the Inframarginal Effect in Alternate Samples

	(1) Main estimate	(2) Extra controls	(3) Waivers, lag unemp.	(4) Excludes recession	(5) Weighted	(6) Avg of coefficients	(7) All data
Panel A. 0–115% FPL							
Income limit (% FPL) / 100	0.064	0.067	0.052	0.064	0.063	0.054	0.072
	(0.057)	(0.056)	(0.055)	(0.060)	(0.073)	(0.065)	(0.048)
Panel B. Any dependents							
Income limit (% FPL) / 100	0.105*	0.112*	0.096*	0.114*	0.123*	0.104*	0.133***
	(0.057)	(0.060)	(0.056)	(0.059)	(0.071)	(0.062)	(0.048)
Observations	705	705	680	628	705	705	1071
N states	45	45	45	45	45	45	51

This table presents Table 1 with different samples, using the specification in Equation (2). See notes to Table 1 for details. Panel A uses the sample of people at 0–115% of the Federal Poverty Line (FPL). Panel B presents estimates for the sample of households with dependents, who are not subject to ABAWDs rules, in households earning 50–115% FPL. The outcome is SNAP enrollment as estimated from the USDA Quality Control data. Standard errors are robust to heteroskedasticity and clustered by state. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table B.2: Pre-Policy State Characteristics

Panel A. By Ever Changed Threshold	No	Yes	p-value
Share of state pop enrolled	0.08	0.07	0.29
Unemployment rate	3.93	3.88	0.85
Average family income in state	51.87	57.42	0.01
Ganong-Liebman Index	0.06	0.10	0.09
Outreach spending	1.64	16.92	0.09
Observations	30	21	
Panel B. By New Eligibility Threshold	$<$ 200% FPL	200% FPL	p-value
Share of state pop enrolled	0.09	0.09	0.79
Unemployment rate	6.04	5.89	0.86
Average family income in state	73.57	64.79	0.08
Ganong-Liebman Index	0.52	0.43	0.48
Outreach spending	54.57	48.82	0.85
Observations	13	17	

In Panel A, we compare states which did and did not ever change their SNAP eligibility threshold in their pre-policy characteristics, as measured in the year 2000 (before any states implemented policy changes). In Panel B, we limit the sample to states which did increase their eligibility threshold and compare those which raised it to 200% FPL to those which raised it to a value below 200% FPL, where the pre-policy characteristics are measured two years before their policy change. The first row shows the share of the state population enrolled in SNAP in the given year. The second row shows the state unemployment rate. The third row shows the average family income (from the CPS). The fourth shows the Ganong-Liebman Index, excluding the BBCE indicator. The final row shows spending on SNAP outreach in the state, where the value is winsorized.

Table B.3: USDA FSPAS Characteristics

	(1) Info types	(2) Stigma types
Enrolled	0.38	0.43
Female	0.76	0.78
White	0.52	0.65
Has kids	0.47	0.53
Age	43.77	39.31
Observations	953	575

The table shows summary statistics for respondents categorized as "information-only" and "stigma-only" in the USDA Food Stamp Program Access Study (approved applicants and eligible nonparticipant samples only).

B.3 Measurement Error Robustness

We study whether measurement error in reported income in the Current Population Survey (CPS) could explain our main results. Figure B.4B shows that the count of people in the CPS earning below 130% FPL does not change discretely around the time of the policy implementation. The figure for people earning 50–115% FPL looks very similar. Especially given that we control for the denominator, it is implausible that state populations grow fast enough only in treated state-years, beginning exactly at the time of the eligibility increase, that this measurement error could explain our event study results. Any threat to identification requires that the mismeasured portion of the denominator grows in a way that is correlated with treatment, beginning precisely at the date of treatment.

To formalize this point, we obtain the following bound on the magnitude of measurement error in the denominator required to explain our results. In state-years with an eligibility threshold above 130% of the FPL, we simulate systematic measurement error in the denominator using an "inflated" denominator that we define as:

$$\text{simulated denominator} := \text{observed denominator} \times \text{inflation factor},$$

where the inflation factor represents the magnitude of simulated measurement error. For instance, an inflation factor of 1.05 represents the case where we replace the treated state-years' denominators as being 5% larger than what we observe in the CPS.

We then estimate Equation (2) with the simulated denominator in treated state-years. We find that the inframarginal effect vanishes only if the denominator in treated state-years is inflated by more than 30% (Figure B.8). Put another way, only when we add an additional 30% of the population to the denominator (and impose that this measurement error only exists in treated state-years) can we eliminate the inframarginal effect. As a benchmark, we note that the average state population between 50 to 115% of the FPL (i.e., the denominator) grew by 26% between 2001 and 2016. Thus the measurement error required to reverse our result would need to be larger than the entire observed population growth in the sample period. It is implausible that *only* treated states are subject to measurement error that is this extreme.

Altogether, while our denominator obtained from the CPS may be subject to some measurement error, it would have to be systematically correlated with treatment to an implausible degree in order to explain our results.

Figure B.8: Simulated Measurement Error

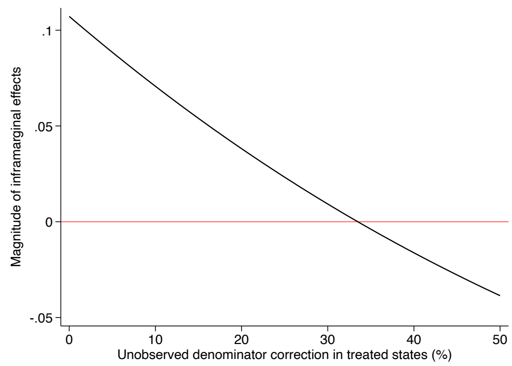

This figure presents a bound on the amount of measurement error in the denominator that would be required to reverse our results. In states where the eligibility threshold exceeds 130% of the FPL, we inflate the observed population between 50–115% of the FPL by the factor on the x-axis. We then present the estimate of the inframarginal effect from Equation (2), estimated using the simulated denominator. Only if the population is inflated by 30% can we reverse the inframarginal effect.

C Online Experiment Appendix

C.1 Auxiliary Experiment

Table C.2 shows the auxiliary experiment is balanced between treatment and control. The results of this second experiment are mixed (Table C.6). We find no evidence for an effect of a belief correction exercise on first-order beliefs. We find a *positive* effect of the belief correction on second-order beliefs: for people whose priors were below the truth, correcting beliefs *raises* the stigma they report (point estimate: 0.069, SE: 0.041, $p = 0.091$).

We note that the treatment effect is positive for people whose beliefs are corrected down (point estimate: 0.018, SE: 0.035). This point estimate is consistent with the results from the high-state treatment. Alternatively, it may suggest that any belief correction may simply cause participants to report more stigma, e.g. because they do not like being corrected after receiving an initial hint. We also present effects with demographic controls (Table C.7), which are similar. In this case, the positive effect on second-order stigma for correcting beliefs upward is very slightly attenuated.

We are more cautious about interpreting the results from auxiliary experiment for the following reasons. First, people who are shown multiple pieces of information might simply end up confused, which could attenuate or undo its effects. Because we did not elicit beliefs after being shown the belief correction, we do not have a way of checking how the correction actually shifted posteriors. The inconclusive results suggest that providing the second piece of information might have had an unintended consequence of causing participants to tune out the second piece of information, perhaps because it was perceived as contradicting the first piece of information.

Second, the auxiliary belief correction only operates on people *after* they have been shown a hint. As a result, because it is cross-randomized, it affects the group of people that do or do not comply with the high or low treatment. The staggered nature of the design complicates this interpretation: people who have low prior beliefs after treatment are a selected group, since they have been exposed to a hint that causes them to update.[56]

Third, the belief-correction treatment, when paired with the high-share treatment, affects people's beliefs about the *distribution* of eligibility thresholds across states. If stigma is linked to people's beliefs about the distribution of eligibility thresholds, it is not clear how the combination of experiments affects stigma.

Altogether, the main experiment provides a somewhat cleaner test of the null hypothesis that stigma plays no role in inframarginal effects. Nevertheless, the inconclusive results from the auxiliary experiment lead us to interpret the experiment with some caution.

[56]Consistent with this point, the positive treatment effect on second-order stigma from correcting beliefs upward attenuates once we add demographic controls (Table C.7).

C.2 Additional Figures

Figure C.1: Visual Depiction of Experiment Design

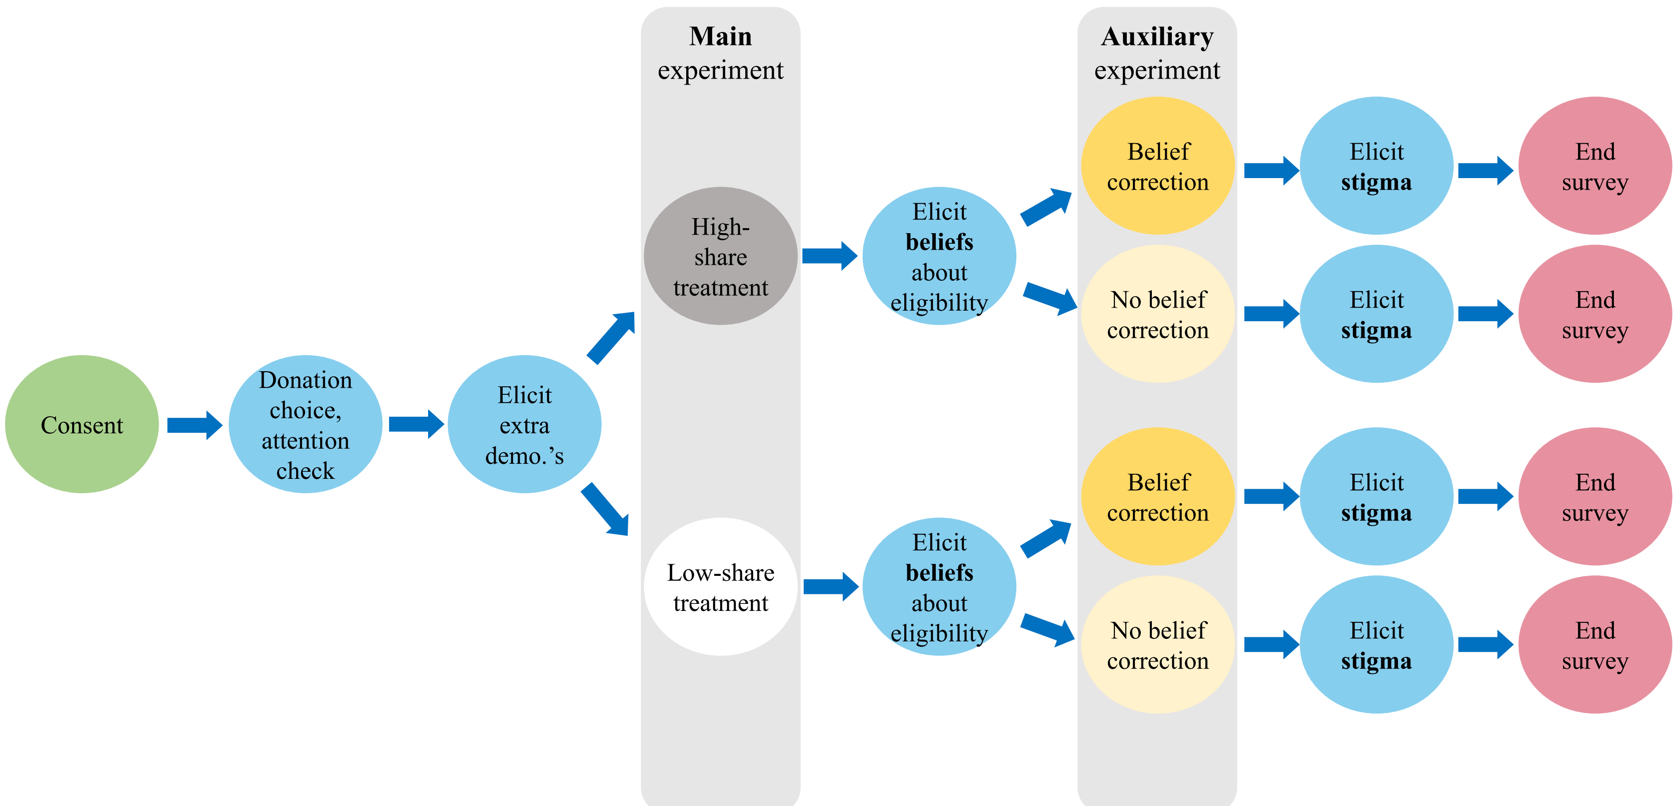

This figure presents the experiment design. The donation choice was to one of four charities (used to incentivize belief elicitation). We elicited several demographics (in addition to those provided by Lucid).

Figure C.2: Effect of High-Share Treatment on Beliefs about Eligibility

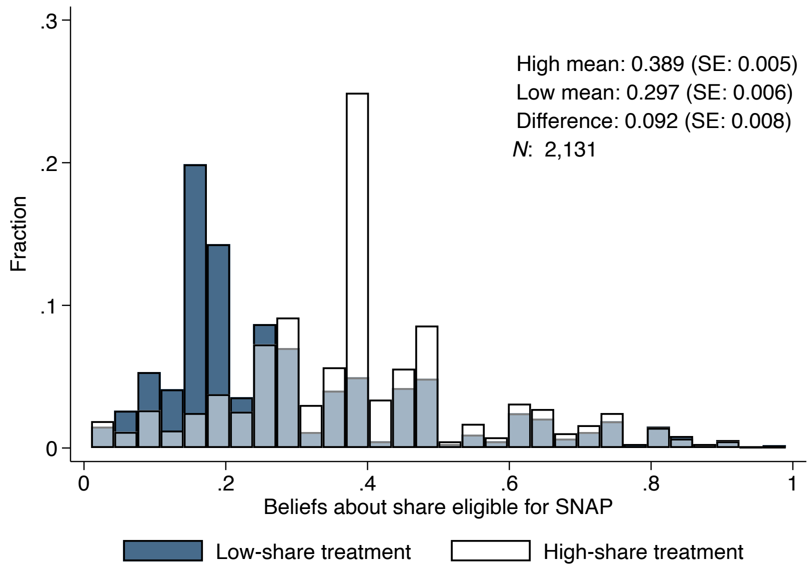

This figure presents the distribution of beliefs from the online experiment, split by treatment group, about the fraction of people who are eligible for SNAP. The y-axis shows the share of people within each treatment group who report a given fraction are eligible for SNAP. The blue bars show the values for the low-share treatment. The white bars show the values for the high-share treatment. The light blue shaded area shows the overlap.

Figure C.3: Correlations Between Stigma Questions

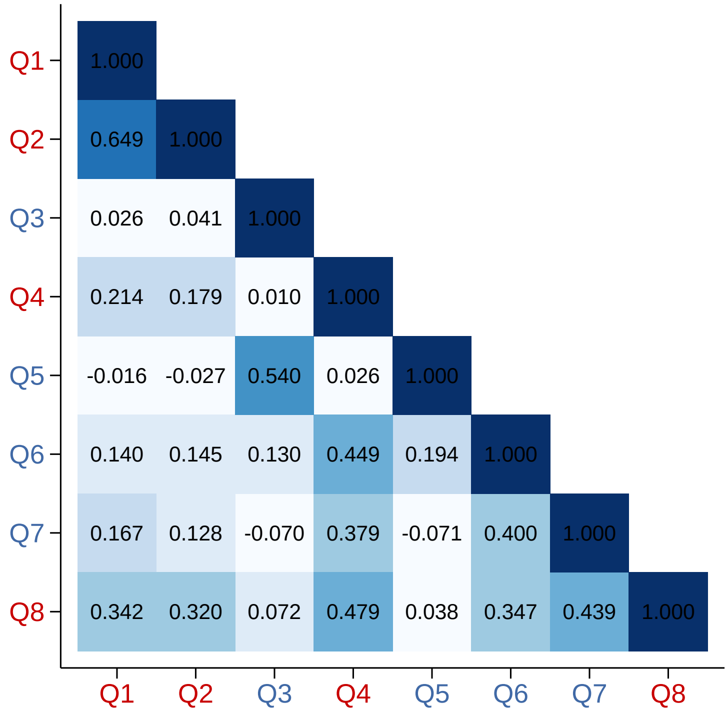

This figure presents correlations between the stigma questions in the order they were elicited. Section 3 provides the question texts. We classify questions 1, 2, 4, and 8 (labeled in red) as first-order stigma. We classify questions 3, 5, 6, and 7 (labeled in blue) as second-order stigma.

Figure C.4: Treatment Effect Heterogeneity

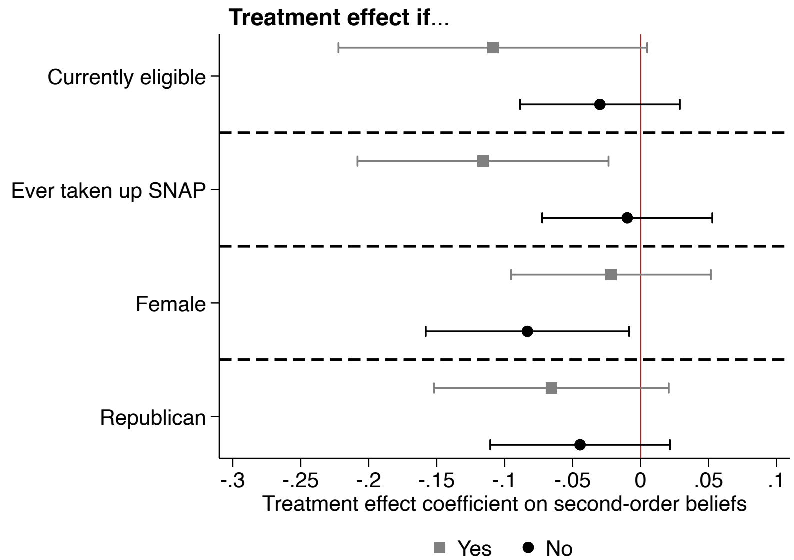

This figure presents treatment effects and 95% confidence intervals of the high-share treatment on the second-order stigma index (Equation (4)), split by demographic group.

C.3 Additional Tables

Table C.1: Experiment Sample Composition and Balance for High vs. Low Treatment

	CPS Sample	Full Sample			Below 130% FPL		
		Low-share	High-share	p-value	Low-share	High-share	p-value
Female	0.517	0.531	0.522	0.658	0.647	0.597	0.248
White	0.776	0.727	0.737	0.623	0.707	0.684	0.583
Hispanic	0.165	0.109	0.112	0.824	0.120	0.160	0.203
At least some college	0.611	0.778	0.772	0.737	0.606	0.612	0.894
Age	47.714	45.679	46.145	0.526	45.036	45.042	0.997
Any Children	0.254	0.537	0.531	0.790	0.618	0.597	0.619
Single	0.291	0.366	0.368	0.927	0.418	0.441	0.594
Household Size	2.296	2.519	2.517	0.973	2.687	2.692	0.970
Democrat	-	0.541	0.517	0.275	0.522	0.490	0.476
On Food Stamps (Currently or Ever)	-	0.383	0.392	0.648	0.627	0.624	0.946
Household Income (000's)	-	59.007	59.941	0.680	15.331	13.431	0.021
Census regions							
Northeast	0.175	0.208	0.191	0.308	0.169	0.095	0.014
Midwest	0.207	0.190	0.198	0.617	0.189	0.209	0.565
South	0.379	0.344	0.372	0.176	0.369	0.441	0.100
West	0.238	0.259	0.240	0.311	0.273	0.255	0.639
Joint F-test p-value				0.941			0.018
Observations				2131			512

Income uses the midpoint of a set of bins and is top-coded at $250,000. Household size is top-coded at 6. The CPS sample uses the 2019 NBER MORGs.

Table C.2: Online Experiment: Randomization Balance for Belief Correction

	CPS Sample	Full Sample			Below 130% FPL		
		No Correction	Belief Correction	p-value	No Correction	Belief Correction	p-value
Female	0.517	0.515	0.537	0.304	0.615	0.626	0.798
White	0.776	0.742	0.722	0.312	0.704	0.687	0.665
Hispanic	0.165	0.109	0.113	0.753	0.130	0.151	0.488
At least some college	0.611	0.783	0.767	0.378	0.615	0.604	0.788
Age	47.714	45.770	46.048	0.705	45.725	44.400	0.391
Any Children	0.254	0.528	0.540	0.559	0.623	0.592	0.473
Single	0.291	0.366	0.368	0.906	0.401	0.457	0.203
Household Size	2.296	2.530	2.507	0.724	2.757	2.626	0.350
Democrat	-	0.525	0.533	0.709	0.490	0.521	0.486
On Food Stamps (Currently or Ever)	-	0.377	0.398	0.329	0.615	0.634	0.665
Household Income (000's)	-	61.526	57.476	0.073	14.787	13.952	0.311
Census regions							
Northeast	0.175	0.199	0.200	0.965	0.162	0.102	0.044
Midwest	0.207	0.208	0.180	0.112	0.215	0.185	0.402
South	0.379	0.361	0.354	0.749	0.413	0.400	0.766
West	0.238	0.232	0.265	0.077	0.211	0.313	0.008
Joint F-test p-value				0.611			0.498
Observations				2131			512

Income uses the midpoint of a set of bins and is top-coded at $250,000. Household size is top-coded at 6. The CPS sample uses the 2019 NBER MORGs.

Table C.3: Online Experiment: Attrition Balance

	Total N	High-share treatment		Beliefs correction	
		All	<= 130% FPL	All	<= 130% FPL
1. Any attrition or drops	**567**	**0.009**	**0.018**	$-$**0.001**	**0.005**
		(**0.016**)	(**0.033**)	(**0.015**)	(**0.033**)
2. Bad priors	237	0.002	0.008	-0.000	0.003
		(0.011)	(0.024)	(0.011)	(0.025)
3. Attrited before share treatment	49	0.004	-0.004		
		(0.005)	(0.010)		
4. Attrited at or after treatment	126	0.006	0.007	-0.001	0.013
		(0.008)	(0.019)	(0.008)	(0.019)
5. Omitted any stigma answers	107	0.002	0.005	-0.006	-0.013
		(0.008)	(0.016)	(0.008)	(0.017)
6. Inattentive	106	0.000	0.021	0.005	0.003
		(0.008)	(0.013)	(0.008)	(0.014)
Observations		2,698	689	2,698	689

This table shows that attrition and drops were balanced across treatment and control. Each row tests for balance between treatment and control on a different dummy outcome. The first column gives the total number of respondents who were dropped for the reason indicated by the row. Note that respondents could be dropped for multiple reasons. The next two columns show balance for the main experiment, where respondents were provided a random hint about the share of Americans eligible for SNAP. The last two columns show balance for the secondary experiment, where respondents beliefs were corrected with the true share. Row 1's outcome is a dummy for attriting or being dropped from the sample. Row 2's outcome is a dummy for providing prior beliefs about the share of Americans eligible for SNAP that were below 1% or above 99%, or skipping this question entirely. Row 3's outcome is a dummy for dropping out of the survey before the treatment screen. The second two columns of this row are empty because individuals who attrited before the treatment screen were not randomized into treatment or control for the beliefs correction. Row 4's outcome is a dummy for attriting at or after the share treatment screen. Row 5's outcome is a dummy for not answering any of the stigma questions. Row 6's outcome is a dummy for failing an attention check. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table C.4: Online Experiment: High-Share Effect on Reported Stigma, without Demographic Controls

	Overall	Subindices	
		First-Order	Second-Order
Under 130% FPL			
High-share treatment	-0.032	0.046	-0.109*
	(0.050)	(0.065)	(0.058)
p-value	0.530	0.485	0.061
Observations	512	512	512
Full Sample			
High-share treatment	-0.013	0.025	-0.050*
	(0.024)	(0.031)	(0.027)
p-value	0.598	0.421	0.061
Observations	2,131	2,131	2,131

The table shows the effect of the "high-share" hint on individuals' level of agreement to statements measuring stigma around food stamps and welfare for individuals under 130% FPL (top panel) and the full sample (bottom panel) (Equation (4)). The estimates are identical to Figure 5. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table C.5: Online Experiment: High-Share Effect on Reported Stigma, with Demographic Controls

	Overall	Subindices	
		First-Order	Second-Order
Under 130% FPL			
High-share treatment	-0.023	0.049	-0.096*
	(0.050)	(0.064)	(0.058)
p-value	0.640	0.448	0.099
Observations	512	512	512
Full Sample			
High-share treatment	-0.016	0.016	-0.048*
	(0.023)	(0.029)	(0.026)
p-value	0.489	0.580	0.072
Observations	2,131	2,131	2,131

The table shows the effect of the "high-share" hint on individuals' level of agreement to statements measuring stigma around food stamps and welfare for individuals under 130% FPL (top panel) and the full sample (bottom panel) (Equation (4)). It is identical to Table C.4 and Figure 5 except we include demographic controls for: an age quadratic, income, political party, gender, region, household size, marital status, having children, being on or ever having been on food stamps, and education and race/ethnicity fixed effects. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table C.6: Online Experiment: Belief Correction, No Demographic Controls

	Overall	Subindices	
		First-Order	Second-Order
Panel A. Priors < Truth			
Beliefs Correction Treatment	0.044	0.020	0.069*
	(0.036)	(0.048)	(0.041)
Observations	868	868	868
p-value	0.218	0.680	0.091
Panel B. Priors $\geq$ Truth			
Beliefs Correction Treatment	0.008	-0.002	0.018
	(0.031)	(0.041)	(0.035)
Observations	1,263	1,263	1,263
p-value	0.800	0.964	0.615

This table shows results from the second experiment embedded in our online survey, where respondents were informed of the true share of Americans eligible for SNAP after previously being asked to report their beliefs (and given a hint, which is the primary experiment discussed in the text). It presents treatment effect estimates from Equation (4). Panel A restricts the sample to those who initially underestimated the eligibility share, so that the treatment should have led them to revise upwards. Panel B restricts the sample to those who initially overestimated the eligibility share, so that the treatment should have decreased their beliefs. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table C.7: Online Experiment: Belief Correction, With Demographic Controls

	Overall	Subindices	
		First-Order	Second-Order
Panel A. Priors < Truth			
Beliefs Correction Treatment	0.036	0.006	0.066
	(0.035)	(0.045)	(0.040)
Observations	868	868	868
p-value	0.301	0.900	0.103
Panel B. Priors $\geq$ Truth			
Beliefs Correction Treatment	0.032	0.034	0.030
	(0.030)	(0.038)	(0.035)
Observations	1,263	1,263	1,263
p-value	0.290	0.375	0.389

This table shows results from the second experiment embedded in our online survey, where respondents were informed of the true share of Americans eligible for SNAP after previously being asked to report their beliefs (and given a hint, which is the primary experiment discussed in the text). It presents treatment effect estimates from Equation (4). Panel A restricts the sample to those who initially underestimated the eligibility share, so that the treatment should have led them to revise upwards. Panel B restricts the sample to those who initially overestimated the eligibility share, so that the treatment should have decreased their beliefs. This table is identical to Table C.6 except we additionally include demographic controls for: an age quadratic, income, political party, gender, region, household size, marital status, having children, being on or ever having been on food stamps, and education and race/ethnicity fixed effects. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table C.8: Online Experiment: Treatment Effect by Belief-Correction Randomization

	(1) First-order index	(2) Second-order index	(3) First-order index	(4) Second-order index
High-share treatment	0.048	-0.019	0.003	-0.081**
	(0.044)	(0.037)	(0.044)	(0.038)
Observations	1050	1050	1081	1081
Sample	Not shown truth	Not shown truth	Shown truth	Shown truth

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

This table presents treatment effects on first- and second-order stigma from Equation (4) the sample by whether the sample's beliefs were not truthfully corrected (Columns 1 and 2) or were truthfully corrected (Columns 3 and 4). *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

Table C.9: Online Experiment: Association Between Take-Up and Stigma

	(1) On SNAP (currently or ever)	(2) On SNAP (currently or ever)	(3) On SNAP (currently or ever)
First-order index	-0.133***		-0.145***
	(0.014)		(0.015)
Second-order index		-0.012	0.044**
		(0.018)	(0.018)
Constant	0.391***	0.388***	0.391***
	(0.010)	(0.011)	(0.010)
Observations	2131	2131	2131

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

This table presents associations between first- and second-order stigma and participants' reports about taking up SNAP (now or in the past). We elicit the take-up questions before treatment. *, **, and *** indicate $p < 0.1$, 0.05, and 0.01, respectively.

D Mechanisms Appendix

This appendix provides information about the measurement and estimation in Section 3.2 and 5.2.

D.1 FSPAS Data

- We use surveys of (1) eligible nonparticipants and (2) successful SNAP applicants from the FSPAS, a study conducted by the USDA in the year 2000. The USDA considered someone an eligible nonparticipant if their household income was beneath 130% FPL and they were not currently enrolled in SNAP. There are 421 successful SNAP applicants and 1,323 eligible nonparticipants.
- **Stigma.** Respondents were considered affected by stigma if they answered "yes" to (agreed with) at least one of the following questions (statements).
  - If they'd ever been enrolled in SNAP:
    * Have you ever avoided telling people you got food stamps?
    * Did you ever go out of your way to shop at a store where no one knew you?
    * Have you ever been treated disrespectfully when using food stamps in a store?
    * Were you ever treated disrespectfully when you told people that you received food stamps?
  - If they'd never been enrolled in SNAP:
    * "If I got food stamps, I might go out of my way so people would not find out."
    * "I might not shop in certain stores because I don't want people there to know I use food stamps."
    * "People in stores would treat me disrespectfully when I use food stamps."
    * "People would treat me disrespectfully if they found out that I got food stamps."
- **Information.** Respondents were considered affected by information barriers if they (a) were in the eligible nonparticipant sample and (b) said "no" to any of the following questions:
  - Had you heard of food stamps or the Food Stamp Program before today's interview?
  - Do you know where you would have to go to apply for food stamps or other assistance?
  - Do you think you may be eligible to receive food stamp benefits?
- **Survey weights.** Each survey in the FSPAS is weighted to be representative of the population the respondents were sampled from. When we combine participants and eligible nonparticipants, we adjust these weights according to the share of Americans who participated in SNAP conditional on being eligible in the year 2000 (estimated in the QC data to be 40%).

D.2 Estimation procedures

We seek to estimate the coefficients β^s and β^i from the equation:

$$\left(\frac{\partial \ln(\text{N enrolled})}{\partial m}\right)_d = \frac{1}{p}\left(\left(\frac{\partial c}{\partial m}\right)_d s_d\right) \times \beta^s + \frac{1}{p}(1-s)_d \times \beta^i + \epsilon_d, \tag{D.1}$$

noting that the $\frac{\partial \ln \text{Number Eligible}}{\partial m}_d$ term in Equation (13) vanishes because our objective is to study inframarginal effects and we assume (and test in Section 2) that the eligibility changes do not coincide with other changes to inframarginal status.

We estimate this equation at the demographic-cell level d. For each demographic cell, we need to estimate three inputs: $\frac{\partial c}{\partial m}_d$, s_d, and $\frac{\partial \ln(\text{N enrolled})}{\partial m}_d$. We use the following process:

1. **Estimating** $\frac{\partial \ln(\text{N enrolled})}{\partial m}_d$. We compute estimates of the state-level population within each demographic cell d. For each demographic cell d, we then estimate $\frac{\partial \ln(\text{N enrolled})}{\partial m}_d$ using the following equation:

$$\ln(\text{N enrolled})_{s,t,d} = \eta_d \ln(\text{share eligible})_{s,t,d} + X'_{s,t,d}\boldsymbol{\phi} + \delta_{s,d} + \gamma_{t,d} + \varepsilon_{s,t,d}, \tag{D.2}$$

instrumenting for $\gamma = \frac{\partial \ln(\text{N enrolled})}{\partial m}_d$ with the BBCE eligibility rate in each state-year. This a demographic-cell level version of Equation (3), estimated with a different outcome variable. The coefficient $\hat{\eta}_d$ corresponds to the desired parameter.

2. **Estimating** $\left(\frac{\partial c}{\partial m}\right)_d$. We assume the change in stigma costs is proportionate to the change in the second-order stigma index measured from the online experiment. Because the experiment is small at the demographic cell level, we estimate one regression:

$$y_i = \beta \mathbb{1}(\text{high})_i + \gamma \mathbb{1}(\text{truth})_i + \mathbf{X}_i \delta + (\mathbf{X}_i \mathbb{1}(\text{high})_i)\, \boldsymbol{\lambda} + \varepsilon_i. \tag{D.3}$$

Equation (D.3) linearly interacts the coefficient for several demographic groups (contained in $\mathbf{X}_i$) to obtain cell-level estimates of $\left(\frac{\partial c}{\partial m}\right)_d$ by summing the relevant entries of $\boldsymbol{\lambda}$ with β. To be concrete, X contains indicators for: female, white, age bins, income groups, and household size. To obtain at the demographic *cell* level, we sum the relevant coefficients for each cell. This approach is less flexible but more precise than fully saturating the model.

3. **Estimating** s_d. We use the FSPAS. We estimate the share of individuals affected by stigma (as the mean of the indicator variable) within each demographic cell, weighted using the FSPAS survey weights described above.

D.3 Additional details

Demographics. We focus on the following demographic variables: female/non-female, white/non-white, age group (0–18, 19–30, 31–65, 66+), household size (1, 2, 3, or 4), and income decile (grouping deciles 40–70 and 70–100). To construct demographic *cells*, we fully interact each variable. For instance, "single white women ages 19–30 in income decile 10" is an example of a demographic cell.

To focus on the group that is most affected by inframarginal effects, our estimates of $\frac{\partial \ln(\text{N enrolled})}{\partial m}_d$ use the population between 50–115% FPL. We cannot precisely limit to this group in the experiment, but we limit that to less than 130% of the FPL.

Bootstrap. We employ a bootstrap to estimate standard errors. When bootstrapping Equation (D.2), we compute a *Bayesian Bootstrap* with weights drawn from Exponential(1). We use a Bayesian Bootstrap for (D.2) because otherwise smaller demographic cells were not drawn in some bootstraps. We use a standard bootstrap for Equations (D.3) and when estimating the share of individuals affected by stigma.

Boostrap bias correction and hypothesis testing. Bootstrap estimates in Table 5 are bootstrap-bias corrected using the following standard procedure. Consider any parameter θ that we bootstrap. Let $\hat{\theta}$ denote the estimate from the data. Let $\hat{\theta}^b$ denote the estimate from bootstrap b. Denote the mean estimate of θ across B bootstraps as $\bar{\hat{\theta}} := B^{-1} \sum_b \hat{\theta}^b$. The bias-adjusted coefficient we present is: $2\hat{\theta} - \bar{\hat{\theta}}$. We compute a standard error by taking the sample standard deviation of bootstrap coefficients. We compute p values by testing the bias-corrected coefficient against the normal distribution.

Precision weighting. The regression used for Figure 6 and Table 5 use previously estimated demographic subgroup effects. Because we estimate these effects with noise, the dispersion in the effects — and thus in the data used to estimate Equation D.1 — will be larger than the true variation. Moreover, effects estimated in small cells will be estimated less precisely than effects estimated in larger cells. To adjust for this, in the binned scatterplots, we weight by the inverse of the product of the variance of the estimates; i.e., we give more weight to cells that are more precisely estimated. In Table 5, we weight by the inverse of the product of the variance of the coefficients estimated from the listed datasets (columns 1 and 2) and show robustness to the variance of the coefficient from the estimated dataset (i.e., not the product) (columns 3–5).

E Welfare Analysis Appendix

E.1 Structural Analysis

To solve for globally optimal solutions (under different parameter values) to the social planner's problem, we require assumptions about non-local behavior. Here, we provide details of these assumptions.

Welfare weights. We assume λ_θ is linear in m and satisfies the value of $\lambda_\theta / \lambda_m$ obtained from the inverse optimum exercise. We assume that $\lambda_0 = 1$. These two assumptions pin down a unique linear welfare weight schedule of inverse optimum welfare weights.

Linear take-up probabilities. We assume the take-up probability is linear in m. We assume a representative part-stigma part-information agent who obeys: $p^i(m) = p_0^i + \frac{\partial p}{\partial m} m + s \frac{\partial p}{\partial B} B$.

Using the values for the elasticities η_m and η_B, we obtain the slope $\frac{\partial p}{\partial m} = \eta_m \frac{p^*}{m^*}$, which then gives η_m^{i*} using that $\eta_m^i = \eta_m (1-s)^{-1}$ by Assumption 2. We obtain $\frac{\partial p}{\partial B} = \eta_B \frac{p^*}{B^*}$.

Obtaining optimal m and B. From the planner's budget and take-up probability, we obtain the average equilibrium SNAP benefit B^*. We invert Equation (8), which, together with the linearity assumptions above, delivers a unique value of optimal m^{opt} and B^{opt}. Intuitively, this approach obtains the values of m and B that satisfy the planner's optimality conditions we derived in Section 4. We solve this problem numerically using Matlab.

E.2 Marginal Value of Public Funds Approach

As a related alternative, we consider the policy of expanding eligibility within the context of its Marginal Value of Public Funds (MVPF) (Hendren, 2016; Hendren and Sprung-Keyser, 2020). This approach lets us relax the assumption that today's policy constitutes the naïve solution to Equation (8). It also permits us to probe other assumptions about agents' behavior and utility.

In this framework, the planner considers the ratio of benefits (willingness to pay for the policy) to the net cost to the government. Because our focus is on the redistributive nature of the policy, we ultimately consider welfare-weighted MVPFs, i.e., welfare impacts (per dollar of government expenditure). Note that equating the welfare-weighted MVPF of raising the means test to that of raising the benefit size recovers our main optimality condition.

We analyze the size of the bias in the welfare-weighted MVPF, which we define as:

$$\text{bias} := 100 \times \frac{\bar{\lambda}^n \text{MVPF}^n - \bar{\lambda}^w \text{MVPF}^w}{\bar{\lambda}^w \text{MVPF}^w}, \tag{E.1}$$

where MVPF^n is the MVPF when the eligibility threshold does not affect inframarginal recipients and MVPF^w is the MVPF when it does. $\bar{\lambda}^n$ and $\bar{\lambda}^w$ correspond to the average welfare weights of the beneficiaries of the policy in each case (denoted $\bar{\eta}$ in Hendren and Sprung-Keyser (2020)). Let Δ be the size of the eligibility threshold increase; for instance, $\Delta = 0.01$ when we study the welfare effect of letting 1 pp more people become eligible.

Derivation. The naïve welfare impact per dollar of government expenditure of an eligibility increase is: $\text{WI}^n := \bar{\lambda}^n \frac{\Delta p_m (1-\gamma s) \text{WTP}}{\Delta p_m \kappa_m} = \bar{\lambda}^n \frac{(1-\gamma s)\text{WTP}}{\kappa_m}$, with $\bar{\lambda}^n = \mu_m \lambda_m$, where μ_m is the marginal utility of income for newly eligible population. The denominator κ_m is the entire fiscal cost per person of providing food stamps to the next share Δ of the population (including the fiscal externality).[57] Assume the WTP to take up the benefit is the same for all θ.

Let WTP^s be the willingness to pay for the reduction in stigma costs from increasing the eligibility threshold to the next share Δ of the population. Let $\tilde{\alpha}$ satisfy $\text{WTP}^s := \tilde{\alpha}\text{WTP}$, where $\tilde{\alpha} < 1$, and μ_B and μ_0 are the marginal utilities of income of previously and newly enrolled, respectively.

[57] We continue to assume that labor supply is fixed and abstract from bunching. To relax this assumption, one could assume the newly eligible are willing to pay only some fraction βWTP and follow this through. This would correspond to "bunchers" having lower WTP for the higher eligibility threshold, since they are already eligible via a distortion in their labor supply. However, note that this would also correspond to a lower κ_m.

Stigma agents who newly take up due to the inframarginal effect are just indifferent due to the Envelope Theorem.[58] Stigma agents who previously took up the benefit have a positive willingness to pay for the reduction in stigma costs. We assume that information agents who newly take up due to the inframarginal effect have full willingness to pay for the benefit. Regardless of type, individuals who are newly eligible for the program under an eligibility expansion see first-order utility gains; information agents again gain the full WTP, and stigma agents are willing to pay $(1-\gamma)$WTP.[59] Suppose there is a share s of stigma agents.

The sophisticated welfare impact per dollar of government expenditure is:

$$\mathrm{WI}^w := \bar{\lambda}^w \frac{\overbrace{\Delta p_m (1-s\gamma)\mathrm{WTP}}^{\text{Mechanical effect for marginal types}} + \overbrace{\int_0^m \mathrm{WTP}(1-s)w^i d\theta}^{\text{Inframarginal effect for info agents}} + \overbrace{\int_0^m \tilde{\alpha} p_{\text{avg}} s \mathrm{WTP} d\theta}^{\text{Reduction in stigma costs for stigma agents}}}{\underbrace{\int_0^m \kappa_\theta w d\theta}_{\text{Cost of inframarginal effect}} + \underbrace{\Delta p_m \kappa_m}_{\text{Cost of expansion to marginal types}}} \tag{E.2}$$

with

$$\bar{\lambda}^w = \frac{\Delta \lambda_m \mu_m p_m (1-s\gamma)\mathrm{WTP} + \int_0^m \lambda_\theta \mu_\theta \mathrm{WTP}(1-s)w^i d\theta + \int_0^m \mu_\theta^E \lambda_\theta \tilde{\alpha} p_{\text{avg}} s \mathrm{WTP} d\theta}{\Delta p_m (1-s\gamma)\mathrm{WTP} + \int_0^m \mathrm{WTP}(1-s)w^i d\theta + \int_0^m \tilde{\alpha} p_{\text{avg}} s \mathrm{WTP} d\theta}$$

where μ_θ is the marginal utility of income for person θ who is not previously enrolled on the program, μ_θ^E is the marginal utility of income for person θ who is an 'always-taker,' t_{avg} is the take-up rate for inframarginal types prior to the eligibility expansion (which we assume is constant across all types), and κ_θ is the total fiscal cost of an additional 1 pp inframarginal take-up of type θ (including the fiscal externality). w is the proportion increase in inframarginal take-up (i.e., the inframarginal effect), and is a weighted average of the effect among information agents and the effect among stigma agents $w = (1-s)w^i + sw^s$.

Assume that all individuals who would newly enroll in the program have marginal utility of income $\mu_\theta = \mu_0$ before the policy change, and individuals who would were previously enrolled have marginal utility of income $\mu_\theta^E = \mu_B$. Define α as $\tilde{\alpha} * \frac{\mu_0}{\mu_B}$.

Noting that the bias can be written as $100 \times \frac{WI^n - WI^w}{WI^w}$, algebra gives:

$$\text{bias} = 100 \times \left(\frac{p_m + \frac{m}{\Delta} w \frac{\kappa_{\text{avg}}}{\kappa_m}}{p_m + \frac{m}{(1-s\gamma)\Delta} \frac{\lambda_{\text{avg}}}{\lambda_m} \left((1-s)w^i + s p_{\text{avg}} \alpha \right)} - 1 \right), \tag{E.3}$$

where $\kappa_{\text{avg}} := \frac{\int_0^m \kappa_\theta d\theta}{m}$, by analogy to λ_{avg}; p_m (p_{avg}) is the take-up rate for those newly eligible (previously eligible); w is the percentage point increase in the take-up rate for information-types (i.e., the inframarginal effect); κ_m (κ_{avg}) is the total fiscal cost of an additional 1 pp of take-up, including fiscal externalities, for those newly eligible (previously eligible); γ represents participation (stigma) costs as a share of WTP to take up the benefit; and α corresponds to the *reduction* in costs when the eligibility threshold rises, as a share of total WTP for the policy.[60]

Note also that, if $s = 0$, $w > 0$, and $\frac{\lambda_{\text{avg}}}{\lambda_m} > \frac{\kappa_{\text{avg}}}{\kappa_m}$, then bias < 0. Intuitively, as long as the planner's valuation of inframarginal types exceeds their fiscal cost, inframarginal effects raise the welfare impact of an eligibility increase. In the stigma case, the planner also values the reduction in costs to inframarginal types.

[58] We assume that there are not utility gains to individuals who are not decision-makers (e.g., the children of SNAP recipients). Otherwise, while newly-enrolled, inframarginal stigma agents have no first-order welfare gains, there would be utility gains from their children.

[59] This is analogous to the γ in the model in Section 4; stigma agents face costs which erode some fraction of their WTP.

[60] The advantage of focusing on the proportion bias in the welfare impact is that the expression does not require an estimate of willingness to pay or separate estimates of the costs κ_{avg} and κ_m. The magnitudes of these costs are difficult to estimate, because SNAP involves many fiscal externalities that plausibly vary by type. This exercise permits us to conduct welfare analysis with only the *ratio* ($\kappa_{\text{avg}}/\kappa_m$).

With this approach, we relax several assumptions imposed in Section 4. While our model in Section 4 defines the gains to inframarginal stigma agents in relation to the size of the inframarginal effects, this exercise decouples them (since α and w enter separately). This provides the flexibility to incorporate welfare gains from decreases in stigma costs even in the absence of an effect of stigma on take-up. Moreover, this expression permits inframarginal program participants to have different costs from participants who are newly eligible.[61] This is at the expense of additional assumptions (e.g., on the size of α), but it is easier to see robustness to those assumptions. We also emphasize the role of the welfare weights here: the policy evaluation is as much about its incidence as it is about its utility gains and fiscal costs.

Parameters. We estimate w using instrumental variables, as in Table 4, where we instrument for the share of a state's population that is eligible using the eligibility threshold; here, however, we regress the take-up rate on the share eligible, instead of estimating an elasticity. The result is $w = 0.0028$: take-up increases by 0.28 percentage points for every 1 percentage point increase in the share eligible. From Section 3, we assume that $w^s = 0$ and $s = 0.4$, such that $w^i = 0.0047$. We continue to assume today's take-up rate, $p_{avg} = p_m = 0.53$, and eligibility threshold $m = 0.27$. We assume $\gamma = 0.5$, analogous to the calibration used elsewhere in the paper. Finally, we use $\frac{\lambda_{\text{avg}}}{\lambda_m}$ derived from inverse-optimum weights, although we note that these employ an assumption MVPFs usually relax — that current policy is optimal under a certain model. However, our results are robust to a range of values for $\frac{\lambda_{\text{avg}}}{\lambda_m}$.

Results. We evaluate the welfare impact per dollar of government expenditure of expanding eligibility by 1 pp, i.e. we set $\Delta = 0.01$. To be conservative, we assume that the willingness to pay for a reduction in stigma costs is small, so we set $\alpha = 0.02$.

We find that, if $\frac{\kappa_{\text{avg}}}{\kappa_m} = 1$, the naïve MVPF can be about 20% below the sophisticated MVPF for the information-only case, with even larger results in the information and stigma case (Figure E.1A). However, the planner may overvalue the welfare impact for larger values of $\frac{\kappa_{\text{avg}}}{\kappa_m}$, say $\frac{\kappa_{\text{avg}}}{\kappa_m} = 10$ (Figure E.1B). This is because with $\frac{\kappa_{\text{avg}}}{\kappa_m} \gg 1$, the cost of new participants who are costly may exceed their value to the planner. Hence the naïve planner sets the eligibility threshold *too high*.[62]

[61] Note that while inframarginal types tend to have higher benefits, the higher benefit may yield a *reduced* fiscal externality because people with higher SNAP benefits receive better educations or are less likely to be incarcerated; Bailey et al. (2020) show that these benefits reduce the denominator of the MVPF for a benefit increase.

[62] Note that if $\frac{\kappa_{\text{avg}}}{\kappa_m} \gg 1$, the MVPF bias is negative for the stigma and information case ($s = 0.5$) and positive for the information-only case ($s = 0$). Here, unlike in the model, the normative conclusion that the planner may wish to raise the eligibility threshold can be stronger if there is stigma.

Figure E.1: Welfare Bias of an Eligibility Expansion using MVPF Framework

(A) $\frac{\kappa_{\text{avg}}}{\kappa_m} = 1$

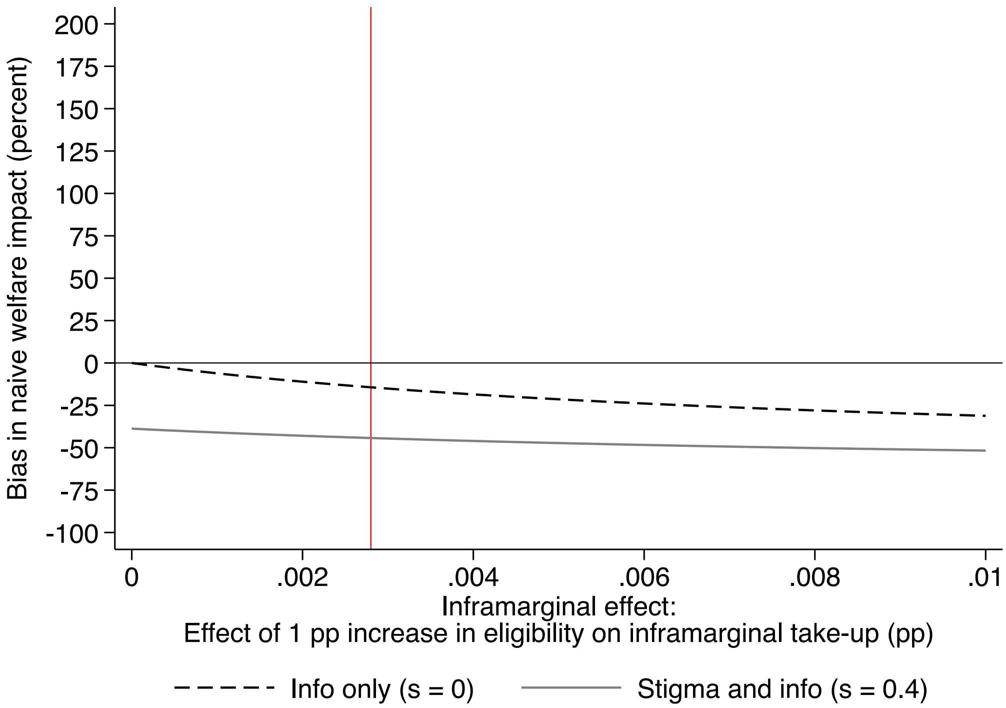

(B) $\frac{\kappa_{\text{avg}}}{\kappa_m} = 10$

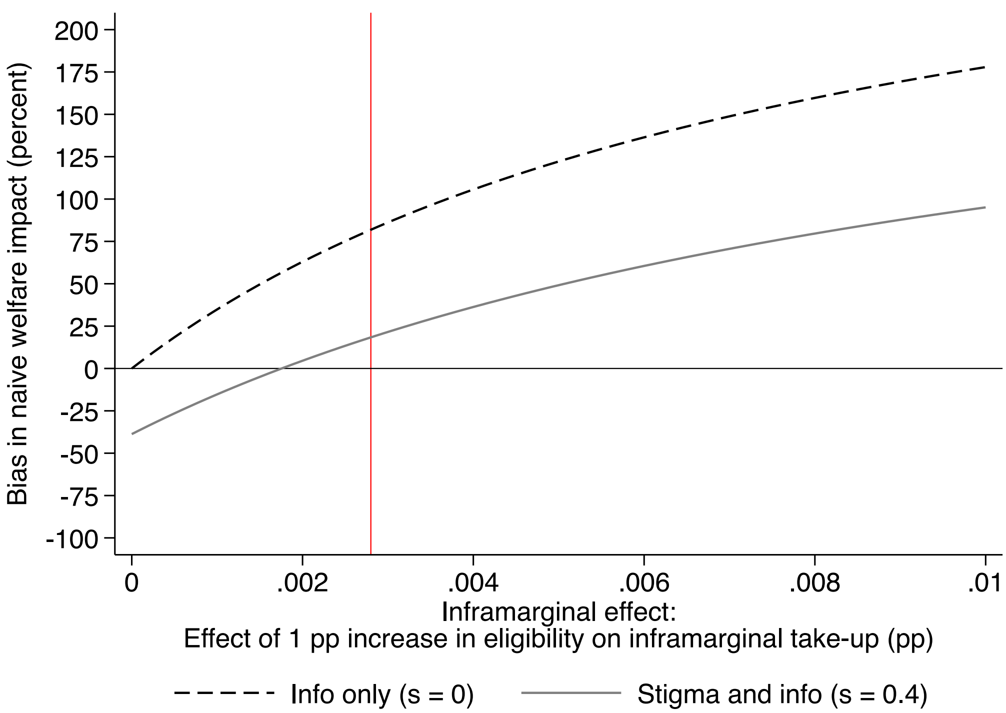

This figure shows the percent bias in the welfare impact per dollar of government expenditure (Equation (E.3)) for an inframarginal to marginal cost ratio of 1 (Panel A) and 10 (Panel B). The vertical red line plots our preferred estimate of the inframarginal effect w in terms of take-up.

E.3 Robustness

Figure E.2: Robustness: Naïve Planner's Biased Risk Aversion

(A) True $\rho = 2$

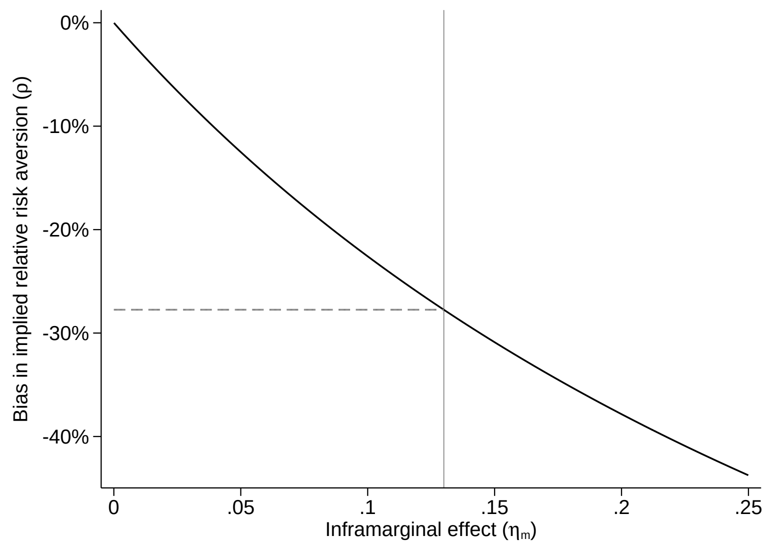

(B) $\eta_B = 0.3$

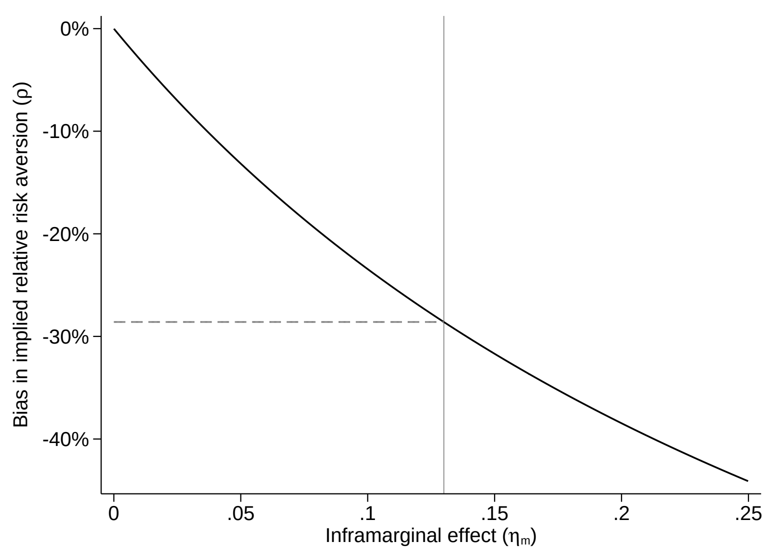

This figure shows the percent bias between the planner's "as-if" risk aversion ($\tilde{\rho}$) and the ground-truth risk aversion (ρ) (black line). It is identical to Figure 7A except it sets $\rho = 2$ (Panel A) or $\eta_B = 0.3$ (Panel B). Negative numbers indicate that the planner is behaving as if people are less risk averse than they really are. Panel A plots the bias as a function of the inframarginal effect; the vertical gray line plots the empirical inframarginal effect presented in Table 6.

Figure E.3: Numerical Simulations: Robustness

(A) Varying the Coefficient of Relative Risk Aversion (ρ)

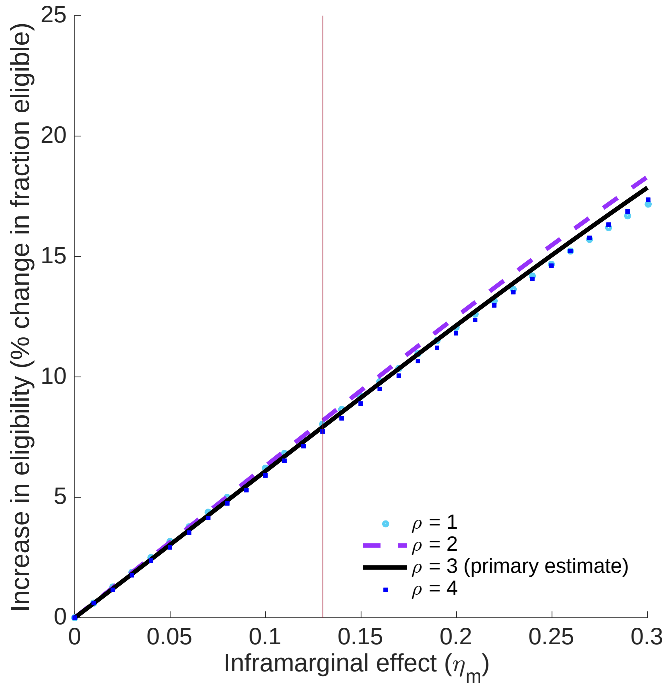

(B) Varying Take-Up Elasticity with Respect to Benefit Size (η_B)

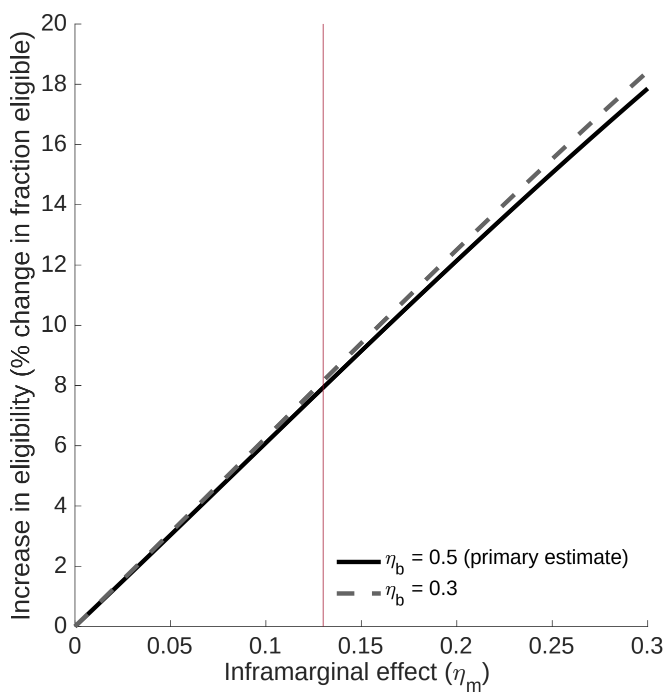

This figure shows the results from our numerical simulation exercise, which uses the optimality condition in Equation (8). It presents changes in the percent of people who are eligible if the planner acknowledges inframarginal effects. It shows robustness to different ρ (Panel A) and take-up elasticities with respect to the benefit size (η_B). Auxiliary parameters are set according to the values in Table 6.

Figure E.4: Numerical Simulations: Optimal Benefit Size

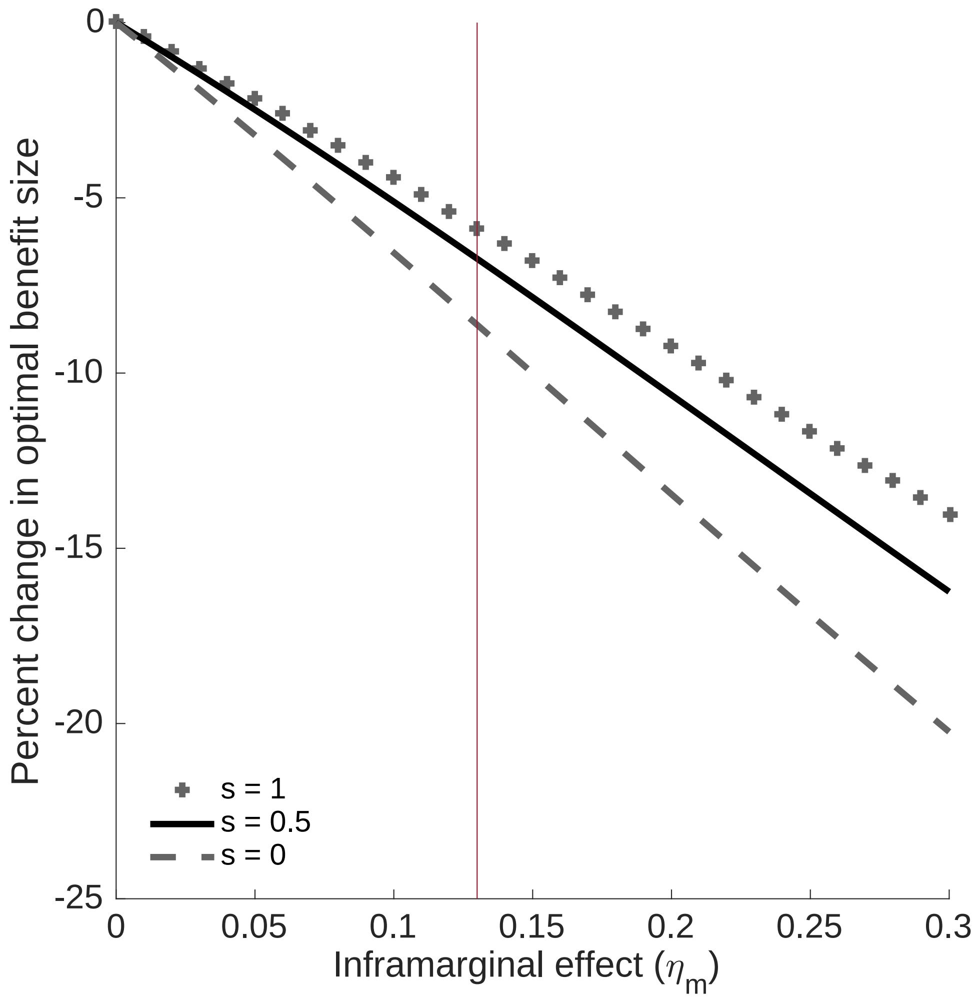

This figure shows the results from our numerical simulation exercise, which uses the optimality condition in Equation (8). It presents the percent change in the optimal benefit size if the planner acknowledges inframarginal effects. Auxiliary parameters are set according to the values in Table 6.

Figure E.5: Numerical Simulations: Robustness to Quadratic Utility

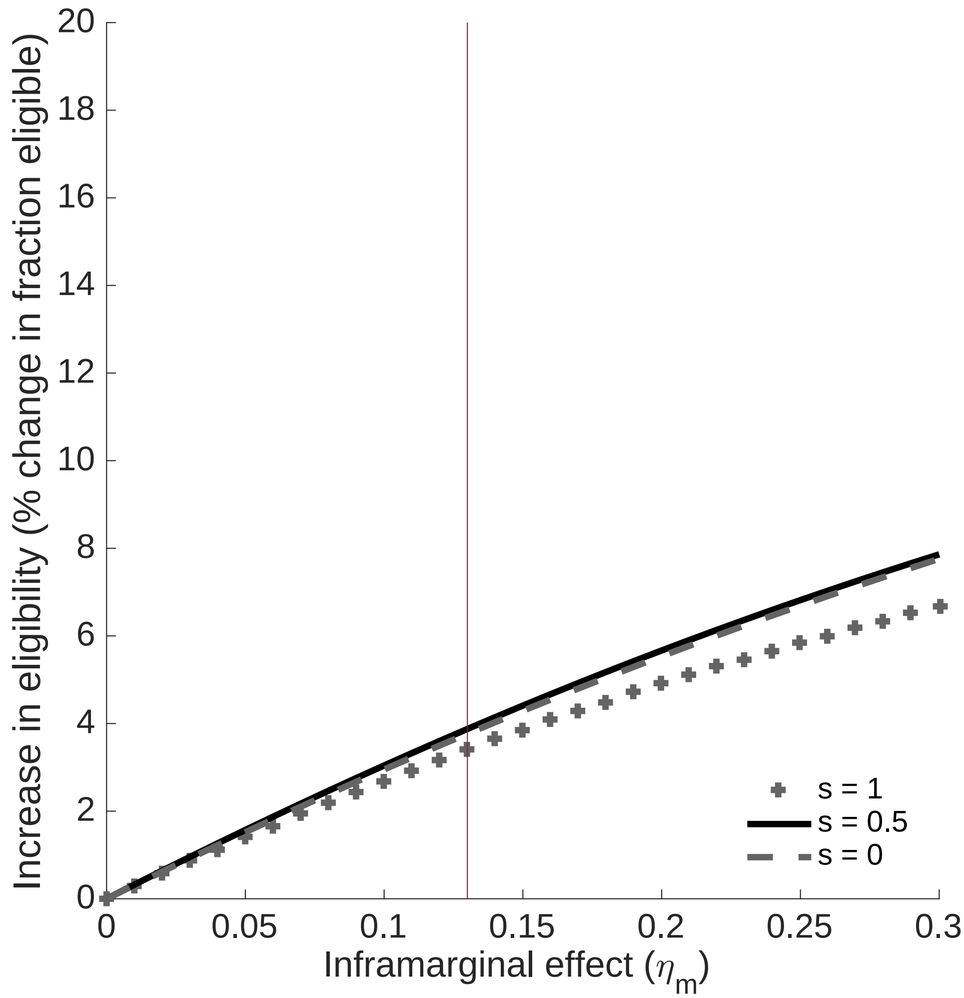

This figure shows the results from our numerical simulation exercise, which uses the optimality condition in Equation (8). It presents the change in the percent of people who are eligible if the planner acknowledges inframarginal effects. It is identical to Figure 8A except the simulations impose quadratic utility with $\rho = 3$ at equilibrium, using Equation (1) with $\eta_m = 0$ to infer the welfare weights.

F Proofs

F.1 Proofs of Propositions 1 and 2.

Proof. Note that Proposition 1 is a special case of Proposition 2. We therefore prove Proposition 2 only.

The planner's problem is:

$$\max_{B,m}\left(sp^s(B,m)\left(\int_0^m \lambda_\theta u(B)d\theta-\int_0^m\int_{c\le u(B)}\lambda_\theta ch(c|c<u(B),m)dcd\theta\right)+(1-s)\left(\int_0^m \lambda_\theta u(B)p^i(m)d\theta\right)\right) \tag{F.1}$$

subject to

$$(1-s)p^i(m)\int_0^m Bd\theta+sp^s(B,m)\int_0^m Bd\theta\le T \tag{F.2}$$

$$m\in[0,1] \tag{F.3}$$

We inspect interior solutions using the Karush-Kuhn-Tucker conditions where the constraint $m\in[0,1]$ is slack. We consider cases in which such an interior solution exists; there are possible corner solutions where $m=1$ (i.e., the program is universal). Proposition 1 and 2 give necessary conditions for local optimality. To obtain that the statement in the proposition is sufficient for a global maximum, it is sufficient to additionally impose that the maximand is concave and the constraint is convex.

The first-order condition for B is:

$$\begin{aligned}&s\left(\frac{\partial p^s}{\partial B}\lambda_{\text{avg}}mu(B)+p^s(B,m)u'(B)\lambda_{\text{avg}}m-\left(\frac{\partial}{\partial B}\int_0^m\int_{c\le u(B)}c\lambda_\theta h(c|m)dcd\theta\right)\right)\\&+(1-s)p^i(m)u'(B)\lambda_{\text{avg}}m=\sigma\left((1-s)p^i(m)m+s\frac{\partial p^s}{\partial B}Bm+sp^sm\right),\end{aligned} \tag{F.4}$$

where σ denotes the Lagrange multiplier, and we note that

$$\int_0^m\int_{c\le u(B)}\lambda_\theta h(c|c<u(B),m)dcd\theta=\frac{1}{H(u(B)|m)}\int_0^m\int_{c\le u(B)}\lambda_\theta h(c|m)dcd\theta. \tag{F.5}$$

Leibniz's rule gives that:

$$\frac{\partial}{\partial B}\int_0^m\int_{c\le u(B)}c\lambda_\theta h(c|m)dcd\theta=u(B)\lambda_{\text{avg}}mh(u(B)|m)u'(B) \tag{F.6}$$

$$=\lambda_{\text{avg}}mu(B)\frac{\partial p^s}{\partial B}. \tag{F.7}$$

We collect terms to obtain:

$$\lambda_{\text{avg}}m(sp^s+(1-s)p^i)u'(B)=\sigma\left(s\frac{\partial p^s}{\partial B}Bm+sp^sm+(1-s)p^im\right). \tag{F.8}$$

We divide by p^sm and rearrange, recalling that $\eta_B=\frac{\partial p^s}{\partial B}\frac{B}{p^s}$:

$$\frac{sp^s+(1-s)p^i}{p^s}u'(B)\lambda_{\text{avg}}=\sigma\left(s\eta_B+\frac{sp^s+(1-s)p^i}{p^s}\right). \tag{F.9}$$

Next we take the first-order condition with respect to m and use the shorthand $E:=E[c|c<u(B)]$ to be

succinct:

$$s\frac{\partial p^s}{\partial m}\left(\lambda_{\text{avg}} m u(B) - \lambda_{\text{avg}} m E\right) + s p^s \left(\lambda_m u(B) - \lambda_m E - \lambda_{\text{avg}} m \frac{\partial E}{\partial m}\right)$$
$$+ (1-s)\left(\frac{\partial p^i}{\partial m}\lambda_{\text{avg}} m u(B) + p^i \lambda_m u(B)\right) = \sigma\left((1-s)\left(\frac{\partial p^i}{\partial m}Bm + p^i B\right) + s\left(\frac{\partial p^s}{\partial m}Bm + p^s B\right)\right). \quad \text{(F.10)}$$

Noting that $\frac{\partial E}{\partial m} = \frac{\partial \gamma}{\partial m}u(B)$, we collect terms to obtain:

$$u(B)\lambda_{\text{avg}}\left((1-\gamma)s\frac{\partial p^s}{\partial m}m + (1-s)\frac{\partial p^i}{\partial m}m - sp^s\frac{\partial \gamma}{\partial m}m\right) + \lambda_m u(B)(1-\gamma)sp^s$$
$$+ \lambda_m u(B)p^i(1-s) = \sigma\left((1-s)\left(\frac{\partial p^i}{\partial m}Bm + p^i B\right) + s\left(\frac{\partial p^s}{\partial m}Bm + p^s B\right)\right). \quad \text{(F.11)}$$

We divide by p^s and p^i to get:

$$u(B)\lambda_{\text{avg}}\left(\frac{(1-\gamma)s\eta_m^s}{p^i} + \frac{(1-s)\eta_m^i}{p^s} - \frac{s\frac{\partial \gamma}{\partial m}m}{p^i}\right)$$
$$+ s\frac{\lambda_m u(B)(1-\gamma)}{p^i} + \frac{\lambda_m u(B)(1-s)}{p^s} = B\sigma\left((1-s)\frac{\eta_m^i+1}{p^s} + s\frac{\eta_m^s+1}{p^i}\right). \quad \text{(F.12)}$$

Then, substituting for σ from Equation (F.9) and rearranging gives:

$$\frac{u(B)\left(\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\eta_m^i + \frac{1-s}{p^s}\right) + \frac{s}{p^i}[(1-\gamma)\eta_m^s - m\frac{\partial \gamma}{\partial m}]\right)}{u'(B)} = \frac{B\left[\frac{(1-s)}{p^s}\left(\eta_m^i+1\right) + \frac{s}{p^i}\left(\eta_m^s+1\right)\right]}{\left(\frac{p^s}{sp^s+(1-s)p^i}s\eta_B + 1\right)}. \quad \text{(F.13)}$$

At this point, rearranging terms and using a Taylor Expansion (see Section F.2) produces the result in Equation 10.

To produce Equation 8, we invoke the following lemma:

Lemma 1. $\eta_m(1-\gamma) - m\frac{\partial \gamma}{\partial m} = \frac{\frac{m}{p(B,m)}\int_0^{u(B)}\frac{\partial H(c|m)}{\partial m}dc}{u(B)}$.

The proof is below. Now we can rewrite the expression as:

$$\frac{\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)u(B) + \frac{1-s}{p^s}\eta_m^i u(B) + \frac{s}{p^i}\frac{m}{p^s}\int_0^{u(B)}\frac{\partial H(c|m)}{\partial m}dc}{u'(B)} = \frac{B\left[\frac{(1-s)}{p^s}\left(\eta_m^i+1\right) + \frac{s}{p^i}\left(\eta_m^s+1\right)\right]}{\left(\frac{p^s}{sp^s+(1-s)p^i}s\eta_B + 1\right)} \quad \text{(F.14)}$$

An important special case is where $p^i = p^s$ and $\frac{\partial p^i}{\partial m} = \frac{\partial p^s}{\partial m}$. Then, multiplying both the numerator and denominator by $p := p^i = p^s$, and noting that $\eta_m^i = \eta_m^s$, we get:

$$\frac{u(B)}{Bu'(B)} = \frac{\eta_m + 1}{(s\eta_B+1)\left(\left(\eta_m + \frac{\lambda_m}{\lambda_{\text{avg}}}\right)(1-s\gamma) - s\frac{\partial \gamma}{\partial m}m\right)}. \quad \text{(F.15)}$$

□

F.1.1 Proof of Lemma 1

Proof. Multiplying both sides by $u(B)$ and recalling that $\frac{\partial E}{\partial m}(c|c < u(B), m) = \frac{\partial \gamma}{\partial m}u(B)$, it suffices to show that $\eta_m u(B)(1-\gamma) - m\frac{\partial E[c|c\leq u(B),m]}{\partial m} = \frac{m}{p(B,m)}\int_0^{u(B)}\frac{\partial H(c|m)}{\partial m}dc$. Below, we show that $m\frac{\partial E[c|c\leq u(B),m]}{\partial m} =$

$\eta_m(u(B) - E[c|c \leq u(B), m]) - \frac{m}{p(B,m)} \int_0^{u(B)} \frac{\partial H(c|m)}{\partial m} dc$, which completes the proof.

$$\begin{aligned}
m\frac{\partial E[c|c \leq u(B), m]}{\partial m} &= m \int_0^{u(B)} c \frac{\partial h(c|m, c < u(B))}{\partial m} dc && \text{(F.16)}\\
&= m \int_0^{u(B)} c \frac{\partial}{\partial m} \left(\frac{h(c|m)}{H(u(B)|m)} \right) dc && \text{(F.17)}\\
&= m \int_0^{u(B)} c \left(\frac{\frac{\partial h(c|m)}{\partial m} H(u(B)|m) - h(c|m) \frac{\partial H(u(B)|m)}{\partial m}}{H(u(B)|m)^2} \right) dc && \text{(F.18)}\\
&= m \left(\frac{1}{H(u(B)|m)} \int_0^{u(B)} c \frac{\partial h(c|m)}{\partial m} dc - \frac{1}{H(u(B)|m)^2} \int_0^{u(B)} ch(c|m) \frac{\partial p^s}{\partial m} dc \right), && \text{(F.19)}
\end{aligned}$$

where p^s is the take-up rate among stigma agents ($p^s = H(u(B)|m)$, the share of stigma agents with stigma costs below the utility benefits of take-up). We apply integration by parts to the first integral. We apply that $\frac{\int_0^{u(B)} chdc}{H(u(B)|m)} = E[c|c < u(B), m]$ to the second integral. Suppressing arguments of h and H to be concise, this yields:

$$\begin{aligned}
& m \left(\frac{1}{H} \int_0^{u(B)} c \frac{\partial h}{\partial m} dc - \frac{1}{H^2} \int_0^{u(B)} ch \frac{\partial p^s}{\partial m} dc \right) && \text{(F.20)}\\
&= m \left(\frac{1}{H} \left(u(B) \frac{\partial H(u(B)|m)}{\partial m} - \int_0^{u(B)} \frac{\partial H(c|m)}{\partial m} dc \right) - \frac{\frac{\partial p^s}{\partial m}}{p^s} E[c|c < u(B), m] \right) && \text{(F.21)}\\
&= \eta_m^s \left(u(B) - E[c|c < u(B), m] \right) - \frac{m}{p^s(B, m)} \int_0^{u(B)} \frac{\partial H(c|m)}{\partial m} dc, && \text{(F.22)}
\end{aligned}$$

recalling that $\frac{\partial p^s}{\partial m} \frac{m}{H(u(B)|m)} = \eta_m^s$. □

F.2 Proof of Taylor Expansion (Equation (8)).

Proof. Throughout the paper, we use the second-order Taylor approximation:

$$u(0) = 0 \approx u(B) - u'(B)B + \frac{u''(B)B^2}{2}, \tag{F.23}$$

which gives

$$u(B) \approx u'(B)B - \frac{u''(B)B^2}{2}. \tag{F.24}$$

We then obtain

$$\frac{u(B)/B}{u'(B)} \approx 1 + \frac{\rho}{2}. \tag{F.25}$$

Note that $u'(B) = \frac{\partial}{\partial B} (u(B) - c)$, so ρ represents the coefficient of relative risk aversion for people who would take up the program if informed. □

F.3 Lemma 2 and Proof

Subsequent proofs invoke the following lemma:

Lemma 2. *If* $\rho \geq 1$, $\frac{\partial}{\partial B} \left(\frac{u(B)/B}{u'(B)} \right) > 0$.

Proof. The quotient rule gives

$$\frac{\partial}{\partial B} \left(\frac{u(B)/B}{u'(B)} \right) > 0 \tag{F.26}$$

iff

$$(u'(B))^2 B - u(B)u'(B) - Bu''(B)u(B) > 0. \tag{F.27}$$

Dividing by $u'(B)$ (which is always greater than 0), we conclude that the left-hand side is always positive as long as

$$u'(B)B + u(B)(\rho - 1) > 0, \tag{F.28}$$

which completes the proof. □

G Theory Extensions

G.1 Endogenous labor supply

Section 4 develops a proposition that gives that B and m satisfy

$$\frac{u(B)}{u'(B)B} = \frac{1+\eta_m}{\lambda_m/\lambda_{\text{avg}} + \eta_m}, \tag{G.1}$$

if $\eta_B = 0$ and $s = 0$.

We show how this expression can be microfounded in a more elaborate environment with endogenous labor supply. We focus on this parsimonious expression, nested by the more general case, for simplicity; this analysis captures many of the relevant insights.

Model environment. There is a continuum of types $\theta \sim F$, where F has support Θ. People earn labor income y from hours worked h, depending on their type θ. Let labor income $y = \theta h(\theta)$; we use this parametric form for simplicity, but the model can easily be generalized. People with labor income below r (the "eligibility threshold") earn a benefit B. People have utility $\tilde{v}(h, B, \theta)$ over labor supply and the benefit amount.[63] This utility function induces an indirect utility function v over labor supply, the benefit amount, and the eligibility threshold:

$$v(h^*(B,r,\theta), B, r, \theta) = \max_h \begin{cases} \tilde{v}(h, B, \theta) & \theta h(\theta) \leq r \\ \tilde{v}(h, 0, \theta) & \theta h(\theta) > r. \end{cases} \tag{G.2}$$

The Envelope Theorem gives the following intermediate results, which we will invoke later:

$$\frac{dv}{dB}(h^*(B,r,\theta), B, r, \theta) = \frac{\partial v}{\partial B} \tag{G.3}$$

$$\frac{dv}{dr}(h^*(B,r,\theta), B, r, \theta) = 0 \text{ if } h^*(B,r,\theta) \neq r/\theta \tag{G.4}$$

Equation (G.4) states that if the benefit constraint does not bind, there is no value to the agent to relaxing the constraint. Intuitively, for people who are very poor or very rich, adjustments to the eligibility threshold have no effect on behavior. However, the existence of a lump-sum benefit and discrete eligibility threshold can induce bunching at the threshold. A small change in eligibility will have first-order effects on utility for bunchers.

Take-up probabilities. Agents are aware of the program with probability $p(r)$ and get $v(h^*(B,r,\theta), B, r)$ if they take up. Otherwise they optimize as if the program does not exist, do not take up the program, and get $v(h^*(0,r,\theta); 0, r, \theta)$ (the "outside option"). Moreover, this outside option does not depend on r: $v(h^*(0,r,\theta); 0, r, \theta) = v(h^*(0,\theta); 0, \theta)$ for all r.

Planner's problem. We begin with a technical assumption. Assume that income $\theta h^*(B,r,\theta)$ is weakly increasing in θ: higher types always earn weakly more labor income even though the existence of the benefit distorts labor supply. This assumption amounts to a standard single-crossing condition: even if the tax system affects labor supply or causes bunching, it will not cause high types to earn strictly less income than low types (or vice-versa).

[63] We can think of utility over the benefit as the indirect utility of the agent's inner problem of allocating the benefit to consumption of various goods.

This assumption yields a threshold type $\tilde{\theta}(B,r)$ such that all $\theta \leq \tilde{\theta}$ will choose a labor supply that is low enough that they will be eligible for the benefit. All types $\theta > \tilde{\theta}$ are not eligible.

Next, we assume that the planner has a budget T which depends on the amount of money raised through taxes on labor income. Assume the income tax schedule is exogenous, but make no other restrictions on this schedule. In that case, we can parameterize T as depending on B and r alone: $T(B,r)$.[64]

Altogether, the planner's problem is:

$$\max_{r,B} \int_0^{\tilde{\theta}(B,r)} \lambda_\theta p(r) v(h^*(B,r,\theta);B,r,\theta) f(\theta) d\theta + \int_0^{\tilde{\theta}(B,r)} \lambda_\theta (1-p(r)) v(h^*(0,r,\theta);0,\theta) f(\theta) d\theta + \int_{\tilde{\theta}(B,r)}^{\infty} \lambda_\theta v(h^*(0,r,\theta);0,\theta) f(\theta) d\theta \tag{G.5}$$

subject to

$$\int_0^{\tilde{\theta}(B,r)} p(r) B f(\theta) d\theta \leq T(B,r). \tag{G.6}$$

Noting that $\int_0^\infty \lambda_\theta v(h^*(0);0,\theta) f(\theta) d\theta$ is a constant, we can re-write the planner's problem as:

$$\max_{r,B} \int_0^{\tilde{\theta}(B,r)} \lambda_\theta p(r) \left(v(h^*(B,r,\theta);B,r,\theta) - v(h^*(0,r,\theta);0,\theta) \right) f(\theta) d\theta \tag{G.7}$$

subject to

$$\int_0^{\tilde{\theta}(B,r)} p(r) B f(\theta) d\theta \leq T(B,r). \tag{G.8}$$

Then, let $V(h^*(B,r,\theta);B,r,\theta) := v(h^*(B,r,\theta);B,r,\theta) - v(h^*(0,r,\theta);0,r,\theta)$ be the net utility gain from taking up the program. Note that for types $\theta > \tilde{\theta}$, $V = 0$: these types choose labor supply that renders them ineligible for the benefit. For other types, $\theta \leq \tilde{\theta}$, $V > 0$ assuming they earn positive utility from the benefit.

Solving for the optimum. Letting σ represent the Lagrange multiplier, take the first-order condition with respect to r:

$$\frac{\partial \tilde{\theta}}{\partial r} \left(\lambda_{\tilde{\theta}} p(r) V(h^*(B,r,\tilde{\theta});B,r,\tilde{\theta}) f(\tilde{\theta}) \right) + \int_0^{\tilde{\theta}} \lambda_\theta \left(\frac{dp}{dr} V(h^*(B,r,\theta);B,r,\theta) + p \left(\underbrace{\frac{dV}{dr} \cdot \mathbb{1}(\theta = \tilde{\theta})}_{\text{As } \frac{dV}{dr}=0 \text{ otherwise, by Equation (G.4)}} \right) \right) f(\theta) d\theta$$

$$-\sigma \left(\frac{\partial \tilde{\theta}}{\partial r} (pBf(\tilde{\theta})) + \frac{dp}{dr} BF(\tilde{\theta}) - \frac{dT}{dr} \right) = 0. \tag{G.9}$$

Take the first-order condition with respect to B:

$$\frac{\partial \tilde{\theta}}{\partial B} \left(\lambda_{\tilde{\theta}} p V(h^*(B,r,\tilde{\theta});B,r,\tilde{\theta}) f(\tilde{\theta}) \right) + \int_0^{\tilde{\theta}} \lambda_\theta p(r) \underbrace{\frac{\partial V}{\partial B}}_{=\frac{dV}{dB}\text{, by Equation (G.3)}} f(\theta) d\theta -$$

$$\sigma \left(\frac{\partial \tilde{\theta}}{\partial B} p(r) B f(\tilde{\theta}) + p(r) F(\tilde{\theta}) - \frac{dT}{dB} \right) = 0. \tag{G.10}$$

[64]Formally, let

$$I(B,r) := \{(\theta h^*(B,r,\theta), \theta h^*(0,r,\theta), \theta) : \theta \in \Theta\}.$$

Here, I is the set of triples of: (i) labor incomes chosen if a given type θ receives the benefit, (ii) labor incomes chosen if the type θ does not receive the benefit, and (iii) the type θ, which then yields a density $f(\theta)$ and a labor supply $h^*(B,r,\theta)$. These values uniquely determine the taxes raised for a generic tax schedule that only depends on labor income, even if there is incomplete take-up of the benefit, assuming the planner knows $p(r)$. This notation shows that we can write $T(B,r) = T(I(B,r))$. Intuitively, holding F fixed, any (B,r) pair induces a distribution of labor incomes chosen across types.

Solving for σ, we obtain:

$$\sigma = -\frac{\frac{\partial\tilde{\theta}}{\partial B}\lambda_{\tilde{\theta}}p(r)V(h^*(B,r,\tilde{\theta});B,r,\theta)f(\tilde{\theta}) + \int_0^{\tilde{\theta}}\lambda_\theta p(r)\frac{\partial V}{\partial B}f(\theta)d\theta}{\frac{\partial\tilde{\theta}}{\partial B}p(r)Bf(\tilde{\theta}) + p(r)F(\tilde{\theta}) - \frac{dT}{dB}} \tag{G.11}$$

Plugging into Equation (G.9) yields:

$$\begin{aligned}
&\overbrace{\frac{\partial\tilde{\theta}}{\partial r}\left(\lambda_{\tilde{\theta}}p(r)V(h^*(B,r,\tilde{\theta});B,r,\tilde{\theta})f(\tilde{\theta})\right)}^{\text{Value of } r\uparrow \text{ to otherwise ineligible people}} + \int_0^{\tilde{\theta}}\lambda_\theta\left(\underbrace{\frac{dp}{dr}V(h^*(B,r,\theta);B,r,\theta)}_{\text{Effect of } r\uparrow \text{ on take-up of inframarginals}} + \overbrace{p(r)\left(\frac{dV}{\partial r}\cdot\mathbb{1}(\theta=\tilde{\theta})\right)}^{\text{Value of } r\uparrow \text{ to bunchers}}\right)f(\theta)d\theta \\
&-\left(\underbrace{\frac{\partial\tilde{\theta}}{\partial B}\lambda_{\tilde{\theta}}p(r)V(h^*(B,r,\theta);B,r,\theta)f(\tilde{\theta})}_{\text{Value of } B\uparrow \text{ to bunchers}} + \overbrace{\int_0^{\tilde{\theta}}\lambda_\theta p(r)\frac{\partial V}{\partial B}f(\theta)d\theta}^{\text{Value of } B\uparrow \text{ to inframarginals}}\right) \\
&\times\frac{\underbrace{\frac{\partial\tilde{\theta}}{\partial r}(pBf(\tilde{\theta}))}_{\text{Mechanical cost of } r\uparrow} + \overbrace{\frac{dp}{dr}BF(\tilde{\theta})}^{\text{Indirect cost of } r\uparrow \text{ from changes in take-up}} - \underbrace{\frac{dT}{dr}}_{\text{Indirect cost of } r\uparrow \text{ from changes in taxes via change in labor supply}}}{\underbrace{\frac{\partial\tilde{\theta}}{\partial B}p(r)Bf(\tilde{\theta})}_{\text{Indirect cost of } B\uparrow \text{ from changes in take-up}} + \overbrace{p(r)F(\tilde{\theta})}^{\text{Mechanical cost of } B\uparrow} - \underbrace{\frac{dT}{dB}}_{\text{Indirect cost of } B\uparrow \text{ from changes in taxes via change in labor supply}}} = 0.
\end{aligned} \tag{G.12}$$

Discussion. Equation (G.12), while involved, captures the following intuitions. At an optimum, the planner equates the following trade-offs.

- Raising r has benefits. First, it brings in more people to the program who were previously ineligible. Second, it has the value of raising take-up among inframarginal types. Third, it has a direct effect on welfare for people who bunch at the eligibility threshold, who can then adjust their labor supply (which was not necessarily at an optimum).
- Raising B has benefits. First, it brings value by affecting bunching. Second, it also has value to inframarginal types who take up the program, because it is a transfer.
- Raising r has costs. First, there is a mechanical cost of bringing more people into the program because more people are eligible. Second, there is an indirect cost of raising take-up. Third, there is an indirect cost of changing people's labor supply, which then affects the income taxes collected.
- Raising B has costs. First, there is a mechanical cost of raising the transfer to people who take up the program. Second, there is an indirect cost of bringing more people into the program via changes in labor supply. Third, there is an indirect cost of changing people's labor supply, which then affects the income taxes collected.

G.1.1 Simplifications

In this subsection, we show how this more general solution nests the solution in the paper.

First, we apply a change of units. Instead of considering raising the eligibility threshold by one dollar of labor income, we raise the eligibility threshold by one quantile of the population that is eligible. Let m represent the share who is eligible for the benefit: $m := F(\tilde{\theta})$.

Use the chain rule to observe that:

$$\frac{\partial p}{\partial r} = \frac{\partial p}{\partial F(\tilde{\theta})} \frac{\partial F(\tilde{\theta})}{\partial \tilde{\theta}} \frac{\partial \tilde{\theta}}{\partial r} = \frac{\partial p}{\partial m} f(\tilde{\theta}) \frac{\partial \tilde{\theta}}{\partial r}. \tag{G.13}$$

Next, we invoke the following assumption:

Assumption S1: No Bunching. Assume that $h^*(\cdot) = \bar{h}(\theta)$ for all B, r, i.e. that the amount of labor supply chosen depends only on one's type.

This assumption has three implications. First, $\int_0^{\tilde{\theta}} \mathbb{1}(\theta = \tilde{\theta}) f(\theta) = 0$, assuming there are no atoms in the type distribution. Second, $\frac{\partial \tilde{\theta}}{\partial B} = 0$. Third, since labor supply is constant for all θ and the budget T only depends on r and B via h, $\frac{dT}{dr} = \frac{dT}{dB} = 0$.

As a result, employing the No Bunching assumption and dividing by $f(\tilde{\theta}) \frac{\partial \tilde{\theta}}{\partial r}$ gives:

$$\lambda_{\tilde{\theta}} p(r) V(h^*(B, r, \tilde{\theta}); B, r, \tilde{\theta}) + \int_0^{\tilde{\theta}} \lambda_\theta \frac{\partial p}{\partial m} V(h^*(B, r, \theta), B, r, \theta) f(\theta) d\theta = \left(\int_0^{\tilde{\theta}} \lambda_\theta \frac{\partial V}{\partial B} f(\theta) d\theta \right) \left(pB + \frac{\partial p}{\partial m} BF(\tilde{\theta}) \right) \frac{1}{F(\tilde{\theta}))}. \tag{G.14}$$

Finally, under the No Bunching assumption, observe that for any fixed (B, r) pair, there exists $\kappa(\theta; B, r)$ such that

$$V(h(B, r, \theta); B, r, \theta) = \kappa(\theta; B) u(B)$$

for some function $\kappa(\theta)$. Put otherwise, because labor supply is fixed at $\overline{h}(\theta)$, r has no effect on utility independently of the benefit B and the type θ. Moreover, fixing B, we can always rescale utility for each type by multiplying by a real number $\kappa(\theta; B)$.

We assume that the function $\kappa(\theta)$ holds locally for all B in a neighborhood of the solution, for all types that are eligible for the benefit:

Assumption S2: Multiplicative Separability. Suppose $V(B, r, \theta) = \kappa(\theta) u(B)$ for all B in a neighborhood of B^* and for all $\theta \leq \tilde{\theta}$.

The Multiplicative Separability assumption states that utility gains from the benefit can be multiplicatively rescaled by the schedule $\kappa(\theta)$. Note that this assumption always holds if utility is homogeneous across types and all types have the same outside option; in that case, $\kappa(\theta) = 1$ for all θ. In the body of the paper, we start directly from that more demanding homogeneity assumption.

For other utility functions, the assumption holds as long as slight changes to the benefit around the optimum do not change the relative differences in the net utility that the different types experience from receiving the benefit. These relative differences are parameterized by the $\kappa(\theta)$ schedule, which must be invariant around the optimum. This assumption fails if, e.g., high types' marginal utility from receiving B diminishes at a faster rate than low types' marginal utility even in a neighborhood around the optimum.

The Multiplicative Separability assumption permits us to rescale differences in net utility with a (re-written) λ_θ welfare weight schedule.

Define $\tilde{\lambda}_\theta := \lambda_\theta \kappa(\theta)$. Moreover, let

$$\tilde{\lambda}_{\text{avg}} = \frac{\int_0^{\tilde{\theta}} \lambda_\theta \kappa(\theta) f(\theta) d\theta}{F(\tilde{\theta})}.$$

Intuitively, these $\tilde{\lambda}_\theta$ weights capture both: (i) the differences in the planner's value for one util given to each type (parameterized via the λ_θ weights), and (ii) the differences in utility each type experiences when given B in benefits (parameterized via the $\kappa(\theta)$ schedule).

Then, working from Equation (G.14), applying the Multiplicative Separability assumption, dividing by p and using that $F(\tilde{\theta}) = m$, we obtain:

$$\frac{u(B)}{u'(B)B} = \frac{1 + \eta_m}{\frac{\tilde{\lambda}_m}{\tilde{\lambda}_{\text{avg}}} + \eta_m} \tag{G.15}$$

for $\eta_m := \frac{\partial p}{\partial m}\frac{m}{p}$, which is the equation we target.

G.2 Additional Discussion of Equation (6)

We begin the additional discussion stating the following lemma, proven in Appendix F.

Lemma. *If $\rho \geq 1$, $\frac{\partial}{\partial B}\left(\frac{u(B)/B}{u'(B)}\right) > 0$.*

Lemma 2 follows from elementary properties of concavity. It establishes that the LHS of Proposition 1, the ratio of the average utility to the marginal utility, is increasing in $u(B)$. Henceforth we assume $\rho \geq 1$. It is useful because it allows us to determine how the planner adjusts B and m if the LHS and RHS are not equated.

To build additional intuition for Equation (6), we consider two sub-cases of Case 1 (without inframarginal effects), i.e. where $\eta_m = 0$.

Case 1a: no stigma, complete take-up. Assume there are no costs ($\gamma = 0$) and there is perfect take-up ($\eta_B = 0$). Rearranging Proposition 1 and applying Proposition 2 gives that at an optimum,

$$1 + \frac{1}{2}\rho \approx \frac{\lambda_{\text{avg}}}{\lambda_m}.$$

The LHS of this expression is the welfare gain from transferring an additional dollar to inframarginal types. The RHS of this expression (which always weakly exceeds 1, since $\lambda_{\text{avg}} \geq \lambda_m$ as λ_θ is decreasing in θ) is the welfare-weight-adjusted cost of taking a dollar away from type λ_m to transfer to inframarginal types. A small increase in m gives $u(B)$ (valued at $u(B)/B$ per dollar) in benefits to people who have λ_m in welfare weight. A small increase in B gives $u'(B)$ to people who have λ_{avg} in welfare weights. Proposition 1 establishes that at an optimum, the planner is indifferent between: (i) relaxing the eligibility criterion by increasing m (and transferring to new people, but reducing the benefit to the inframarginal types), and (ii) transferring a bit more by increasing B (giving $u'(B)$ to people with weights λ_{avg}). This tradeoff is at the core of many public discussions of social welfare programs.

Case 1b: incorporating costs. Now consider the case where benefit size affects take-up probability ($\eta_B > 0$) because there are costs ($\gamma > 0$). Observe that $\eta_B > 0$ tends to reduce the RHS. Intuitively, if $\eta_B > 0$, the planner must consider that raising the benefit for inframarginal types will boost take-up. People who newly take up the benefit are just indifferent to doing so, by an envelope condition, but they have a fiscal externality. For large η_B, the planner raises B. However, if γ is large, that serves as a force against raising B: large γ implies that most of the additional gain from take-up is soaked up by costs.

G.3 Discussion of Assumption 1

G.3.1 Necessary Condition for Proposition 3

Assumption 1 states that the change in the eligibility threshold reduces the average stigma costs among the fraction of people who take up the program. Assumption 1 is difficult to validate empirically without granular information on the treatment effect of changing the eligibility threshold on people's perceived stigma cost at every part of the stigma cost distribution.

First we show that this assumption is sufficient but not necessary. Equation (G.28) from the proof of Proposition 3 gives that the necessary and sufficient condition is:

$$\begin{aligned}&\left(\frac{1-s}{p^s}\left(\eta_m^i + 1\right) + \frac{s}{p^i}\left(\eta_m^s + 1\right)\right)\left(\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right)\\ &< \left(\frac{1-s}{p^s} + \frac{s}{p^i}\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s + \frac{1-s}{p^s}\eta_m^i - \frac{sm}{p^i}\frac{\partial\gamma}{\partial m} + \frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right).\end{aligned} \tag{G.16}$$

As long as Equation (G.16) holds, it is true that for all Ξ, $m^w > m^n$. Put another way, Equation (G.16) is a necessary condition that encodes the combination of Assumption 1 and either condition (i) or condition (ii) in Proposition 3. Thus, Equation (G.16) is weaker than Assumption 1 and condition (i) or condition (ii).

Equation (G.16) encodes the observation that as $s \to 0$, Proposition 3 always holds, because the necessary condition then reduces to:

$$\left(\eta_m^i + 1\right)\left(\frac{\lambda_m}{\lambda_{\text{avg}}}\right) < \left(\eta_m + \frac{\lambda_m}{\lambda_{\text{avg}}}\right), \tag{G.17}$$

which always holds since $\frac{\lambda_m}{\lambda_{\text{avg}}} \leq 1$. Intuitively, because information-only types capture the full benefit of the program, the fully naïve planner undervalues the social value of raising m more with more information-only types. As a result, she can tolerate a larger violation of $\frac{\partial \gamma}{\partial m} > 0$.

We also note that for various configurations of $\lambda_m/\lambda_{\text{avg}}$, γ, and $\frac{\partial \gamma}{\partial m}$, as well as the other parameters, the necessary condition may hold. For instance, as $\gamma \to 0$, the necessary condition always holds. Intuitively, as stigma agents become more like information agents, we no longer need a separate condition governing the behavior of η_m^i and η_m^s.

G.3.2 Discussion of Assumption 1

How could Assumption 1 fail? Suppose there are no information-only types ($s = 1$). Suppose moving the eligibility threshold reduces costs for people who are just indifferent to taking up the program (i.e., for whom $c \approx u(B)$). Suppose it has no effect on people for whom $c < u(B)$. Then, changing the eligibility threshold will first-order stochastically reduce the cost distribution. However, $\frac{\partial \gamma}{\partial m}$ will perhaps counterintuitively *rise*. Intuitively, the average cost conditional on taking up the program will feature a larger density at $c \approx u(B)$.

For a concrete example of the assumption failing, suppose $c \sim H_{\text{pre}} = N(1, \sigma)$ for known σ. Suppose $u(B) = 2$. Suppose raising m changes moves all costs larger than 2 to be at 2:

$$H_{\text{post}} = \begin{cases} N(1,\sigma) & c \leq 2 \\ c = 2 & \text{otherwise} \end{cases} \tag{G.18}$$

where we denote this truncated distribution by H_{post}. Changing the eligibility threshold induces a first-order stochastic reduction in the cost distribution. It raises the share of people in the population who take up the program. However, it also raises the average cost conditional on taking up the program. The average cost before raising m is $E[c|c \leq 2, H_{\text{pre}}] \approx 0.71$, whereas the average cost after raising m is $E[c|c \leq 2, H_{\text{post}}] \approx 0.91$.

However, violations of Assumption 1 are unlikely in practice. To see why, note that the counter-example above requires a large change in the cost distribution *only* for draws of the cost distribution that are about as large as $u(B)$. If raising m also affects the draws of the cost distribution for $c < u(B)$, that serves as a force pushing $\frac{\partial \gamma}{\partial m}$ downward.[65]

Second, Equation (G.16) shows that the necessary and sufficient condition for Proposition 3 to fail is much weaker than $\frac{\partial \gamma}{\partial m} \leq 0$.

Third, the specific counter-example changed the *shape* of the cost distribution. H_{pre} is normal; H_{post} is a truncated normal. We develop propositions showing that for the normal and exponential cost distributions, any any change in the (unconditional) mean costs that maintains the distributional family from which the costs are drawn will feature $\frac{\partial \gamma}{\partial m} < 0$.

Proposition 4. *Let $c \sim N(\mu(m), \sigma)$ with $\mu'(m) < 0$. Then $\frac{\partial \gamma}{\partial m} < 0$.*

We prove Proposition 4 in Appendix F. A change in m reduces the mean (unconditional) cost but the cost distribution remains normal. Then the change in ratio of costs to benefits, conditional on taking up the program, will shrink in m; i.e., Assumption 1 holds. We develop a similar proposition if costs are exponentially distributed:

[65] Note that we suppose all types receive draws from the same cost distribution. Thus, the violation of Assumption 1 is *not* that changing m only affects costs for $\theta = m$ at the marginal of eligibility. Rather, Assumption 1 is only likely to fail if changing m affects people for whom $c \approx u(B)$, i.e. they are indifferent to signing up (regardless of their income).

Proposition 5. *Let $c \sim Exp(\theta(m))$, where $1/\theta$ is the mean of the exponential distribution Exp and $\theta'(m) > 0$. Then $\frac{\partial \gamma}{\partial m} < 0$.*

Note that $\theta'(m) > 0$ implies the average unconditional cost $1/\theta$ falls in m, so Proposition 5 is qualitatively similar to Proposition 4.

G.4 Formal Statement of Proposition 3.

Fix a vector of parameters $\Xi = (p^i, p^s, \lambda_\theta, s, \gamma, \eta_B, u(\cdot))$. Notice that, for any Ξ, any given η^i_m, η^s_m, and function $\frac{\partial \gamma}{\partial m}(B, m)$ induce a pair $(B^*(\eta^i_m, \eta^s_m, \frac{\partial \gamma}{\partial m}), m^*(\eta^i_m, \eta^s_m, \frac{\partial \gamma}{\partial m}))$ that satisfies Equation (8).

We call $(B^n, m^n) := (B^*(0,0,0), m^*(0,0,0))$ the *naïve* choice of (B, m): this is the choice of eligibility threshold and benefit size if (i) the planner neglects inframarginal effects arising from either agent, and (ii) does not realize that the eligibility threshold affects stigma. Call $(B^w, m^w) := (B^*(\eta^i_m, \eta^s_m, \frac{\partial \gamma}{\partial m}), m^*(\eta^i_m, \eta^s_m, \frac{\partial \gamma}{\partial m}))$ the *sophisticated* choice of (B, m).

We make several assumptions to rule out edge cases. First we assume, (i) $\lambda_m/\lambda_{\text{avg}} < 1$ at the naïve solution. This implies there exists some point up to the naïve planner's choice of m at which the welfare weight schedule is strictly decereasing. We require (ii) $\gamma > 0$, that stigma costs are positive for stigma agents (if they exist). We also require (iii) $\rho \geq 1$.

Finally, as discussed in the body and this appendix, we impose that (iv) Assumption 1 holds.

Under these assumptions we can show the following:

Proposition 6 (Formal statement of Proposition 3)**.** *If $\eta^i_m > 0$ or $\eta^s_m > 0$, then $m^w > m^n$ for all Ξ as long as $\eta^s_m \leq \eta^i_m$ (condition (i)). Moreover, there exists $\varepsilon > 0$ such that $m^w > m^n$ for all Ξ as long as $\eta^s_m \in [\eta^i_m, \eta^i_m + \varepsilon]$ (condition (ii)).*

The proof is in Appendix G.5. This version of Proposition 3 is slightly more general than the version stated in the body. As in the body, one hypotheses that delivers the sharp policy implication is if stigma agents are less elastic than information agents (condition (i)). However, we also add a second condition: each vector Ξ yields an interval $\eta^s_m \in [\eta^i_m, \eta^i_m + \varepsilon]$ for $\varepsilon > 0$ in which the proposition still holds (condition (ii)). The utility of having condition (ii) as an alternative is that then the statement holds for some $\eta^s_m > \eta^i_m$ for all parameterizations. These conditions are also sufficient but not necessary.

G.5 Proofs in Extensions

G.5.1 Proof of Proposition 4

Proof. It suffices to prove that $\frac{\partial}{\partial m}\left(E[c|c \leq u(B), \mu(m)]\right) < 0$. First, let $\chi(Z) := \phi(Z)/\Phi(Z)$ for normal PDF ϕ and normal CDF Φ. Equation (3) in Sampford (1953) gives that

$$0 < \frac{\partial}{\partial Z}\left(\frac{\phi(Z)}{1 - \Phi(Z)}\right) < 1 \tag{G.19}$$

for all Z. Thus

$$-1 < \frac{\partial \chi}{\partial Z} < 0 \tag{G.20}$$

since the normal PDF is even and $1 - \Phi(Z) = \Phi(-Z)$.

The usual properties of the normal distribution give:

$$E[c|c \leq u(B), \mu(m)] = \mu(m) - \sigma \frac{\phi(Z(m))}{\Phi(Z(m))} \tag{G.21}$$

for $Z(m) := (u(B) - \mu(m))/\sigma$.

The chain rule gives

$$\frac{\partial}{\partial m}\left(\frac{\phi(Z(m))}{\Phi(Z(m))}\right) = -\frac{\partial \chi}{\partial Z}\frac{\mu'(m)}{\sigma}. \tag{G.22}$$

Then evaluating Equation (G.21) at the bounds in Equation (G.20) gives

$$\mu'(m) < \frac{\partial}{\partial m}\left(E[c|c < u(B), \mu(m)]\right) < 0. \tag{G.23}$$

□

G.5.2 Proof of Proposition 5

Proof. It suffices to prove that $\frac{\partial}{\partial m}\left(E[c|c < u(B), \mu(m)]\right) < 0$. The mean of the truncated exponential distribution is:

$$\mu(\theta(m)) = \frac{1}{\theta} - u(B)\left(\exp(\theta u(B)) - 1\right)^{-1} \tag{G.24}$$

for $u(B) > 0$. This function is monotonically decreasing for all $u(B)$ (Al-Athari, 2008). □

G.6 Proof of Proposition 3/Proposition 6

Proof. We want to show that the naïve planner would raise the eligibility threshold m and lower the benefit size B. First, we note that for a given (B, m) pair, the budget constraint ensures that raising B requires lowering m, and raising m requires lowering B. Thus, it is sufficient to argue that the naïve planner sets B too high. Consider the following rearrangement of Equation (F.13):

$$\frac{u(B)}{Bu'(B)} = \frac{\frac{(1-s)}{p^s}\left(\eta_m^i + 1\right) + \frac{s}{p^i}\left(\eta_m^s + 1\right)}{\left(\frac{p^s}{sp^s+(1-s)p^i}s\eta_B + 1\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s + \frac{1-s}{p^s}\eta_m^i - \frac{sm}{p^i}\frac{\partial\gamma}{\partial m} + \frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right)}. \tag{G.25}$$

Using Lemma 2, we note that the LHS of Equation (G.25) is increasing in B. Thus, noting that the naïve planner solves Equation (G.25) for B^n and the sophisticated planner solves the equation for B^s, we want to show:

$$\frac{u(B^s)}{B^s u'(B^s)} - \frac{u(B^n)}{B^n u'(B^n)} < 0, \tag{G.26}$$

and substituting Equation (G.25), we have that we want to show:

$$\begin{aligned}&\frac{\frac{(1-s)}{p^s}\left(\eta_m^i + 1\right) + \frac{s}{p^i}\left(\eta_m^s + 1\right)}{\left(\frac{p^s}{sp^s+(1-s)p^i}s\eta_B + 1\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s + \frac{1-s}{p^s}\eta_m^i - \frac{sm}{p^i}\frac{\partial\gamma}{\partial m} + \frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right)} \\ &- \frac{\frac{(1-s)}{p^s}(0+1) + \frac{s}{p^i}(0+1)}{\left(\frac{p^s}{sp^s+(1-s)p^i}s\eta_B + 1\right)\left(\frac{(1-\gamma)s}{p^i}0 + \frac{1-s}{p^s}0 - \frac{sm}{p^i}0 + \frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right)} < 0.\end{aligned} \tag{G.27}$$

Observe that

$$\left(\frac{p^s}{sp^s + (1-s)p^i}s\eta_B + 1\right) > 0.$$

Cross-multiplying, it is then sufficient to show:

$$\begin{aligned}&\left(\frac{1-s}{p^s}\left(\eta_m^i + 1\right) + \frac{s}{p^i}\left(\eta_m^s + 1\right)\right)\left(\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right) \\ &< \left(\frac{1-s}{p^s} + \frac{s}{p^i}\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s + \frac{1-s}{p^s}\eta_m^i - \frac{sm}{p^i}\frac{\partial\gamma}{\partial m} + \frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i} + \frac{1-s}{p^s}\right)\right).\end{aligned} \tag{G.28}$$

Under Assumption 1, $-\frac{sm}{p^i}\frac{\partial \gamma}{\partial m} \geq 0$, so it is sufficient to show that:

$$\begin{aligned}&\left(\frac{1-s}{p^s}\left(\eta_m^i+1\right)+\frac{s}{p^i}\left(\eta_m^s+1\right)\right)\left(\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i}+\frac{1-s}{p^s}\right)\right)\\&<\left(\frac{1-s}{p^s}+\frac{s}{p^i}\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s+\frac{1-s}{p^s}\eta_m^i+\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i}+\frac{1-s}{p^s}\right)\right).\end{aligned} \tag{G.29}$$

Rearranging gives that this condition is equivalent to:

$$\left(\frac{1-s}{p^s}\eta_m^i+\frac{s}{p^i}\eta_m^s\right)\left(\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i}+\frac{1-s}{p^s}\right)\right)<\left(\frac{1-s}{p^s}+\frac{s}{p^i}\right)\left(\frac{(1-\gamma)s}{p^i}\eta_m^s+\frac{1-s}{p^s}\eta_m^i\right) \tag{G.30}$$

$$\iff \frac{\frac{1-s}{p^s}\eta_m^i+\frac{s}{p^i}\eta_m^s}{\frac{(1-\gamma)s}{p^i}\eta_m^s+\frac{1-s}{p^s}\eta_m^i}<\frac{\frac{1-s}{p^s}+\frac{s}{p^i}}{\frac{\lambda_m}{\lambda_{\text{avg}}}\left(\frac{(1-\gamma)s}{p^i}+\frac{1-s}{p^s}\right)}. \tag{G.31}$$

The statement holds strictly if $\eta_m^s=\eta_m^i$ as long as $\lambda_m<\lambda_{\text{avg}}$, by factoring the LHS and canceling. Moreover, holding η_m^s fixed, the LHS is strictly decreasing in η_m^i. Thus, if the statement holds for $\eta_m^s=\eta_m^i$, it also holds for $\eta_m^i<\eta_m^s$. This shows that the desired statement holds under condition (i). To argue that the desired statement holds under condition (ii), notice that the LHS is strictly increasing in η_m^s, holding η_m^i fixed. As a result, there exists $\tilde{\eta}_m^s>\eta_m^i$ such that the statement holds with equality. Since the LHS is increasing in η_m^s, the statement holds strictly for $\eta_m^s<\tilde{\eta}_m^s$. Thus, there exists an interval $\eta_m^s\in[\eta_m^i,\eta_m^i+\varepsilon]$ for $\varepsilon>0$ such that Equation (G.31) holds, which completes the proof. $\square$