Benefits of Design: The Impact of Shifting Administrative Burdens from Citizens to their Bank

Olivia Bergman[*], and Michael Hiscox[†]

Working Paper Draft: March 14, 2022[‡]

Abstract

Can simplifying the experience of applying for government benefits change attitudes and take-up? In two large field experiments with an Australian bank, we use design, data and digital infrastructure to shift administrative burdens away from citizens. First, an experiment among 195,414 low-income people shows simplification of a means-tested benefit substantially increases satisfaction with both the bank and government. Transactions and linkages to government records show take-up effects are more modest, suggesting redesigning how citizens experience policies can reshape attitudes toward government, independently of take-up effects. Second, in a randomized national roll-out to 7,720,009 people, we show this simplification scales successfully and cost-effectively. Broadly, we argue that besides asking how policy content affects attitudes—the "who gets what"—examining "when" and especially "how" citizens experience policies helps explain interactions with, and attitudes toward, government. These experiments also illustrate how private institutions can play a role in such policy feedback loops.

[*]Postdoctoral Fellow, University of Chicago Booth School of Business; Corresponding Author (olivia.bergman@chicagobooth.edu)

[†]Clarence Dillon Professor of International Affairs, Department of Government, Harvard University

[‡]We gratefully acknowledge the Commonwealth Bank of Australia for their phenomenal partnership in executing these experiments start to finish, and the NSW Department of Planning, Industry, and Environment for sharing additional data. See final page for detailed acknowledgements. We also thank Harvard University's Sustainability, Transparency, and Accountability Research Lab for support. AEA pre-registration: www.socialscienceregistry.org/trials/3083/history/30864. IRB: Harvard CUHS, 04/16/2018, IRB18-0546.

1 Introduction

When we think of politics, we might picture heated legislative debates or campaign signs in yards. But every year, the average citizen is more likely to pay taxes, wait in line at the post office, or fill out a government form. Much research focuses on formal political processes, but in most if not all countries, participation in dramatic political events is less frequent than mundane experiences with policies.

In this paper, we examine how one such everyday policy experience—applying for a government benefit—shapes views of government. Together with the Commonwealth Bank of Australia (CBA), we implement two field experiments to test what happens when frictions are removed from benefit application experiences.

In the first experiment with 195,414 bank customers, we focus on the Low Income Household Rebate, provided by the state government of New South Wales to help reduce electricity costs for low-income households. We redesign and simplify the application experience with the help of the bank's data and widely used mobile app.[1] Concretely, we use bank data to identify people who are likely eligible, contact them about the $285 government benefit, and randomly assign them to one of two application experiences. In the 'business-as-usual' group, people get directed to a standard application page, modeled after the existing government page: the process looks complicated and has several steps where citizens might "drop off". In the treatment group, people get directed to a simplified page, which we have designed to have clearer formatting and a more straightforward application process. The goal of the treatment is to reduce administrative burden: the effort required to find out about a benefit, determine one's eligibility, and apply (Herd and Moynihan 2018). Since we cannot change or eliminate any of the formal application rules, our goal is to shift as much of the administrative burden as possible away from citizens and onto the bank.

In contrast to economic evaluations of efforts to reduce administrative burden, we consider not only financial outcomes but also attitudinal consquences. Administrative burden is a common feature of interactions between citizens and the state, and it is often viewed as the result of a technical policy trade-off: between the goal of maximizing program take-up (i.e., making sure the right people get it) and program integrity (i.e., making sure the wrong people do not get it). But administrative burden can also be the result of differing political views, ultimately encoded in implementation design rather than policy content. Since administrative burden is a product of political processes, it could also have political

[1] At least 40 percent of the customer base uses the bank's mobile banking application. Smartphone penetration in Australia is high at 89 percent, including among older bank customers (Deloitte 2018). Conducting the experiment via the mobile app makes it possible for us to communicate with people in their daily life, and to experimentally alter content and track customer activity.

consequences. In addition to any gatekeeping function, administrative burden is experienced in ways that might have effects that are distinct from take-up. We hypothesize that changing administrative burden can also change important attitudes.

We find that intervening to reduce administrative burden changes a number of behaviors and attitudes. We first show that simplifying the application page results in 3 times higher likelihood of initiating the process of claiming this government benefit (14.7 percent versus 4.8 percent on the standard page). Then, using detailed bank data and linkage to government records, we find that the simplified treatment increases take-up by 0.8 percentage points (or 2.1 percent) relative to the standard government-like page, and 0.9 percentage points (or 2.7 percent) relative to the pure control group. While these effects could seem small in absolute terms, the intervention is highly cost-effective: we estimate that each dollar the bank spent on the build generates at least $106 (yearly) for customers.

Finally, though effects on take-up are modest, reducing burdens substantially increases satisfaction with all involved parties. Those with a simplified experience were 45.5 and 44.4 percent more likely to submit surveys rating the bank and government positively, respectively. In the sample of survey respondents, we present suggestive evidence that the simplified design generates satisfaction with the bank and government over and above self-reported claiming.

Taken together, the first experiment provides new empirical evidence that the design of the policy experience—beyond content, and without large financial effects—can shape views of government. As such, it complements the approach and insightful analyses of policy feedback scholars before us (Schneider and Ingram 1993; Soss 1999; Mettler and Soss 2004), by advancing the theory that we can change the nature of the citizen experience, in a way that affects political attitudes, using the tools of policy design.[2]

In a national scale-up, with 7,720,009 customers, we evaluate the decision to build on the first findings and create a permanent digital platform with over 200 simplified benefits. For a year, platform invites get randomly assigned. Focusing on the same type of benefit studied in the first experiment, we find that getting invited to the platform lowers average monthly energy bills by $1.32, and more broadly that the platform generates simplified government experiences for over 2 million customers over the course of the first 14 months.

It is not insignificant that these interventions took place with the help, data infrastructure, and digital platform of a large bank. This provides intriguing evidence that private actors can help reduce frictions embedded in citizen-state interactions, in a way that is consequential for take-up and political attitudes, particularly among vulnerable, often hard-to-

[2]Here we group several factors together—which different sub-fields categorize as belonging to design, implementation and administration—under the umbrella term 'policy design', here defined as the factors determining how the policy ultimately looks and feels to citizens interacting with it.

reach groups. Importantly, burdens were reduced without compromising program integrity: the policy content, including eligibility criteria and application rules, remained the same. Government can draw on such evidence when considering new ways of delivering services, and in identifying stakeholders for public-private partnerships to shift burdens. Finally, we hope this illustrates the value of bridging scholarship on public policy, public administration and political behavior, for understanding how citizens form views underpinning democracy.

The paper proceeds as follows: in section 2, we review relevant theories and concepts from the literature on the costs of applying for government benefits, the relationship between policy experiences and political attitudes, and which actors can shift administrative burden. Section 3 introduces the study setting and provides background on the Australian low-income household rebate. Section 4 details the empirical design and implementation of the first field experiment, and Section 5 presents the first set of results. Section 6 presents the design and results of the scale-up experiment. The last section identifies limitations and future avenues for research, and concludes with a discussion of the potential perils and promises of involving the private sector in facilitating citizen-state interactions.

2 Theory and Literature

2.1 The Costs of Applying for Government Benefits

Most countries provide some social safety net: a system of government benefits to help citizens in need pay for food, housing, health care, and other basic living expenses. These benefit programs can be designed in different ways; we observe lots of variation in the way social benefits are delivered across the world, within countries, and over time (Currie 2004).

Such design choices are consequential, both for citizens and administrators. Unless government benefit programs are universal, eligibility has to be verified. Even universal programs, unless they have automatic enrollment, have to be applied for. This generates some form of administrative burden—defined as the effort required of citizens to find out about a benefit, figure out if they are eligible, how to apply, and then manage to stay enrolled (Herd and Moynihan 2018)—in almost all government benefit programs.

Many administrative burdens exist in service of program integrity: the tasks were created at least in part to ensure that only individuals who truly qualify receive benefits (Nichols and Zeckhauser 1982; Sunstein 2018). But the complexity and frequency of such tasks can mean that hurdles intended to exclude the ineligible often exclude many eligibles who happen to be forgetful, disorganized, or overwhelmed (Mullainathan and Shafir 2013). As such, the choice can be cast as one between designing the program in a way that results in (a) some ineligible people receiving a benefit or (b) some eligible people not receiving a benefit (Sunstein 2018). This is of course a trade-off: if enrollment by ineligibles is high, then government revenues are diverted from more productive uses. If enrollment among eligible individuals is low, then the program is falling short of reaching its goal of helping the target group (Currie 2004).

We know burdens can matter for take-up of government benefit programs.[3] Adding burdens that increase costs of applying for benefits—whether in terms of time, effort, or money—has been shown to decrease program participation. For example, closing local Social Security field offices, where applications can be submitted, reduced take-up of Disability and Supplemental Security Income (Deshpande and Li 2019); similarly, openings and closings of Women, Infants and Children local program offices affected take-up (Rossin-Slater 2013); and even minimal monthly premiums for Medicaid reduces enrollment (Dague 2014).

There is also evidence of the opposite, that removing administrative burdens can increase take-up. For example, personalized help with student aid applications increased college attendance (Bettinger et al. 2012); application assistance for seniors increased Supplemental Nutrition Assistance Program (SNAP) enrollment (Finkelstein and Notowidigdo 2019), and reducing the informational complexity of forms increased claiming of the Earned Income Tax Credit (EITC) (Bhargava and Manoli 2015).

But do burdens affect people differently? In addition to viewing burdens as a way to exclude ineligibles, there is a long economics tradition of viewing burdens as a way to target public resources. The logic is that burdens help deter those who are only marginally needy from applying, while still allowing those truly disadvantaged to apply and get the benefits they need (Nichols and Zeckhauser 1982; Besley and Coate 1992; Alatas et al. 2016). However, this idea that burdens are a true test of need—that the more somebody needs help, the more effort they will exert to get it—does not always bear out empirically. In fact, research in behavioral economics has shown that burdens may have exactly the opposite targeting effect to what these classical theories of ordeal mechanisms assume, and thus discourage the neediest applicants (Bertrand et al. 2004; Mani et al. 2013; Bhargava and Manoli 2015).

There are other distributive aspects of administrative burdens as well. In general, higher-income people tend to have less contact with the types of means-tested government programs that impose the most burdens on beneficiaries. In the US, for example, the types of programs that benefit higher-income people tend to require less paperwork (Badger and Sanger-Katz 2020). Some argue that when beneficiaries lack political power, or are plagued by negative stereotypes, or both, programs tend to be designed in ways that create more burden and fewer benefits (Schneider and Ingram 1993). When high-income people do encounter burdens, they can often pay someone else (e.g., a lawyer, accountant, or software) to handle these hassles.

[3] Take-up is defined as the proportion of eligible people who access a program.

The net effect is that the unequal impact of burdens can compound: not only do low-income people encounter more burdens, and lack the means to outsource them; poverty can also exact a kind of cognitive tax, which decreases mental bandwidth and can make it hard to deal with precisely these kinds of tasks (Mullainathan and Shafir 2013).

Lastly, the overt goal of protecting program integrity can provide political cover: burdensome tasks required to access programs can be used as a deliberate political tool to halt participation, even among needy eligibles. Partisan gridlock in US Congress has made administrative burdens an increasingly critical policy tool to decrease the size of safety net programs without going through the legislature (Herd and Moynihan 2018).

In sum, there is a body of research suggesting that program participation rates vary at least partly based on how much work it takes to get and stay enrolled, and that these administrative burdens can be addressed to change take-up of benefits. But beyond take-up, we argue that administrative burdens are likely to affect the nature of the citizen experience and their ensuing attitudes, which can both be politically consequential. As Mettler and Soss point out, 'program evaluation' is often left to social scientists in fields other than political science and "policy analysts routinely examine the social and economic consequences of [government] programs, yet their political effects continue to be widely ignored" (2004: 59). Existing evaluations tend to not take into account the attitudinal externalities of administrative burdens. We turn to these next.

2.2 Policy Experiences and Political Attitudes

How are citizens' attitudes toward government influenced by their encounters with specific government programs? A large literature on policy feedback establishes links between policy experiences and political attitudes (Schneider and Ingram 1993; Soss 1999; Mettler and Soss 2004). The central argument is that policies serve as sources of information and meaning, and that features of policy design structure program experiences in ways that teach alternative lessons about the nature of government. "Citizens encounter and internalize the messages not only through observation of politics and media coverage but also through their direct, personal experiences with public policy" (Schneider and Ingram 1993, 340).

The conventional way of testing whether experiencing a policy changes political attitudes has been to "compare how experiences and interpretations vary across users of a small number of distinct programs, or between a single program's beneficiaries and a parallel group of non-beneficiaries" (Mettler and Soss 2004, 64). For example, a study of welfare programs in the 1990s found that beneficiaries interpret their experiences with welfare bureaucracies as evidence of how government works more generally (Soss 1999, 363). In interviews, Social

Security Disability Insurance (SSDI) beneficiaries described a program with that was ultimately responsive to their demands, whereas beneficiaries of Aid to Families with Dependent Children (AFDC) faced more frequent and intrusive interactions, and took away negative lessons about government's responsiveness to people like them (Soss 1999). Another study compared beneficiaries of the EITC, beneficiaries of either AFDC or Temporary Assistance to Needy Families (TANF), and non-beneficiaries of either policy, concluding that "reactions to public policies emanate not just from the amount of resources received but also from the experiences associated with applying for benefits and receiving them" (Shanks-Booth and Mettler 2019, 310).

Whereas all these scholars find that beneficiaries of different programs take away quite different messages about the government, they also agree that a major challenge is disentangling the effects of policy design or program experiences from other factors. The relationship between policy usage and political attitudes is complex, and without exogenous variation it is hard to ascertain whether differences arise because of differences in policy experiences, or because of pre-existing characteristics of beneficiaries of different programs. Pre-existing characteristics might influence both which programs people apply for and their subsequent attitudes. (See Lerman (2019) for evidence that partisanship and pre-existing beliefs affect both policy take-up and views of privatization.) One reason it has been so hard to isolate the effects of policy design on attitudes has been a dearth of data.[4] Data that would permit us to isolate the effects of policy design "remain extremely rare" (Mettler 2018, 24).

A recent survey experiment provides evidence that the perceived burdens —how hard it looks from the outside to apply and get enrolled in a program—might shape views of the program among non-beneficiaries (Keiser and Miller 2020).[5]In contrast, we examine how experiencing these burdens can impact general government attitudes among the large target populations of such programs.

[4]There are few (American) surveys that ask about both policy experiences and political attitudes, and no suitable panel data (Mettler 2018).

[5]Means-testing may be socially divisive to the extent that it divides society into those who give and those who receive. But this survey experiment found that Republicans were more supportive of the TANF program and its beneficiaries when exposed to information about how hard it is to get enrolled (Keiser and Miller 2020). This result might stem from beliefs that burdens are a true test of need, or the related view that burdens are a test of personal responsibility. In that sense, the difficulty of accessing programs could partly be the point for some voters and politicians (Badger and Sanger-Katz 2020), which suggests a potential tension between increasing support among (i) beneficiaries and (ii) non-beneficiaries, of government policies.

2.3 Who Can Shift Administrative Burden?

If we want to explore shifting administrative burdens away from citizens, there are three broad alternatives: the state, non-profits, or for-profits.[6]

The state can take on some or all of the burdens, for example through outreach, simplifying forms, setting up online applications, allowing for presumptive eligibility, or auto-enrollment. A prime US example is Social Security, which is administratively very complex, but almost all tasks embedded in the program are completed by government administrators.[7] In Sweden, the state pre-fills tax returns for all citizens, using information it collects from third parties. For US Medicaid, people can often enroll based on eligibility for other means-tested programs. Thus, administrators can choose to identify and retrieve information from other government databases, rather than making citizens gather and (re)submit it.

Non-state actors can also take on administrative burdens. Within the non-profit category, some organizations were started specifically to help people navigate the bureaucracy involved in applying for benefits. Others, such as churches and labor unions have another core mission, but often help their community navigate the social safety net. Non-profit hospitals and care organizations also play a role in the US: because they receive federal funds, they have to give care to patients who cannot pay. This creates incentives to help patients enroll in public health insurance programs, to reduce the cost of charity care (Herd and Moynihan 2018).[8]

But, whereas non-profits are often recommended as partners for bureaucrats to collaborate with on public outreach and take-up efforts, for-profits rarely seem to be. We can think of at least three concerns with encouraging for-profits to help citizens navigate burdens, which are worth considering in turn. First, for-profits could have a vested interest in keeping things complex. In the US, most are familiar with tax preparation companies, which help citizens navigate the complex web of tax filing, for a fee. Indeed, the industry has succesfully lobbied to stop the government from offering its own free software (or simply pre-filling forms), in exchange for promising to provide this free service to low- and middle-income Americans themselves (Reid 2017). Second, if for-profits are encouraged to reduce burdens, a concern is that they will find ways to make money off people. Fueling this concern, ProPublica recently revealed that the tax preparation industry has long misled taxpayers—who qualify to file for free—into paying for the service (Kiel and Elliott 2020).

[6] Let us assume that factors like program eligibility criteria are fixed, so that it is not an option to reduce or eliminate burden by changing such formal rules.

[7] Individuals do not need to collect and provide verification of their earnings over their lifetimes to get social security benefits. Enrolling is easy, and take-up is almost 100 percent (Herd et al. 2013).

[8] One testament to this is that low-income children and pregnant women are often referred to as 'conditionally covered', even when currently unenrolled in Medicaid, since we can expect them to be enrolled by the hospital once medical care is needed (Cutler and Gruber 1996).

A third concern with encouraging for-profits to facilitate take-up is that they might discriminate by only selectively helping people. For example, in the US, most Health Maintenance Organizations (HMOs) are for-profits. Their many points of contact with beneficiaries make them promising candidates for facilitating both enrollment and re-enrollment in Medicaid and Medicare Advantage programs. But, because HMOs receive a capitated payment for each beneficiary, there is concern they have strong incentives to engage in 'cream-skimming' and enroll only the healthiest people in their programs (Herd et al. 2013). As such, relying on for-profits to shift burdens might exacerbate the already unequal distribution of burdens.

Despite these concerns, any simplistic view that for-profits should stay out of citizen-state interactions, could result in missed opportunities to improve citizens' most frequent experiences of the social contract. [9] Shifting burdens away from citizens takes both political will and administrative capacity—including technology, data, expertise, comunication channels (Herd and Moynihan 2018)—and the state could lack either or both. The private sector might have the up-front resources to develop a proof-of-concept, which could shift political will, or demonstrate the cost-effectiveness of investing in administrative capacity. The government could draw on such evidence when considering new ways of delivering services, and in identifying stakeholders to form public-private partnerships with.

In the case of the Australian bank, it has the means, motive and opportunity to play such a role. First, it already has convenient channels of communication, knows a lot about its customers, and continuously gathers high quality data for its routine operations. Second, it cares about the financial health of its customers,[10] and naturally its own image, so the incentives are aligned with increasing take-up and creating experiences with less friction. Lastly, the bank already has the infrastucture and capacity to build, test and evaluate solutions. In the Discussion section, I will circle back to the topic of the role of private actors in shifting administrative burden, and discuss future causes for concern and optimism.

2.4 Shifting the Burden, Measuring Policy Feedback

In order to isolate the effects of policy design from potentially confounding factors—such as citizens with different characteristics self-selecting into different policies—we partner with the bank to design and implement an experiment that allows us to monitor how citizens'

[9] For instance, several US startups successfully use technology to improve experiences with government, such as one that enables recipients of food stamps (SNAP) to check their balance using an app, instead of making lengthy phone calls (Riley 2017). Consider the alternative: until recently, the official California food stamps (CalFresh) website was unusable on cellphones. Many governments lag behind the private sector in digital capacity, which disproportionately hurts the most needy: 10 percent of Americans only have internet access via smartphones (Pew 2019) and this is especially true for low-income people.

[10] Even in the most cynical of views, customers who have more money in their accounts have more money to repay their interest or debts, for example, so the bank is incentivized to help them secure financial support.

attitudes and take-up behavior shift in response to a randomly assigned experience of a government policy. This strategy complements the approach of policy feedback scholars before us. By gathering new data, we can test one central hypothesis in the literature: that an experience with a specific policy can shape more general political dispositions. Our main policy feedback hypothesis is that reducing administrative burden will lead to more positive attitudes toward government.

3 Setting and Background

3.1 The Bank

The Commonwealth Bank of Australia (henceforth CBA or 'the bank') is the largest retail bank in Australia, serves as the main financial institution for a third of Australians, and had more than 10 million retail customers at the time of the experiment. CBA was founded in 1911 by the Australian government, and thus has an unusual history for a commercial bank: for 50 years it was granted special authority to execute functions typically only bestowed upon central banks. CBA was fully privatized in 1996.

3.2 'Electricity Benefit': The Low Income Household Rebate

In this paper we focus on one means-tested government benefit—the Low Income Household Rebate, henceforth referred to as the electricity benefit—provided by the state government of New South Wales (NSW) to help reduce the cost of electricity for low-income residents. The modal and maximum electricity benefit received is $285 per year,[11] which on average covers around 18 percent of the the cost of beneficiaries' annual electricity bills (DPIE 2019). The other Australian states have similar discount programs, as cost of living is a salient political issue and energy costs have risen steadily since 2006.[12]

The electricity benefit is designed not as a direct payment, but as a discount applied to the household electricity bill.[13] This type of design is common in Australia, where many government benefits are means-tested and come in the form of either tax credits or discounts on services.[14] The average take-up rate of the electricity benefit is an estimated 72 percent,[15]

[11] Amounts refer to Australian dollars.

[12] For a small apartment, the combined cost for water, electricity, and gas is $220 per month. In 2017-18, the average annual electricity bill, for the population already receiving the benefit, was $1,695 (DPIE 2019).

[13] Granted on a daily basis (rather than a $285 lump sum): a $0.78 discount for each day included on an electricity bill. Most Australian companies bill quarterly, so the maximum discount each quarter is $71.76.

[14] In addition, delivery of benefits is often delegated to the non-government sector (AIHW 2017).

[15] An estimated 320,000 eligible people in NSW did not claim the electricity benefit in 2018 (DPIE 2019).

comparable to rates for means-tested US programs frequently targeted with take-up efforts: for example, take-up for SNAP was 85 percent in 2016 (CBPP 2019) and take-up of the EITC was 71–86 percent in 2018, depending on the state (IRS 2019). In NSW, hundreds of thousands low-income households are leaving the electricity benefit money on the table, despite reporting they would struggle to handle an unexpected expense.[16]

3.3 Eligibility and Application Process

At first glance, applying for the electricity benefit seems relatively easy. People need to call their electricity company and provide information to confirm eligibility: being at least 18 years old, living in the state of NSW, having an electricity account in their own name, and having a concession card—a card issued by the government to those with limited incomes (e.g. due to unemployment, old age, disability, sickness, Veteran status, disasters, and caring for family members). The estimated 28 percent of eligible households who have not taken up the benefit have already gone through arguably the hardest part of the process: getting a concession card.[17] But, for some reason, they have not claimed the electricity benefit.

Potential Barriers

There are many reasons why people might fail to claim government benefits they are entitled to. First, a lack of information—about the existence of the benefit, about their own eligibility, or how to apply. But even if people are aware and applying seems worthwhile, the presence of frictions, small and large, can trip up potential applicants before the finish line.

First, to apply for a benefit you have to know it exists. Even though various websites contain information about the electricity benefit, such as those of electricity companies and the different government agencies, 68 percent of our survey respondents were not aware that the benefit existed prior to our experiment (Appendix Figure C2).[18] Once you are aware of its existence, you have to believe you are eligible. Eligible people then need to figure out exactly how to apply, and decide that the effort of applying will likely be worth the trouble.

At first glance, this type of cost-benefit judgement seems straightforward for the electricity benefit: getting $285 anually should be worth becoming informed and calling one's

[16]Research commissioned by Budget Direct Home Insurance, 2019.

[17]Once you have a concession card, it can be used to get access to a range of benefits with similar eligibility criteria. As such, a concession card serves to eliminate redundant verifications of eligibility criteria, which can include income, assets, status of dependents, residency and age. Some concession cards can be applied for directly, others are automatically sent to recipients of certain "qualifying benefits" that make them eligible for a corresponding concession card (Department of Social Services, Australian Government, 2020).

[18]There are many studies corroborating that we cannot assume people are aware of the benefits available to them. For example, in the US, a survey of likely eligible SNAP non-participants found that about half were not aware of their eligibility (Bartlett et al. 2004). See Finkelstein and Notowidigdo (2019) for a review.

electricity company. But, even if applying is actually simple, people might perceive it as hard, for example if long wordy instructions make it seem complicated. We know that small barriers can have outsized impacts due to cognitive biases and lack of bandwidth, especially among those with lower incomes (Mullainathan and Shafir 2013). It is also well-documented that part of the observed gap between intention and action can be attributed to simple barriers, such as having to remember to mail a separate form or look up a phone number (Thaler and Sunstein 2021). Awareness, information, and incentives can help people form intentions, but behavior change might also depend on actually removing barriers and creating channels that make it easier to choose and act (Daminger et al. 2015).

4 Empirical Design and Data

In partnership with a large retail bank in Australia, we design potential solutions to help overcome barriers to take-up, and implement a randomized evaluation of the intervention to encourage eligible individuals to apply for the electricity benefit.

4.1 Design of the Intervention

The intervention uses the bank's data and digital infrastructure to (i) identify people who are likely eligible for the electricity benefit, (ii) contact them to provide this information, and (iii) create a more straightforward application experience, with fewer potential barriers to take-up. The goal is test the effects of shifting more of the administrative burden—the learning, compliance and psychological costs associated with claiming a government benefit (Herd and Moynihan 2018)—away from citizens.

Set-Up and Randomization

The intervention was live between May 4 and June 21, 2018. We randomized our study population of 195,414 individuals into nine equally-sized experimental arms: one pure control and eight treatment groups.[19] The treatment groups are a combination of four invitations and two application pages (4x2).[20] For our purposes, we can think of the randomization as happening in two stages. During the first stage, 12 percent (n=23,832) is set aside for

[19]The probability of being assigned to each experimental group had to be expressed in whole numbers, due to the Adobe software platform used to randomly assign customers as they login, resulting in each treatment group containing approximately 11 percent and the control group 12 percent of the sample.

[20]To avoid overwhelming electricity retailers with unexpectedly high call volumes, participants were enrolled in the experiment according to a phase-in scheme: 5 percent of the total sample was randomized as described into the nine experimental groups on May 4, and on the following nine business days 10 percent of the total sample was added each day, until the last 5 percent of the sample was added on May 21, 2018.

the pure control—individuals in this group received no intervention at all. The rest of the sample (n=171,582) is randomly assigned to four equally-sized treatment groups—each given a different Invitation. During the second stage, those responding positively to their invitation, by selecting "Tell me more" (n=31,299), are randomly assigned and directed to one of two Application Pages. (Those who select "No thanks", or ignore the invitation twice, are never shown an application page.)[21]

Summary Statistics and Balance

As expected given randomization, Table 1 shows that the sample is balanced across experimental groups on a range of covariates measured prior to the intervention.[22]

Appendix Table B1 summarizes a range of characteristics of the study population, measured prior to the intervention: the average age is 39 years, 59 percent is female, 85 percent uses this bank as their primary one, and 13 percent has been in arrears—defined as failing to pay a debt to the bank—in the 6 months leading up to the experiment. $395 is the median monthly combined bank account balance,[23] but some customers are heavily in debt, pulling the average down to $-15,154. Hence our study population is quite a lot worse off financially than the population at large (the average Australian saves an estimated $427 per month (Canstar 2019)).

[21] Because everyone was randomly assigned an application page at login, we know which page those who do not respond to the invitation would have seen had they shown interest, so we can perform Intent-to-Treat analyses if we want to use the full sample.

[22] Table 1 shows the three main arms. Appendix Table B2 shows balance across all nine sub-treatments.

[23] Across a customer's savings, transaction and credit card accounts, but excluding mortgages and loans.

Table 1: Balance of Study Population Across Main Experimental Groups

	Experimental Groups			Balance p-values		
	Standard	Simplified	Control	(1) vs. (2)	(1) vs. (3)	(2) vs. (3)
	(1)	(2)	(3)	(4)	(5)	(6)
Demographics+						
Age	39.2	39.2	39.0	0.341	0.337	0.740
Years with bank	16.6	16.7	16.5	0.677	0.216	0.130
Female	0.593	0.594	0.586	0.862	0.065	0.050
Primary bank	0.850	0.848	0.846	0.228	0.111	0.423
History of arrears	0.131	0.129	0.129	0.135	0.297	0.957
History of hardship	0.000	0.000	0.001	0.999	0.638	0.637
Bank Products						
Credit card	0.286	0.287	0.283	0.784	0.521	0.411
Deposit & trans.	1.950	1.951	1.941	0.884	0.280	0.231
Lending & loan	0.071	0.071	0.071	0.850	0.936	0.837
Term deposit	0.039	0.040	0.039	0.412	0.835	0.450
Home loan	0.123	0.123	0.128	0.895	0.134	0.159
Other product	0.437	0.439	0.431	0.650	0.227	0.132
Account Balances ($)						
Monthly combined	-15,001	-15,046	-16,095	0.908	0.070	0.082
Observations (N)	85,782	85,800	23,832			(Total: 195,414)

This table shows pre-treatment characteristics. Columns (1)–(3) each represent one of the three main experimental arms, and (4)–(6) show the p-value for differences between arms, as indicated. Panel 1 shows means or shares; Panel 2 shows mean number of bank products; Panel 3 shows the combined monthly balance for savings, transaction and credit card accounts (from Jan 2017–April 2018; winsorized at 2.5th and 97.5th percentiles). Postcode is excluded but balanced.

4.2 Defining the Study Population

The study population in Table 1 consists of bank customers who satisfy two sets of criteria. First, they must be users of the bank's mobile app who login during the study period.[24] Second, they must be 18 years or older, live in NSW, and have received a government benefit payment (a "concession") to one of their bank accounts within the last year.[25]

We consider those who satisfy these criteria likely eligible for the electricity benefit. This is an imperfect proxy, because bank data cannot reveal the final eligibility criterion: whether

[24]This means our results can be reliably generalized to digitally-active banking citizens—those who use a digital banking product—which comprise a growing majority of this bank's customers (51 percent during the study period, 53 percent in December 2019). Overall, 98.9 percent of Australians have a bank account (FDIC 2017), 89 percent own a smartphone (Deloitte 2018), and a fourth of the population uses this bank.

[25]This last criterion is a conservative measure of having a concession card, since it is possible to have one without having received a benefit payment to a CBA bank account recently.

someone is the electricity account holder at their primary residence. Whereas eligibility for the electricity benefit cannot be completely observed within bank data, the bank can do at least as good a job as any single government agency presently can.[26] Importantly, the bank also has frequent opportunities to organically interact with people and can do so at low cost.

Our inability to fully observe eligibility criteria is not unduly concerning, since it will affect people in randomly assigned arms equally. However, it does mean we are likely to contact some who are not eligible (e.g. have no electricity account). Relatedly, we cannot observe whether individuals are already receiving the electricity benefit prior to our intervention. Although it might be annoying for these groups to be contacted, the invitation is light-touch, easy to turn down, and the eligibility criteria appears immediately at the top of both application pages. Moreover, we are only providing information already widely available on government websites. Thus, when constructing the study sample, we choose (i) trying to reach all possible eligibles given the information at hand, instead of the alternative of (ii) not annoying already-takers (or non-eligibles). Encouragingly, we found little evidence of annoyance in our data.[27] [28]

4.3 Treatments

Stage 1: Invitations

The purpose of the invitations is to notify people about the electricity benefit they may be eligible for.[29] The four invitations are designed to make slightly different appeals, while keeping either the headline or second row constant, for comparability across treatments. This allows for testing whether the invitation is more effective in generating interest when it emphasizes that claiming is easy (invitation 1), when it is more transparent why you are targeted with the invite (2), or when it states the value of the benefit (primarily 3 vs. 4). Appendix Figure A5 shows what each of these invitations looked like within the app.[30]

[26]Few government departments and agencies routinely share data with each other (but see DPMC (2020) for an 'Open Government' plan). That means, for instance, that an agency that knows which citizens have a valid concession card does not necessarily know whether they have an electricity account.

[27]In early pilot testing using bank call centers, conducted before experiment 1 was designed, we saw no evidence of annoyance among already-takers or non-eligibles. In experiment 1, we hope already-takers realize they already have the benefit and select out of engaging with the invitation. If they do not, it should become immediately obvious on the application page that this concerns the electricity benefit.

[28]With respect to power considerations, since we know that baseline take-up prior to the study period is 72 percent on average across the state, we assume that at most 28 percent of the sample (n=54,715) could move from non-claimants to claimants of the electricity benefit during the study period.

[29]The tentative language reflects that we predict eligibility to the best of our ability given available data.

[30]Each headline has larger font size than the rest of the invitation, and might attract more attention.

Invitation Treatments:

1. Low Effort \$: “Get a \$285 electricity rebate You may be eligible for an annual rebate— it only takes a few minutes to claim”	**2. Transparency \$:** “Get a \$285 electricity rebate As a concession card holder in NSW you may be eligible for a rebate”
3. Generic \$: “Get a \$285 electricity rebate You may be eligible for a rebate”	**4. Generic:** “Get an electricity rebate You may be eligible for a rebate”

Below the invitation, two buttons are displayed: “No thanks” and “Tell me more”.

Stage 2: Application Pages

Those who respond to their assigned invitation by clicking “Tell me more” move onto a randomly assigned application page, which they are redirected to within the mobile app. Appendix Figure A6 shows what each of these application pages look like.

Standard

The first possible experience is a standard government application page, designed to as closely as possible mimic ‘business-as-usual’, and thus serves as a proxy for the government’s existing online page describing how to claim the electricity benefit.[31]

Simplified

As an alternative experience, we designed a simplified application page, with a briefer set of instructions for how to claim the benefit and clearer formatting. We also added “click-to-call” buttons for the seven most common electricity retailers, which we pre-programmed to enable individuals to immediately contact the right person at each retailer’s office to ask for the benefit. This stands in juxtaposition to the standard page, which contains the necessary information, but requires at least the additional steps of independently finding and calling the number for one’s electricity retailer, and navigating through the phone tree.

[31] We reconstruct the information from the government page within our own environment so we can track each person’s activity on the application page.

Figure 1: Sample Invitation (Generic \$) and Application Page (Simplified)

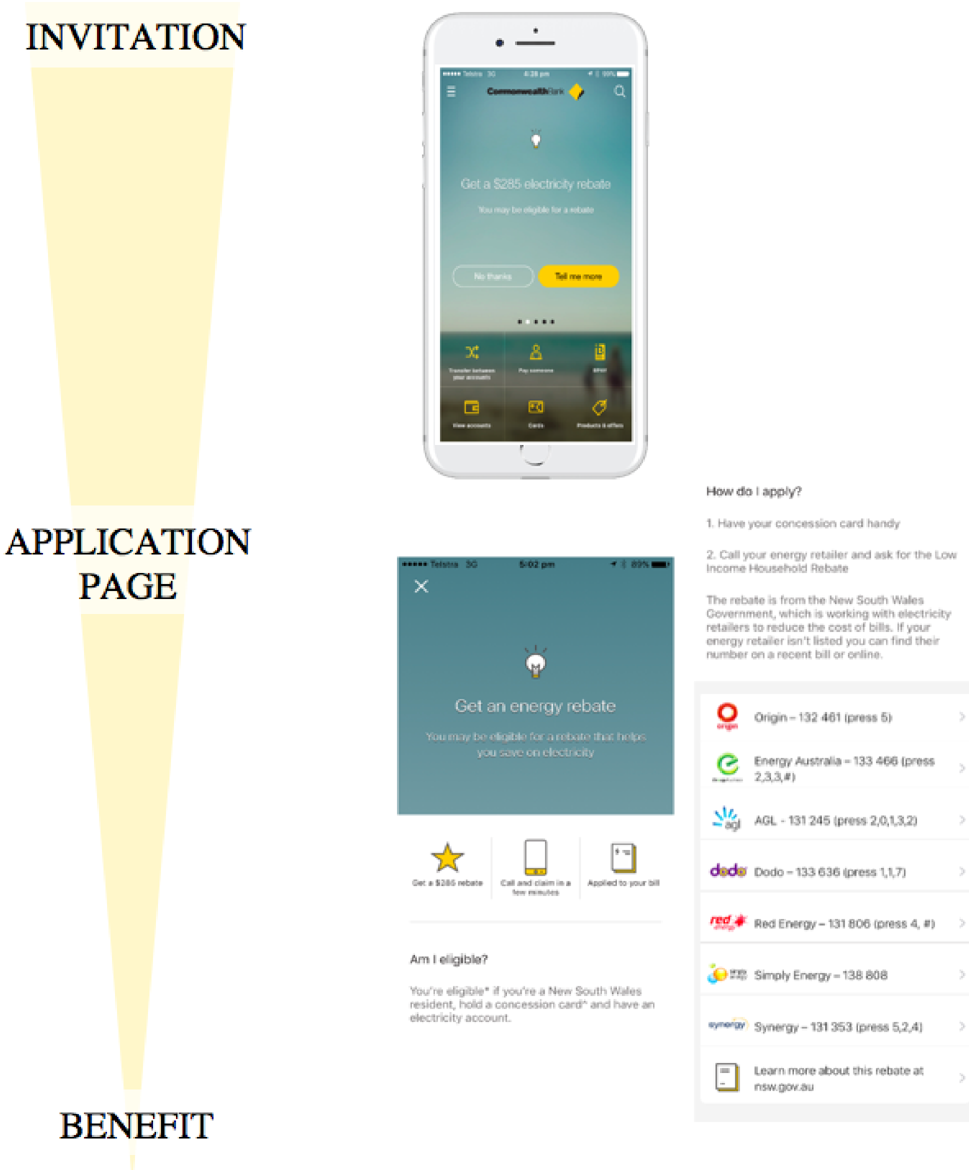

Figure 1 shows the design of the experience for an individual assigned to invitation 3 (Generic \$) and the simplified application page. See Appendix A for full treatment design details.

4.4 Outcomes Data

To evaluate the intervention, we measure the following dependent variables. Appendix Figure A4 illustrates when these main outcomes are observed within the experimental timeline.

Showing Interest in the Invitation

We measure whether individuals click "Tell me more", "No thanks" or ignore the invite.

Taking Steps Toward Claiming

We measure whether individuals scroll down to read the entire application page; whether they click on the link to go to the government's page for more information; and whether they "click-to-call" any of the electricity companies (available on simplified page only).

Attitudes

In contrast to most evaluations of efforts to increase program take-up, we measure attitudes. A post-treatment survey—designed to gather information about attitudes toward the bank and government—was sent to treatment group individuals on their next login after the intervention.[32] [33] This allows us to test the idea that citizens update their views of government as a result of their policy experience. Specifically, for this main hypothesis we use a question about satisfaction with the government. See Appendix Figure A8 for the survey instrument.

Financial Impact

We use two data sources to evaluate financial effects. First, bank records provide account balances for all customers, and direct payments of electricity bills for a sub-sample of 64,036 customers who pay theirs via this bank. Second, goverment records linked to bank data provide a ground truth measure of take-up for 15,656 customers of one electricity company.

5 Results

5.1 Do Invitations Generate Interest in Learning More?

Are people interested in getting information about the electricity benefit via their bank? We test this with the following specification:

[32] The survey is sent to people through the mobile app. The same trigger rule is used to invite everyone, but when they actually see the survey invite depends on their next login after finishing the intervention. Here "finishing" is defined as either of the following: clicking yes "Tell me more" on the invitation and moving on to the application page; clicking "No thanks" on the invitation; or logging in twice to the app and ignoring the invitation (i.e. taking no action). On average, people respond to the survey 8.4 days after finishing the intervention, with a range of 0-63 days (see Appendix Figure A7).

[33] We also designed a control version of the survey, sent to individuals assigned to the control group on a roll-out schedule matching the timing and volume of treatment survey invitations. Unfortunately, we cannot use control survey responses for comparison, due to decisions taken during implementation. Appendix section A.2.1 provides more information and demonstrates that both the control invite and the control survey itself resemble a treatment condition, rather than a treatment survey comparison.

$$Interest_i \quad = \quad \alpha + \beta Invitation_i + \epsilon_i \tag{1}$$

where $Interest_i$ is the binary outcome for clicking yes "Tell me more", $Invitation_i$ indicates assignment to a treatment message for individual i, and α and ϵ_i represent the constant and error terms, respectively.

First, we find that people are interested in information about the benefit: the proportion clicking "Tell me more" ranges from 11.9 to 21.4 percent depending on the wording of the invitation (see Table 2). Across the four invitations overall, 18.2 percent want to learn more.[34] In fact, this was the most popular 'messaging campaign' at the bank to date.[35] We find little evidence of distrust, or that people find it strange their bank is contacting them about a government benefit: only 5 percent of survey respondents report they did "not trust the message or process" as the reason they did not manage to claim the benefit; and 91 percent say they would find it valuable if the bank told them about similar benefits in the future (see Appendix Figures C1—C2).

Second, the most important factor for generating higher interest is including the $285 benefit amount in the invitation. As shown in Table 2, showing the dollar-value increases click-through by 66-80 percent.[36] This suggests there is at least some cost-benefit calculation going on in individuals' minds when they are considering whether seeking more information is worthwhile. This is in line with the economic theory that a person's incentive to obtain information about a program is influenced by the size of the benefit relative to the effort (transaction cost) associated with applying.[37]

Third, we also examine how the invitation wording matters for subsequent claiming behavior (the topic of next section). We find that a higher fraction of those in the Transparency 'channel'—those who are interested enough to visit an application page after reading the Transparency invitation—take steps toward claiming than those arriving on the application page via the other three invitations. (See Appendix section C.1.1: the rate is about 13 percent, or 1.2 percentage points, higher in the group showing interest in Transparency compared to the Generic invitation, p=0.024). One plausible interpretation is that the

[34] Of the total n=171,582 assigned to treatment, n=31,299 or 18.2 percent clicked "Tell me more".

[35] For comparison, the overall proportion of people who clicked "Tell me more" on the electricity benefit message is over twice the average click-through rate of other bank messages during the same period, which were 7.9 percent on average for Retention campaigns (messages that help customers optimize use of their products) and again 7.9 percent for Acquisition campaigns (messages that promote banking products).

[36] All three comparisons with the 'Generic' invitation are significant at the $p<0.01$ level. The lower bound of 66 percent is the comparison between 'Generic' and 'Low Effort', 67 percent is the comparison with '$ Transparency', and the upper bound of 80 percent is the comparison with '$ Generic'.

[37] E.g., people entitled to larger benefits are more likely to know about food stamps (Daponte et al. 1999).

Transparency invitation provides more up-front information pertinent to eligibility, which helps filter out people who will not try to apply and attract those who might.

Table 2: Clicking 'Yes' in Response to the Invitation

	Dependent variable:
	Yes, "Tell me more"
Invitation:	(1)
Low Effort \$	0.198***
	(0.003)
Transparency \$	0.199***
	(0.003)
Generic \$	0.214***
	(0.003)
Generic	0.119***
	(0.003)
Pure Control (intercept)	−0.000
	(0.002)
Observations	195,414
R^2	0.036
Adjusted R^2	0.036
Residual Std. Error	0.360 (df = 195409)
F Statistic	1,805.061*** (df = 4; 195409)

This table shows results from a linear probability model of a binary outcome variable for clicking yes, "Tell me more", on treatment indicators for each of the randomly assigned invitations. The pure control group is the omitted category, which is mechanically zero. The sample includes everyone included in the study. Every estimate is the proportion of individuals assigned to each invitation that wants more information, and the differences between each pairwise comparison are statistically significant (p=0.00) except for 'Low Effort \$' and 'Transparency \$', which are statistically indistinguishable (p=0.47). The overall rate of those assigned to treatment who want more information is 18.2 percent (31,299/171,582). *p<0.1; **p<0.05; ***p<0.01

5.2 Does Simplifying the Process Affect Steps Toward Claiming?

Once individuals are on the application page, do they take the next steps toward claiming the government benefit? We test this with the following specification:

$$Action_i = \alpha + \beta Page_i + \epsilon_i \tag{2}$$

where $Action_i$ is the binary outcome for clicking on the government link or click-to-call, $Page_i$ indicates assignment to an application page for individual i, and α and ϵ_i represent the constant and error terms, respectively.

Table 3: Steps toward Claiming, by Application Page

	Dependent variable:	
	Clicking 'Call' or 'Government Link'	
	(1)	(2)
Simplified	0.018***	0.099***
	(0.001)	(0.003)
Standard (intercept)	0.009***	0.048***
	(0.0004)	(0.002)
Observations	171,582	31,299
R^2	0.005	0.028
Adjusted R^2	0.005	0.028
Residual Std. Error	0.132 (df = 171580)	0.291 (df = 31297)
F Statistic	794.178*** (df = 1; 171580)	894.951*** (df = 1; 31297)

This table shows results from a linear probability model of binary outcome variables for taking steps toward claiming that we can observe in our application page environment—clicking either the government link (available on both pages) or the call button (simplified page only)—on a treatment indicator for the simplified page, with the standard page as the omitted category. In Column (1) the sample includes all people who were randomly assigned to any of the invitation treatment groups, and in (2) the sample includes all people who clicked 'Yes' in response to their invitation and actually experienced their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

We find that people take action on both application pages, but critically, that simplifying the experience substantially increases this probability: those assigned to the simplified application page are 304 percent (9.9 percentage points) more likely than those on the standard government-like page to take observable next steps toward claiming, as seen in Table 3.[38]

Of course, we do not measure claiming behavior outside of our experimental environment—so we do not capture, for example, whether people assigned to the standard page look for more information or call their electricity companies after logging out of the app. But we have some evidence suggesting that those experiencing the standard page are less likely to read all the information provided.[39] If fewer people read the application instruc-

[38]In absolute numbers, 2,289 (out of the 15,595) people who saw the simplified page took steps toward claiming, versus 752 (out of 15,704) people who saw the standard page.

[39]We know that on the standard application page, by design, more scrolling is required to reach all the information about how to apply (see Appendix Figure A6). Scrolling down to the footer is a reasonable

tions on the standard government-like page, we should expect fewer of them to take steps toward applying on a later occasion as well. If unobserved behavior translates into take-up, we should capture it in our financial data, where we observe all experimental groups equally.

5.3 Financial Effects

Did the intervention affect take-up? We use two sources of data to evaluate financial effects: bank records, including account balances and direct payments of electricity bills; and goverment records, linked to bank customer data to provide a ground truth measure of take-up.

A few notes: because of the design of the electricity benefit, there is no direct measure of receiving the electricity benefit in the bank data (i.e., there is no direct deposit that we can identify as a separate transaction). Rather, we must infer the presence of the benefit: via higher account balances or lower electricity bills. Because of these limitations, we also secure and analyze official data on take-up for a sub-sample of the study population. With the help of the New South Wales Department of Planning, Industry and Environment (DPIE),[40] and in collaboration with the bank, information from bills for a sample of 15,656 customers, who get billed by one of the largest electricity companies, are matched with government records on take-up. These matches are provided to us at the experimental group level, to protect customer anonymity. Though aggregated and available for a smaller set of the study population, these data are the direct measures of take-up.

Bank Records

Bank Account Balances

Account balances are available for all customers, before and after the intervention. However, two aspects of these data could decrease our ability to detect an effect: first, if people simply spend any discount they get on something else, in the same time period as they get the benefit, we would not observe a monthly difference in their account balances. Second, account balances are noisy (see Appendix Tables C4–C6).[41] This means that even in the most extreme scenario—where the intervention causes full take-up in the simplified group (an extra $24 per month for the modal beneficiary) and no take-up in the standard group—we would likely be unable to detect that effect: even with customer and month fixed effects, and restricting the sample to customers with less extreme balances, the standard error (71) is

proxy for looking at the entire page, and those experiencing the simplified application page are 42 percent more likely to scroll down to the footer (Appendix Table C1).

[40]DPIE is the government agency in charge of reimbursing electricity companies for the electricity benefit.

[41]The data had more noise than expected; this is the first trial in this type of population using such data.

larger than twice the maximum expected average monthly effect size (24×2) (see Appendix Table C4. Therefore, we focus instead on the sub-sample of electricity bills.

Electricity Bills Paid via the Bank

Electricity bills are observable for the customers who pay theirs via this bank. The electricity bill data have a few advantages compared with the account balance data, despite covering a sub-sample of the experimental population.[42] First, bill amounts are more stable over time, increasing precision of estimates. Second, the effect is observed even if the money is immediately spent on something else (other than electricity), as a discount of \$0.78 for each day included on the bill.[43]

We test whether the intervention had an effect on customers' electricity bill amounts, using the following specification:

$$ElectricityBills_{it} = \alpha_i + \gamma_t + \beta Treat_{it} + \epsilon_{it} \tag{3}$$

where $ElectricityBills_{it}$ is the bill amount and $Treat_{it}$ the treatment status for customer i at month t, respectively. α_i and γ_t are customer and month fixed effects, ϵ_{it} the error term.

First, we find no statistically significant difference in average bill amounts between those assigned to the simplified and standard page (Table 4).[44] In the main specification, we restrict the sample to customers with bills from the 7 electricity companies listed on the simplified page, since we expect to see the strongest effect among those benefiting from the click-to-call options.[45] In Table 4, we focus on the time window that allows each person to receive at least two (quarterly) bills after the intervention, but results remain substantively similar and statistically unchanged for different time windows (see Appendix Table C7).

[42] We observe bills for a third of the sample (n=64,036) who use this bank to pay (what looks like) their electricity bills. We cannot perfectly classify transactions as electricity charges, but these data represent our best attempt based on transaction details such as the name of the company, amounts and frequency.

[43] Of course, if people who get the discount increase their electricity consumption as a result, their bills could look the same pre- and post-intervention, even if they claim the benefit. To detect no effect in such a scenario would mean customers are increasing consumption by the equivalent of 0.78 cents per day, which seems like a quite unlikely scenario.

[44] This finding holds when comparing the pure control with any treatment group (Appendix Table C9), suggesting that information alone was not enough to increase take-up.

[45] The results remain substantively and statistically similar in samples not restricted to specific electricity companies or less extreme bill amounts, though estimates become less precise (see Appendix Table C8).

Table 4: Effect of Simplified Treatment on Electricity Bill Amounts

	Dependent variable:
	Bill Amount
Simplified	−0.580
	(1.608)
Customer Fixed Effects	Y
Month Fixed Effects	Y
Standard Page Mean	180.599
SE	(1.078)
Observations	67,594
Customers	20,061
Time Window	12 months
R^2	0.00000
F Statistic	0.130

This table shows results from a regression model of a continuous variable for electricity bill amounts on a simplified treatment status indicator, and customer and month fixed effects. The sample includes bills for 4 months pre and 8 months post intervention; and is restricted to customers with bills from the 7 electricity retailers listed on the simplified page, and bill amounts up to $500. *p<0.1; **p<0.05; ***p<0.01

Second, we note that even with customer and month fixed effects the estimates are imprecise, but back-of-the-envelope calculations illustrate that any effects on take-up are likely modest. If we took the point estimate (−0.580) in Table 4 as given, algebraically that would mean an additional 0.8 percentage points of take-up in the simplified group post-intervention (with some assumptions).[46] The upper bound of the effect likewise corresponds to at most a 3 percentage point increase in take-up relative to the standard treatment.[47]

Government Records

We obtain ground truth data on take-up by matching official government benefit records with bank customer records. This matched sample consists of 15,656 customers with at least one

[46] We can translate the −0.580 into a take-up estimate because we have a good idea of the 'true' effect size in those who claim: we know that bills are quarterly, and assuming the modal benefit size, each successful take-up would generate an expected electricity discount of $71 per quarterly bill. The back-of-the-envelope calculation is simply solving for $f = 0.008$ in $[-71 \times f + 0(1-f) = -0.580]$, where f is the fraction of additional bills in the simplified group that get the electricity benefit applied after the intervention. We back out the full −$71 effect for the fraction f of the treatment group, which gets diluted by people in the treatment group who do not claim ($(1-f)$, with $0 effect) thus lowering the average effect to $−0.580.

[47] The lower bound of the 95 percent confidence interval for the coefficient in Table 4, $-2.188 = -(0.580 + 1.608)$, would translate into an upper bound of a fraction of 3 percentage points more bills in the simplified group getting the electricity benefit applied, relative to take-up in the standard group.

bill for one large electricity company.[48] We do all analysis at the level of the experimental group, not the individual-level, to allow for data sharing without compromising individual identities. We have aggregate data from the department (DPIE) on the fraction of customers in each experimental group that gets the benefit applied to their electricity bills, before and after the intervention. This allows us to directly assess the effects of treatment on take-up for this subset of the study population.[49]

First, we generate a fact about the targeting strategy used to define the overall study population: linkage indicates 39.0 percent were already receiving the benefit in the pre-period before the intervention. In contrast, had we selected a sample at random from the customer base, that would have yielded only around 13 percent already-takers.[50] Our main goal is reaching eligibles who are not yet claiming, but confirming this prevalence of eligibles in the study population suggests the inclusion criteria help target a group of customers who (i) had much higher rates of eligibility and benefit receipt than we would have reached in a non-targeted campaign (39 vs. 13 percent); but also who (ii) were not already receiving the benefit at high enough rates that the trial would be hampered by a ceiling effect.

Second, we test the effect of the intervention on take-up, and find that the simplified treatment increases take-up by 0.8 percentage points (2.1 percent) relative to the standard treatment, as seen in Table 5.[51] This is almost exactly the effect size estimated in the analysis of the full sample of bills observed in bank data. But because there is much less measurement error here, these results are more precisely estimated.

Third, comparing take-up in the pure control group with the two treatment groups reveals that information without simplification had little to no effect. The difference of a 0.1 percentage point increase in take-up in the standard treatment relative to the pure control is not statistically significant ($p=0.76$), suggesting that reaching out with targeted information—but without using design to further simplify a standard government-like page—is not enough to move take-up. The difference of a 0.9 percentage point increase in take-up in the simplified treatment group relative to the pure control is statistically significant ($p<0.01$), suggesting that the take-up effect is driven by actually experiencing the simplified application page. Logically, since the standard treatment had no significant take-up effect compared

[48] DPIE receives information from all electricity companies on which customers have been granted the discount. However, because of differences in how billing information is generated, shared and stored across systems, we are only able to accurately match bank and government records from one electricity company.

[49] This sub-sample constitutes 8.0 percent of the total study population (15,656/195,414=0.0801).

[50] In the general NSW adult population, 18 percent are eligible, and take-up is around 72 percent.

[51] Applying the effect size of 0.08 percentage points to the entire n=85,800 sample assigned to simplified treatment translates into an upper bound of 686 new people getting the benefit, for a total of $195,510, as a result of the intervention. In this sub-sample of n=15,656 customers paying bills from one electricity company, it means 55 new claimers and a total of $15,675.

to pure control, any effect of simplified treatment has to be explained not by information but by the only difference with the standard page: simplification.

Table 5: Take-Up for Customers of one Electricity Company, per Government Records

	Take-Up Difference	*p*-value
Simplified vs. Control	0.009 [0.004; 0.015]	p<0.01
Simplified vs. Standard	0.008 [0.005; 0.011]	p<0.01
Standard vs. Control	0.001 [−0.005; 0.007]	p=0.76

This table shows results from comparing the difference in proportions of take-up, before and after the intervention, between three experimental groups: the simplified page, the standard page and the pure control. Columns 1 and 3 show the estimated difference in take-up between groups, after controlling for baseline take-up pre-intervention. The 95 percent confidence intervals for the differences in proportions are shown in brackets. Columns 2 and 4 show the p-values of the tests of the differences in proportions. The sample includes 15,656 customers who pay bills via the bank to one specific electricity company, with data reported as the fraction of take-up at the experimental group level pre- and post-intervention. Data were provided by the NSW Department of Planning, Industry and Environment.

In sum, using data from government records on take-up, linked to bank customer data for a sub-sample of 15,656 citizens, we find that the simplified treatment increases take-up by 0.8 percentage points (2.1 percent) relative to the standard page and 0.9 percentage points (2.7 percent) relative to the pure control group.[52]

5.4 Cost-Effectiveness

While the absolute effect on take-up is small, it is important to consider cost-effectiveness.[53] Most interventions to drive take-up require humans to spend time helping someone fill out

[52] Percentage increases are calculated relative to post-intervention take-up rates in the comparison groups.

[53] Here we consider the bank's costs (in the form of building costs) and benefits (in the form of more money for their customers). One could also consider government's costs and benefits: the state government of NSW funds the electricity benefit. In terms of marketing costs, this intervention by CBA is costless for the government. There are costs from processing applications and paying out additional benefits, but the former are small and borne by electricity companies, the latter are welcomed by the government as part of its expressed goal to increase take-up for needy households.

forms or submit an application (e.g., Finkelstein and Notowidigdo (2019)). By contrast, these interventions are delivered on a digital platform, where scale-up is effectively costless.[54]

These interventions are highly cost-effective, given the observed effects on take-up. First, in the initial experiment alone the intervention paid for itself, even with conservative assumptions of benefit (see footnote 51). We estimate that creating this specific intervention—a targeted message, leading to a simplified page with a click-to-call function—took 117 person-hours and cost \$14,584 at internal transfer prices (opportunity cost).[55] Second, simply extending the one-time intervention to a larger, readily available sample who meet the same criteria as the initial study population, yields an expected \$106 for customers (yearly) by each dollar CBA spent on the build.[56] Third, in the national scale-up, when customers have more than one chance to visit the simplified page and use the click-to-call function, take-up is higher, resulting in much higher cost-effectiveness: for each dollar CBA spent, \$2,125 was generated for customers in the first year.[57] Lastly, also worth noting is that the national scale-up is expected to generate better, simplified policy experiences for millions of citizens. As detailed in the next section, these improved policy experiences—even absent take-up—can be intrinsically consequential for attitudes toward government.

5.5 Does the Experience Affect Attitudes?

Finally, we turn attention to whether the policy experience affects attitudes, by examining the information gathered via the post-treatment survey.[58] Here we test if reducing administrative burden for one policy leads to more positive attitudes toward government in general.

[54] Unlike many take-up efforts, this intervention relies on digital infrastructure and privileged data; once in place, these allow the experience to be extended at low or no marginal cost. We field the interventions using the bank's existing infrastructure (customer data, mobile app, marketing engine), routinely used by nearly every bank division; no fraction of that fixed cost is included here. (Marginal costs of maintaining the codebase for the intervention are negligible, so we ignore them in back-of-the-envelope calculations.)

[55] The cost of running the intervention as an academic experiment was three times higher (\$43,750). CBA estimates that building the intervention and implementing the experiment with us took 350 hours shared by 4 staff (data engineer, digital marketing specialist, content writer, and behavioral science manager).

[56] Cost-effectiveness calculated as follows: 1) 601,000 people$\times$0.009 treatment effect=5,409 new benefit-takers. 2) 5,409 people$\times$\$285=\$1,541,565 generated yearly. 3) \$1,541,565$\div$14,584=\$106 yearly for customers by each dollar spent on the build. 4) Including experimentation costs yields \$35 yearly for each dollar spent.

[57] \$30,993,415 in energy savings generated for n=1,956,655 (\$1.32 monthly), at a building cost of \$14,584.

[58] As expected, the response rate was low—0.6 percent in the entire sample, and 2.2 percent among those who saw an application page—but in line with other comparable bank surveys, which have a 0.4 percent response rate on average. Only those who login to the app again post-intervention, but before the survey period ends on July 21, see their survey invite. The response rate is calculated based on individuals who saw their survey invite: 950$\div$155,120 for the entire sample, 620$\div$27,894 for those who saw an application page.

Survey Sample and Selection

First, a note on the survey data: the whole sample is invited to take the survey, but the information we get is from people who select into responding. Therefore, we cannot assume these respondents are a random sample—or perfectly representative—of the entire population assigned to standard or simplified treatment.[59] Moreover, to make valid inferences, we need to be as sure as possible that we are comparing survey responses among individuals who are on average equal except for treatment. We find that those assigned to the simplified treatment are 15.8 percent more likely than those assigned to the standard treatment to respond to the survey (p=0.02).[60] This means survey respondents differ in these two treatment groups at least in the sense that they decide to respond at different rates. While respondents in both groups look similar on observable pre-treatment characteristics (see Appendix Table B8), we do not know how they differ on unobservable characteristics.

We take extra precautions to deal with the difference in response rates, and create a main dependent variable that lets us avoid simply analyzing outcomes from the sub-sample of survey *respondents*, which could suffer from selection bias. Because of randomization, we know that those assigned to the standard and simplified pages are balanced on all pre-treatment covariates, and we expect them to be balanced on unobservables as well. In the main analysis, we thus focus on outcomes among all individuals who were randomly assigned to either treatment page—all of whom had an equal opportunity to respond to the survey.[61]

Dependent Variable Construction

Since the empirical strategy outlined above implies comparing outcomes among respondents and non-respondents alike, we construct outcome measures that incorporate the choice to respond (or not) to the survey: each dependent variable combines the survey outcome of interest—satisfaction, dissatisfaction, wanting future interventions —with a binary indicator for submitting a survey response.

Besides ensuring that the only difference between groups is their randomly assigned treatment status, this strategy makes sense from a purely substantive point of view as well: a plausible interpretation of the significant difference in response rates between the simplified

[59]See Appendix Table B6 on how survey respondents differ from non-respondents on observables. These differences affect the extent to which survey findings generalize to the entire study population.

[60]The response rate by application page treatment is as follows: in the simplified group 508÷77,523 (0.66 percent) respond to the survey, and in the standard group 442÷77,597 (0.57) respond, for a difference of 0.9 percentage points, or a 15.8 percent increase in the simplified relative to the standard treatment group.

[61]We know that everyone had the same opportunity to choose to respond to the survey. By comparing survey outcomes between the entire groups randomly assigned to different application pages, we ensure that we are comparing individuals that are on average equal except for treatment.

and standard treatment groups is that the simplified experience made customers more willing to respond to a survey, which is notoriously difficult. Survey response is a meaningful outcome, and is routinely interpreted as a measure of (dis)satisfaction. We capture that effect, and examine whether the increased willingness to respond should be characterized as satisfaction or dissatisfaction, by constructing the following main outcome: a binary indicator for submitting a positive (negative) survey response.[62]

Attitudes Toward the Bank and Government

We test whether the intervention affects attitudes using the following main specification:

$$SurveyResponse_i \quad = \quad \alpha + \beta Page_i + \epsilon_i \tag{4}$$

where $SurveyResponse_i$ is the binary indicator for submitting a positive (negative) survey response, $Page_i$ indicates assignment to an application page for individual i, and α and ϵ_i represent the constant and error terms, respectively.

We find that experiencing the simplified page makes individuals on average 45.5 percent (0.5 percentage points) more likely to submit a survey that rates the bank positively, and 44.4 percent (0.4 percentage points) more likely to submit a survey that rates the government positively. As expected, the satisfaction results are robust to including the entire treatment group in the sample, but are driven by those who actually experienced their application pages (Appendix Table C10).[63] Table 6 shows that the extra survey responses submitted by those who experienced the simplified page are postive, not negative. Individuals thus give credit to both the bank and government for making the application process more straightforward.[64] These findings constitute causal evidence that changing the policy experience can change satisfaction with government, as measured up to 63 days after treatment.[65]

[62]Here, positive is defined as above the mean satisfaction score among survey submitters, and negative as below. For satisfaction with the bank, this means positive is a score of 8 and above (on a scale of 1–10), which is similar to what the bank considers good in routine measurements of Net Promoter Score (NPS), which measures customer experience of a brand. For satisfaction with government, positive is a score of 6 and above, which again corresponds well with real-world benchmarks: in Australia, in this period baseline satisfaction with the government was about 20 percent lower than satisfaction with businesses (Edelman 2018). Results are robust to different threshold specifications.

[63]The increase in bank satisfaction, caused by the simplified treatment translates into an NPS-like increase of 31 percent. (Using the standard NPS formula—except that detractors range from 1–6, and we ask about 'satisfaction' with the bank rather than intent to 'recommend to a friend or colleague'—gives an NPS of 23.2 among respondents assigned to the standard group and 30.5 among those in the simplified group.)

[64]These are large effects, especially given the general climate around the time of the intervention: 56 percent of Australians considered the government the most broken institution (Edelman 2018), and satisfaction

Table 6: Satisfaction with the Bank and the Government, by Application Page

	Dependent variable: submitting a survey with			
	Positive Rating (Satisfaction)		Negative Rating (Dissatisfaction)	
	Bank	**Government**	Bank	Government
	(1)	(2)	(3)	(4)
Simplified	0.005***	0.004***	0.0003	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
Standard (intercept)	0.011***	0.009***	0.005***	0.007***
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	31,299	31,299	31,299	31,299
R^2	0.0005	0.0004	0.00000	0.0001
Adjusted R^2	0.0005	0.0003	−0.00003	0.0001
F Statistic (df = 1; 31297)	15.596***	11.052***	0.124	2.689

This table shows results from a linear probability model of binary outcome variables for submitting a survey that rates the bank (1)(3) or government (2)(4) positively or negatively—on a treatment indicator for the simplified page, with the standard page as the omitted category. Here, positive (negative) means a rating above (below) the mean satisfaction score, but results are robust to other specifications. The sample includes all people who visited their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

Money Versus Experience? Possible Mechanisms

By what mechanism does the intervention affect attitudes? A few things are plausible: people get the $285, or people have a better experience, or some combination of both. We propose several interpretations of the results, and use the survey responses to probe furter.

The first possible interpretation is that the intervention increases take-up, which directly increases satisfaction. We can think of this as the "people like money" explanation. The logic is as follows: simplifying the process translates into higher take-up, which leads to more money (or anticipation of more money).[66] People feel content about the money, and attribute this future windfall not just to their own success managing to claim, but to the efforts by both the bank and government. As discussed in Section 5.3, we find a 0.8 percentage point

with banks was historically low following the "Finance Royal Commission" (Roy Morgan Research 2019).

[65] On average, people responded to the survey 8.4 days after the intervention. See Appendix Figure A7 for the full distribution of time between intervention and survey responses.

[66] Few people will have received their discount when taking the survey, though they are told during the phone call whether their application is successful: on average, survey responses were submitted 8.4 days after the intervention, whereas electricity bills where any discount would be applied are typically issued later, on a fixed schedule every three months. But anticipating saving money could increase satisfaction.

increase in take-up, and many of these new claimers could be survey respondents.

The other possibility is that designing an experience with fewer frictions affects satisfaction independently of financial rewards. We can think of the two sub-explanations under this heading as variants of "the experience matters". On one hand, we could imagine a story involving less time and mental effort: that the simplified experience leads to higher satisfaction (relative to the standard experience) because there are fewer hassles imposing taxes on people's time and cognition. Even if nobody got the benefit as a result of the intervention, those who saw the simplified page might still feel relatively more satisfied because they wasted less energy learning about the policy, understanding how to apply, and so on.

Alternatively, people may form impressions about the functioning of institutions through their interactions with them. People experiencing the simplified page could come to think the bank and government are more efficient than they previously thought, or update positively about their helpfulness, after learning about their efforts with programs—like the electricty benefit—designed to help people (e.g., even if they themselves are not eligible).[67]

To probe these mechanisms further, we can look at the submitted survey responses. There are two caveats to keep in mind: first, because people select into answering the survey, respondents from the two treatment groups could still differ on unobservables, despite being statistically balanced on observable pre-treatment characteristics (Appendix Table B8).

Second, these are self-reported take-up claims, so even though it is not clear why respondents would lie, we cannot validate them. Thus, these measures are not at all preferable to the analyses presented above, which quantify take-up in the entire (unselected) sample. We use them here only to explore the relationship between (self-reported) claiming and satisfaction, and are willing to tolerate sample selection for the purposes of hypothesis generation.

Examining the 950 survey responses reveals a few pertinent facts: first, respondents assigned to the simplified page are not significantly more likely to self-report claiming (see Appendix Table C12–C13). Second, as shown in Table 7, among all respondents, (self-reported) claiming of the benefit is associated with higher satisfaction with both the bank and government. This is perhaps not very surprising, and supports the "people like money" explanation. Third however, among those who experienced the simplified page, there is additional satisfaction—on the order of magnitude of 26–29 percent for the government and bank, respectively—that is unexplained by self-reported claiming. This supports the explanation that the experience itself matters.

[67]Those experiencing the simplified page seem more likely to read the whole page (Appendix Table C1), suggesting they are more likely to digest all information (without spending more time: Appendix Table C2).

Table 7: Relationship Between Satisfaction and Reported Claiming

	Dependent variable: positive rating	
	Bank Satisfaction	**Government Satisfaction**
	(1)	(2)
Claimed (self-report)	0.251***	0.223***
	(0.037)	(0.035)
Simplified	0.096***	0.068**
	(0.031)	(0.030)
Intercept	0.331***	0.258***
	(0.024)	(0.024)
Observations	950	950
R^2	0.058	0.046
Adjusted R^2	0.056	0.044
F Statistic (df = 2; 947)	29.114***	23.089***

This table shows results from a linear probability model of binary outcomes for rating the bank (1) or government (2) positively—on a treatment indicator for the simplified page (with the standard page as the omitted category), and an indicator for saying one successfully managed to claim the benefit. As before, positive means a rating above the mean satisfaction score, but results are robust to other specifications. The sample includes all survey respondents. *p<0.1; **p<0.05; ***p<0.01

6 National Scale-Up

Based on the first experiment, CBA decided to build "Benefits Finder", a permanent platform to help connect customers with any government benefits they are not yet claiming. This is a large-scale extension: the platform also uses bank data to target messages, predict eligibility, and simplify application processes, but for over 200 benefits across Australia, including support payments later introduced in response to the Covid-19 crisis. In contrast to the first intervention, the platform gives customers repeated opportunities to act: it is permanent, searchable, and discoverable from many sources; it also allows customers to actively confirm eligibility criteria and schedule reminders to start a claim. Since the launch, over 2 million people have visited the platform, and more than 1.5 million applications have been initiated (CBA 2021b). The number of monthly applications initiated on the platform doubled after the Covid-19 outbreak (CBA 2020), suggesting the usefulness—especially in times of crisis—of having simplified information gathered in one place.

6.1 Intervention Design and Data

Set-Up and Randomization

The national scale-up intervention was live between July 23, 2019 and Sep 10, 2020. Customers were assigned to either treatment, with 80 percent probability, or control, with 20 percent probability.[68] Subsequently, digitally-active bank customers over 18 years old were invited via the app to use the platform. Those assigned to control got no invite, but in contrast to experiment 1, control group customers could benefit from treatment if they found the platform organically via the bank's website or app menus. The final study population consists of the 7,720,009 customers (6,6311,834 in treatment and 1,408,175 in control).

Summary Statistics and Balance

Table 8 shows that experimental groups are reasonably balanced on a range of pre-treatment covariates; any differences are substantively small but statistically significant given the size of the sample. In the analyses below, we show results with and without covariate controls.

[68] To facilitate any future data-sharing by partners such as government agencies, who navigate complicated privacy concerns, treatment was randomly assigned at the postcode-level (2,537 postcodes).

Table 8: Balance of the National Study Population

	Treatment (1)	Control (2)	P-value (3)	St. Diff. (4)
Demographics+				
Age	42.25	43.01	0.00	-0.05
Years with bank	16.81	17.32	0.00	-0.04
Female	0.49	0.49	0.00	-0.00
Concession card	0.16	0.15	0.00	0.01
NSW	0.35	0.35	0.50	-0.02
VIC	0.28	0.26	0.00	0.00
QLD	0.18	0.18	0.00	-0.03
WA	0.08	0.10	0.00	0.00
TAS	0.02	0.02	0.00	0.01
SA	0.05	0.05	0.00	-0.00
ACT	0.02	0.02	0.00	0.05
NT	0.01	0.01	0.00	-0.07
Bank Products				
Small business	0.00	0.00	0.00	0.01
Credit card	0.44	0.46	0.00	-0.02
Savings account	0.86	0.86	0.00	-0.01
Transaction account	1.14	1.14	0.05	0.00
Number of customers	6,311,834	1,408,175	total:	7,720,009)

This table shows pre-treatment characteristics. Column (1) represents Treatment and column (2) Control. Column (3) shows the p-value for differences between these two experimental groups. Column (4) shows standardized differences (difference in means divided by control standard deviation). Each covariate shown as a mean or share, except 'credit cards', 'savings' and 'transaction account', which indicate mean number of these products.

This study population differs from that of the first experiment in a few ways. In this extension there are customers from all states, and because we did not make having a concession card an inclusion criterion, this population is more gender-equal (compared to 59 percent female), slightly older (by 3 years on average) and on average better-off financially.

Treatment

Figure 2 shows the design of the "Benefits Finder" platform studied in experiment 2. This intervention builds on the best-performing treatments from experiment 1. First, customers assigned to treatment get sent a platform invitation via the bank's app. Those who visit the platform then go through a few quick steps: the customer is asked to answer five simple eligibility questions, after which each customer is shown a personalized list of benefits, where clicking on any suggested benefit leads to a simplified application page. Just like in experiment 1, the electricity benefit application page has instant click-to-call functionality. Expanded features include the choice to get a reminder to come back and apply later (CBA 2021a). See Appendix D for full treatment design details.

6.2 Results

Engagement and Simplified Experiences

We find that customers across Australia are interested in using the platform: 1,190,275 people (18.9 percent of the treatment group) click on the invitation, which is similar to the effect size in the first experiment. A total of 2,162,768 people visit the platform. This is higher than the number engaging with their invite, showing that people also find the platform on their own via the app or website menus.[69]

On the platform, a total of 1,308,656 people (20.7 percent) from the treatment group answer the five questions and are shown a personalized list of benefits. 352,637 of them use the click-to-call function (5.6 percent of the total treatment group).[70] For comparison, in the first experiment, 2,289 people (2.7 percent of total assigned to treatment) used the click-to-call function. In this scale-up design, the sample is larger and people have many opportunities to visit the simplified page and use its functionality; here observable engagement for the final step of applying for energy benefits is about twice as high.

[69] The total 2,162,768 customers who visit the platform comprises 157,222 people assigned to control and 2,005,546 assigned to treatment. From the control group, 11.2% (n=157,222) organically found their way to the platform. From the treatment group, a total of 31.8% (n=2,005,546) visited the platform: 18.9% (n=1,190,275) via clicking on the invitation link, and an additional 12.9% (n=815,271) organically without clicking on the invitation link.

[70] An additional 12,738 people assigned to the control group (0.9 percent of the total control group) used the click-to-call functionality.

Figure 2: “Benefits Finder” Platform: The Electricity Benefit

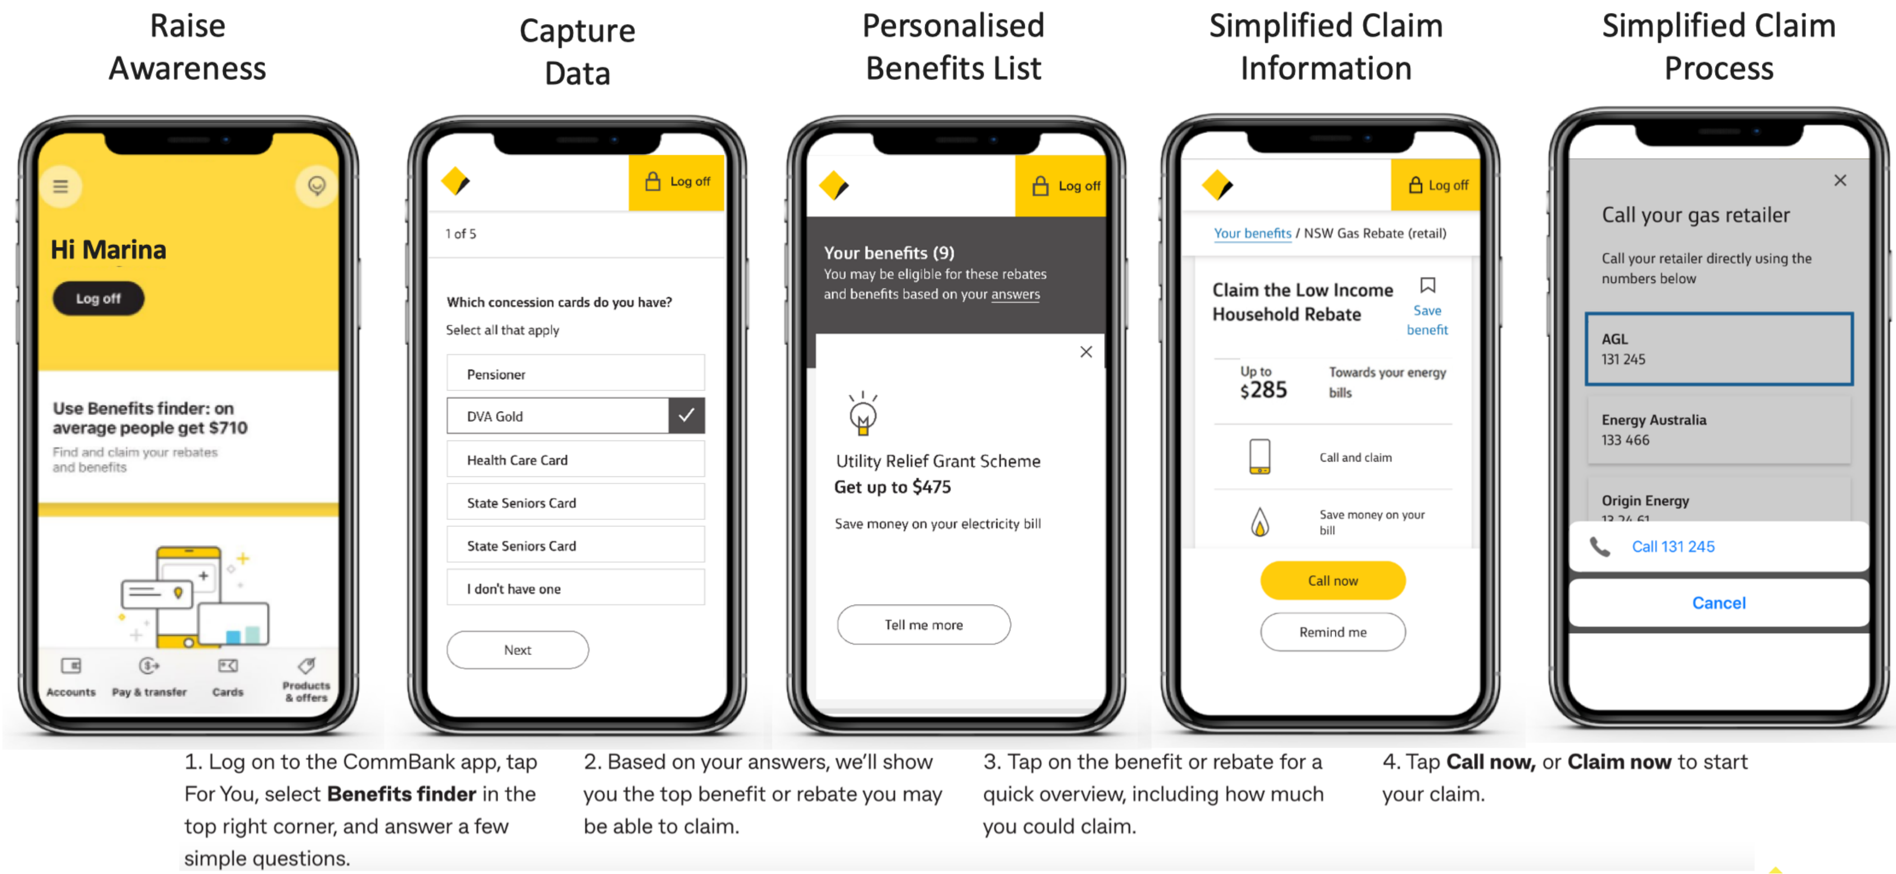

Financial Effects: Energy Bills

We focus on take-up of energy benefits, which include the NSW electricity benefit and its equivalents in other states as well as various gas discounts, and analyze financial impact in a sub-sample of 2,220,834 customers for whom the bank observes energy bills.[71] Because of the size of the dataset, we obtain data in the following format: average monthly energy bill amounts pre- and post-intervention, with the former based on the 6 months leading up to the intervention and the latter based on the 14 months after.

We test if the platform invitation had an effect on energy bills, using the specification:

$$EnergyBills_i = \beta Treat_i + \beta Covariates_i + \epsilon_i \tag{5}$$

where $EnergyBills_i$ is the average monthly bill amount post-intervention, $Treat_i$ the treatment status of getting an invitation for customer i, $Covariates_i$ pre-intervention covariates, and ϵ_{it} the error term.[72]

We find that getting invited to the platform lowers average monthly energy bills by $1.32, in our preferred specification.[73] (We find that actually visiting the platform lowers average monthly energy bills by $6.39, in our preferred specification, as seen in Table E2.)[74]

As seen in Table 9, across different specifications, the point estimates of the effect of getting invited are all negative and range between one and six dollars saved per month on energy. This monetary effect (of $1.32) is 2.27 times larger than the take-up effects observed in the first experiment (−1.316 vs. −0.580). This could be explained by the platform having not just one but several energy benefits, such as gas discounts; and this scale-up giving people more opportunities and a longer timeframe to use the simplified application pages.

[71] Bills are observed for 2,220,834 customers. 1,956,655 are in treatment (82.2 percent) and 423,833 in control (17.8 percent); this breakdown is similar to that in the whole study population (81.8 percent treatment and 18.2 percent control).

[72] We use robust standard errors clustered at the state-level, which is the largest administrative level the computational infrastructure in which we perform our analyses allows for.

[73] In column (4) we control for average energy bills in the pre-intervention period and all other pre-intervention covariates. We prefer this specification given the substantively small but slight imbalances between experimental groups (Table 8).

[74] $1.316 is the intention-to-treat (ITT) effect; which scaled by the difference in compliance rates yields a local average treatment effect (LATE) of $6.39 ($1.316 \div 0.206 = 6.388$).

Table 9: Effect of Platform Invitation on Energy Bill Amounts

	Dependent variable:			
	Energy Bill Amount			
	(1)	(2)	(3)	(4)
Treatment (Platform Invitation)	-6.329^{***}	-2.311^{***}	-2.887^{**}	-1.316^{**}
	(2.345)	(0.816)	(1.323)	(0.542)
Pre-Intervention Bill Average		0.821^{***}		0.794^{***}
		(0.033)		(0.028)
State Controls	Y	Y	Y	Y
Other Covariate Controls	N	N	Y	Y
Observations	2,220,834	2,220,834	2,220,834	2,220,834

This table shows results from a regression model of a continuous variable for average monthly energy bill amount post-intervention on a treatment status indicator (invited to the platform). This is the intention-to-treat effect (ITT). Standard errors are robust and clustered at the state-level. Columns (2) to (4) show additional controls: a customer's pre-intervention monthly bill average, and a vector of other pre-intervention controls: gender, age, years with the bank, concession card, small business, and bank products held (credit card, savings- and transaction accounts). Sample is restricted to customers with energy bills observable in bank data; outcomes are based on bill amounts for 6 months pre- and 14 months post- intervention. *p<0.1; **p<0.05; ***p<0.01

7 Discussion and Conclusion

Summary of Findings

The results presented in this paper suggest that small changes to the way citizens engage with a specific policy can change their attitudes toward government, at least short-term. Through field experiments—conducted in collaboration with the Commonwealth Bank of Australia, to leverage the data infrastructure and digital capacities of a large bank to increase take-up of government benefits—we uncover five main findings.

In experiment 1, we first find that low-income customers are highly interested in these types of innovations: with 18.2 percent of targeted individuals wanting more information about the electricity benefit, this was the bank's most popular messaging campaign to date, and 91 percent of survey respondents want similar information in the future. Second, reducing even small application burdens results in higher engagement: those randomly assigned to a simplified application page were 3 times as likely (14.7 percent versus 4.8 percent on the standard government-like page) to take observable steps toward claiming.

Third, analyzing bank data linked to government records, we find that the simplified treatment increases take-up by 2.1 percent relative to take-up in the standard group. While these effects might seem small in absolute terms, they are highly cost-effective. Fourth, though overall effects on take-up are modest, reducing burdens substantially increases satisfaction with all parties involved. Those with a simplified experience were 45.5 and 44.4 percent more likely to submit surveys rating the bank and government positively, respectively. Taken together, the first field experiment provides evidence that how the policy experience is designed—beyond content, and without large financial effects—matters for policy feedback.

Fifth, we show that the intervention can be successfully scaled-up: we find that during the first year, getting invited to a permanent "'Benefits Finder" platform featuring many simplified benefits, lowers average monthly energy bills by $1.32; and more broadly the platform generates better and simplified government experiences for over 2 million people.

Limitations and Challenges

There are, of course, limits to what we can learn from these experiments; we discuss them here and identify avenues for further research. First, we could only measure political attitudes with one survey item.[75] Is satisfaction with the government—in this context, the state government of NSW—a good enough proxy for more general attitudes toward government? We believe the measured satisfaction with the government is highly correlated with the theoretical concept of interest here—how citizens evaluate and feel about the government—but future research, using additional measures, can shed more light on which specific aspects of citizens' complex attitudes toward government tend to get affected by policy experiences.

Second, the absence of ground truth take-up data at the individual-level makes it challenging to disentangle the effect of getting the extra money (the actual discount) from the effect of the experience. The simplified page was designed to change both the experience and the likelihood of success in claiming the benefit. As such, the observed effects on attitudes could stem from either component. Given the design and data available, we cannot conclusively adjuducate between the following possible mehanisms: (i) getting the discount—either in practice but more likely in expectation[76]—makes people give more credit to the government; or (ii) the experience has an independent effect on attitudes—for example, people infer lessons about the workings of government from interacting with this specific policy; or (iii) some combination of both the discount and experience affect satisfaction.

[75]The bank does not usually, if ever, ask customers about their political attitudes. By design, the satisfaction with government survey question is otherwise identical to the satisfaction with bank question, which in turn mimics the structure of the Net Promoter Score (NPS) question routinely asked by the bank.

[76]Even those approved instantly have inevitable delay in actual receipt until the next (quarterly) bill.

Nonetheless, the evidence at hand points toward the interpretation that the experience itself, not just money, makes those on the simplified page more satisfied. For instance, survey respondents assigned to the simplified experience report additional satisfaction that is not explained by (self-reported) claiming. We interpret this as follows: even though those experiencing the simplified page are more satisfied overall with both the bank and government, their satisfaction appears to be driven to a lesser extent by the money than for respondents who had the standard government-like experience. Perhaps a less cumbersome experience counterbalances any dissatisfaction resulting from not getting the benefit.

Lastly, the first experiment demonstrates attitude effects in the relatively short-term. Attitudes were measured on average 8.4 days after treatment, but with submission of survey responses ranging from 0-63 days after experiencing the treatment (see Appendix Figure A7). Whereas some experimental treatment effects have shown to be short-lived, others, "particularly those with informational treatments, may have more enduring effects" (Keiser and Miller 2020, 145). Understanding how durable these effects are, and whether they would differ in T+1 and beyond, after repeated interactions facilitated by the bank, requires further research. But even if the effects of policy design are small in any one instance, they may operate cumulatively: if policy design shapes daily experiences across policy areas and over time, they could have a large total effect on citizens' overall attitudes toward government.

Revisiting the Role of For-Profits

These field experiments also shed light on the role private actors can play in attempts to shift administrative burdens away from vulnerable, typically hard-to-reach populations. Both directly, as in these interventions, and indirectly, by generating evidence on what can work more broadly. It is not common practice for banks to help customers navigate the social safety net.[77] Given the poor reputation of banks in many countries, some concerns are likely top-of-mind for many who hear about this initiative. It is worth discussing such concerns in turn, but also considering the potentially positive effects of for-profits reducing frictions between citizens and the state.

How do concerns about for-profits trying to shift administrative burdens apply to CBA's initiative to try to reduce frictions for its customers? First, does the bank have an interest in keeping claiming difficult, or making sure the government does not itself simplify the claiming process? In practice, CBA is now collaborating with the government on other efforts to reduce frictions for citizens qualifying for government support. In theory, a bank

[77]To the best of our knowledge, there are no interventions similar to the ones explored in this paper. However, with the recent onset of the Covid-19 crisis, banks in many countries are being asked to serve an intermediary role between small businesses and governments, to facilitate access to stimulus packages.

does not need to provide these types of services to be profitable, and to the extent they want their customers to have more money in the bank, it does not matter if the government eventually shifts all of the administrative burden itself, as long as take-up is enhanced.[78]

Second, is the bank using the existing barriers to make money? Unlike the tax preparation industry, there is no evidence that CBA is using the opportunity to simplify claiming to try to charge customers for unnecessary services.[79] Third, CBA does not appear to have incentives to selectively help people. The impact of efforts to increase take-up will only have unequal effects to the extent that the bank may have more information about certain groups than others, thus being better able to guess their eligibility.

In this particular context, concerns might also be raised about privacy, as well as whether the bank is taking credit for programs that the government is paying for. On privacy, besides the survey, the bank did not collect any additional data on its customers for this intervention. Rather, data it already possesses were put to use to help low-income customers. On credit, rather than making the government look bad, this intervention provides clear evidence that it also benefits from increased credit. Though more research is needed to understand how credit attribution might shift over time, the evidence thus far suggests banks can help 'grow the pie' of satisfaction and draw attention to the government's role providing a safety net.

Finally, is it surprising that banks are not already connecting their customers with government benefits in the US? Not when considering that American banks have largely stopped serving low-income people (Baradaran 2013). Whereas 98.9 percent of all Australians have a bank account, 25 percent of Americans are either unbanked or underbanked (FDIC 2017).[80] But, like in Australia, almost all Americans have a cellphone and most have smartphones (Pew 2019). One last policy implication then, is that the barrier to trying similar interventions in the US likely lies in citizens' lack of access to banks, not technology.

[78]This can be contrasted with companies that help citizens navigate burdens, but whose business idea revolves around making money on ads or from government contracts, which means they might go out of business if the government takes action that obviates the need for simplification by intermediaries.

[79]Of course, if you are of the opinion that banks always try to make money off its customers, then once the customers have more money in their accounts, there is more to take advantage of. However, that seems like a concern about banking practices more generally, which is unlikely to be resolved by discouraging interventions that might help people get government support.

[80]Underbanked means a household has a bank account but has to go outside of the regulated, government-sponsored banking system to a fringe provider for money orders, check cashing, or payday loans. Baradaran argues that given the unique features of banks, including the dependence on their sponsoring government, there is an implicit social contract between banks and the state (2013). In theory, the "public benefit" tests already incorporated in much of banking regulation could be enforced so that banking systems—if they want to keep their government supports—would need to be accessible to all Americans (Baradaran 2013).

References

AIHW (2017). *Australia's Health 2017: The Thirteenth Biennial Health Report*. The Australian Institute of Health and Welfare. https://perma.cc/3YYZ-77W4.

Alatas, V., R. Purnamasari, M. Wai-Poi, A. Banerjee, B. A. Olken, and R. Hanna (2016). Self-Targeting: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy 124*(2), 371–427.

Badger, E. and M. Sanger-Katz (2020). "Could you Manage as a Poor American?" New York Times, Jan 28. https://perma.cc/3P64-BKXA.

Baradaran, M. (2013). Banking and the Social Contract. *Notre Dame Law Review 89*, 1283.

Bartlett, S., N. R. Burstein, and M. S. Andrews (2004). *Food Stamp Program Access Study: Eligible Nonparticipants*. Economic Research Service Washington, DC.

Bertrand, M., S. Mullainathan, and E. Shafir (2004). A Behavioral-Economics view of Poverty. *American Economic Review 94*(2), 419–423.

Besley, T. and S. Coate (1992). Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs. *The American Economic Review 82*(1), 249–261.

Bettinger, E. P., B. T. Long, P. Oreopoulos, and L. Sanbonmatsu (2012). The Role of Application Assistance and Information in College Decisions: Results from the H&R Block FAFSA Experiment. *The Quarterly Journal of Economics 127*(3), 1205–1242.

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

Budget Direct Home Insurance (2019). "The Financial Strain on Australian's Pockets", October 2. https://perma.cc/422D-C3M9.

Canstar (2019). "How Much Do Australians Save On Average?", April 17. https://perma.cc/5247-KMQ5.

CBA (2020). "Demand Spikes for CBA's Benefits Finder as Aussies Seek Out Digital Welfare Support During Coronavirus Crisis", April 14. https://perma.cc/K54R-934L.

CBA (2021a). "Benefits Finder: Find Benefits & Rebates", Oct 29. https://perma.cc/EGQ5-L3LH.

CBA (2021b). "More than 1.5m Claims Commenced to Uncover Government Benefits and Rebates", Aug 27. https://perma.cc/GT83-P8Z2.

CBPP (2019). "Policy Basics: The Supplemental Nutrition Assistance Program (SNAP)", Center on Budget and Policy Priorities, June 25. https://perma.cc/BPK2-2W49.

Currie, J. (2004). The Take Up of Social Benefits. Technical report, National Bureau of Economic Research.

Currie, J. and F. Gahvari (2008). Transfers in Cash and in-Kind: Theory Meets the Data. *Journal of Economic Literature 46*(2), 333–83.

Cutler, D. M. and J. Gruber (1996). Does Public Insurance Crowd out Private Insurance? *The Quarterly Journal of Economics 111*(2), 391–430.

Dague, L. (2014). The Effect of Medicaid Premiums on Enrollment: A Regression Discontinuity Approach. *Journal of Health Economics 37*, 1–12.

Daminger, A., J. Hayes, A. Barrows, and J. Wright (2015). Poverty Interrupted: Applying Behavioral Science to the Context of Chronic Scarcity. In *Ideas*, Volume 42, pp. 1–49.

Daponte, B. O., S. Sanders, and L. Taylor (1999). Why Do Low-Income Households Not Use Food Stamps? Evidence from an Experiment. *Journal of Human Resources*, 612–628.

Deloitte (2018). “9 out of 10 Australian Citizens now Own a Smartphone”, Dec 18. `https://perma.cc/S8RX-QPKU`.

Deshpande, M. and Y. Li (2019). Who is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy 11*(4), 213–48.

DPIE (2019). *NSW Energy Rebates Summary Report*. NSW Department of Planning, Industry and Environment, Australia.

DPMC (2020). *Open Government Partnership Australia*. Department of the Prime Minister and Cabinet, Australia. `https://perma.cc/AAK5-LC6V`.

Edelman (2018). *2018 Edelman Trust Barometer-Australia Results*. Edelman. `https://perma.cc/RUE7-94DG`.

FDIC (2017). *FDIC National Survey of Unbanked and Underbanked Households*. Federal Deposit Insurance Corporation, Washington, DC. `https://perma.cc/WRQ4-ERFF`.

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

Herd, P., T. DeLeire, H. Harvey, and D. P. Moynihan (2013). Shifting Administrative Burden to the State: The Case of Medicaid Take-Up. *Public Administration Review 73*(s1), S69–S81.

Herd, P. and D. P. Moynihan (2018). *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation.

IRS (2019). “EITC Participation Rate by States”, Oct 8. `https://perma.cc/D2SA-C62P`.

Keiser, L. and S. Miller (2020). Does Administrative Burden Influence Public Support for Government Programs? Evidence from a Survey Experiment. *Public Administration Review 80*(1), 137–150.

Kiel, P. and J. Elliott (2020). “TurboTax and Others Charged at Least 14 Million Americans for Tax Prep that Should Have Been Free, Audit Finds”, ProPublica, Feb 5. `https://perma.cc/LK8X-LMJ9`.

Lerman, A. E. (2019). *Good Enough for Government Work: The Public Reputation Crisis in America (and What We Can Do to Fix It)*. University of Chicago Press.

Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). Poverty Impedes Cognitive Function. *Science 341*(6149), 976–980.

Mettler, S. (2018). *The Government-Citizen Disconnect*. Russell Sage Foundation.

Mettler, S. and J. Soss (2004). The Consequences of Public Policy for Democratic Ditizenship: Bridging Policy Studies and Mass Politics. *Perspectives on Politics 2*(1), 55–73.

Mullainathan, S. and E. Shafir (2013). *Scarcity: Why Having Too Little Means So Much*. Macmillan.

Nichols, A. and R. Zeckhauser (1982). Targeting Transfers Through Restrictions on Recipients. *The American Economic Review 72*(2), 372–377.

Pew (2019). “Mobile Fact Sheet”, June 25. `https://perma.cc/3NYF-FZSG`.

Reid, T. (2017). *A Fine Mess: A Global Quest for a Simpler, Fairer, and More Efficient Tax System*. Penguin.

Riley, T. (2017). “Startups Are Finally Taking On Food Stamps”, Wired, Sep 6. `https://perma.cc/YEJ4-DTA9`.

Rossin-Slater, M. (2013). WIC in Your Neighborhood: New Evidence on the Impacts of

Geographic Access to Clinics. *Journal of Public Economics 102*, 51–69.
Roy Morgan Research (2019). “NPS Rating of Big Four Banks Declines During Finance Royal Commission”, April 1. https://perma.cc/J6SM-7FZY.
Schneider, A. and H. Ingram (1993). Social Construction of Target Populations: Implications for Politics and Policy. *American Political Science Review 87*(2), 334–347.
Shanks-Booth, D. and S. Mettler (2019). The Paradox of the Earned Income Tax Credit: Appreciating Benefits But Not Their Source. *Policy Studies Journal 47*(2), 300–323.
Soss, J. (1999). Lessons of Welfare: Policy Design, Political Learning, and Political Action. *American Political Science Review 93*(2), 363–380.
Sunstein, C. R. (2018). *The Cost-Benefit Revolution*. MIT Press.
Thaler, R. H. and C. R. Sunstein (2021). *Nudge: The Final Edition*. Penguin.

Appendices

Appendix A Electricity Benefit: Design and Implementation

A.1 Intervention Design

Figure A3: Experimental Design

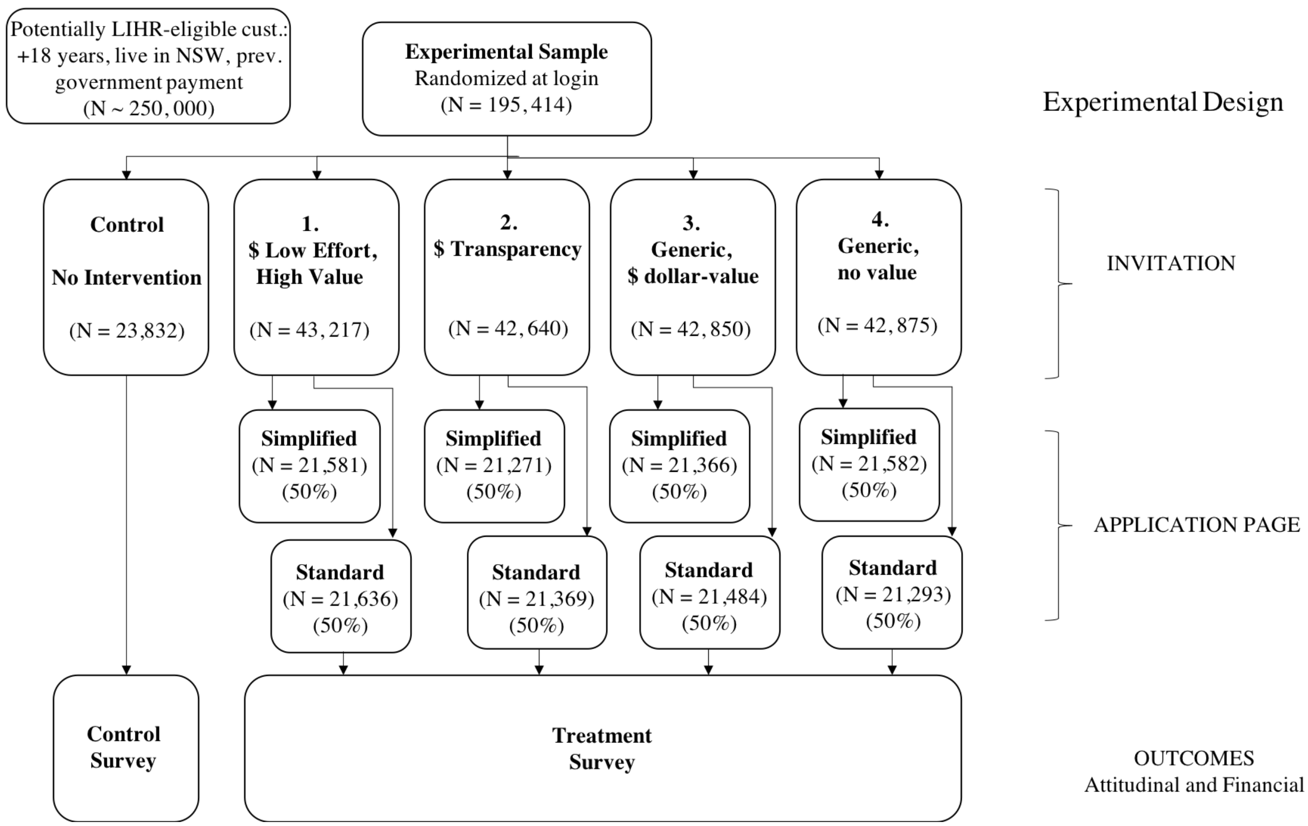

Figure A4: Timeline and Outcome Measurements Experiment 1

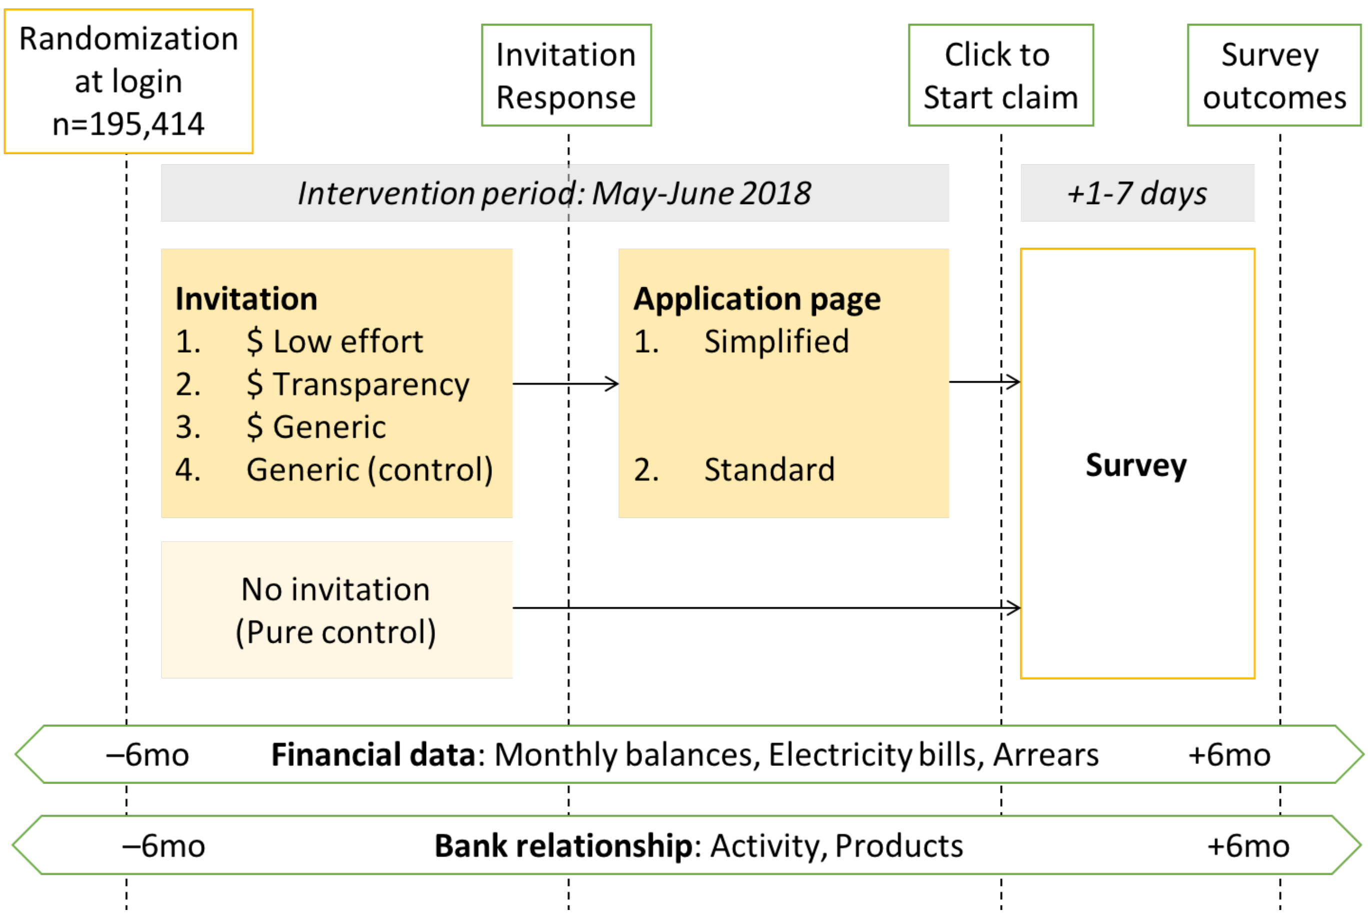

Figure A5: Informational Invitations

1: Low effort $ | 2: Transparency $ | 3: Generic $ | 4: Generic

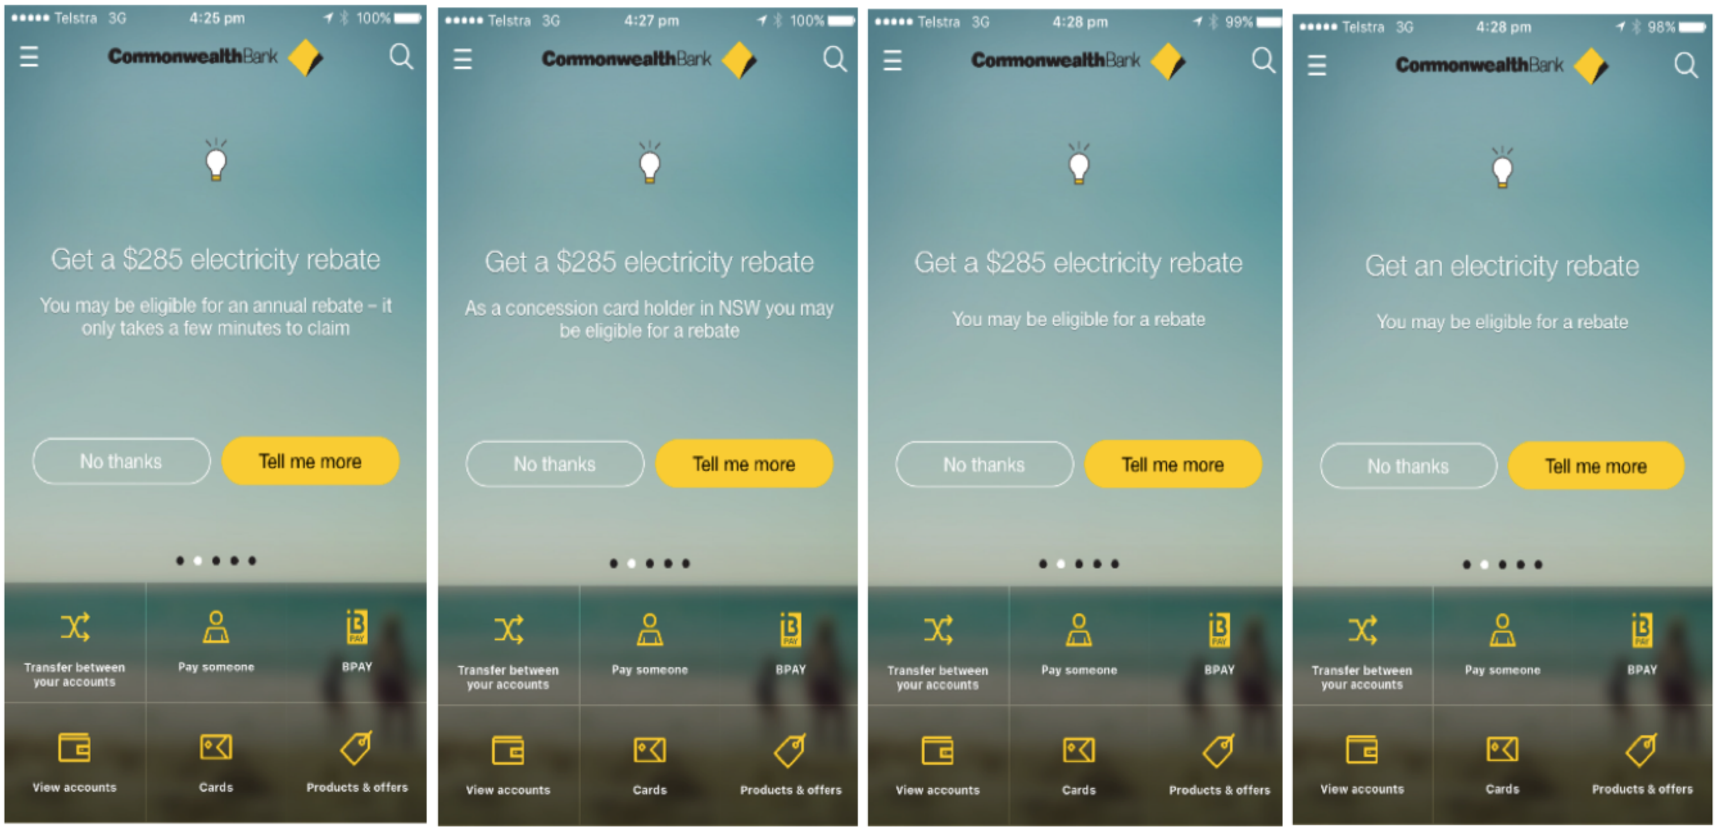

Figure A6: Application Pages

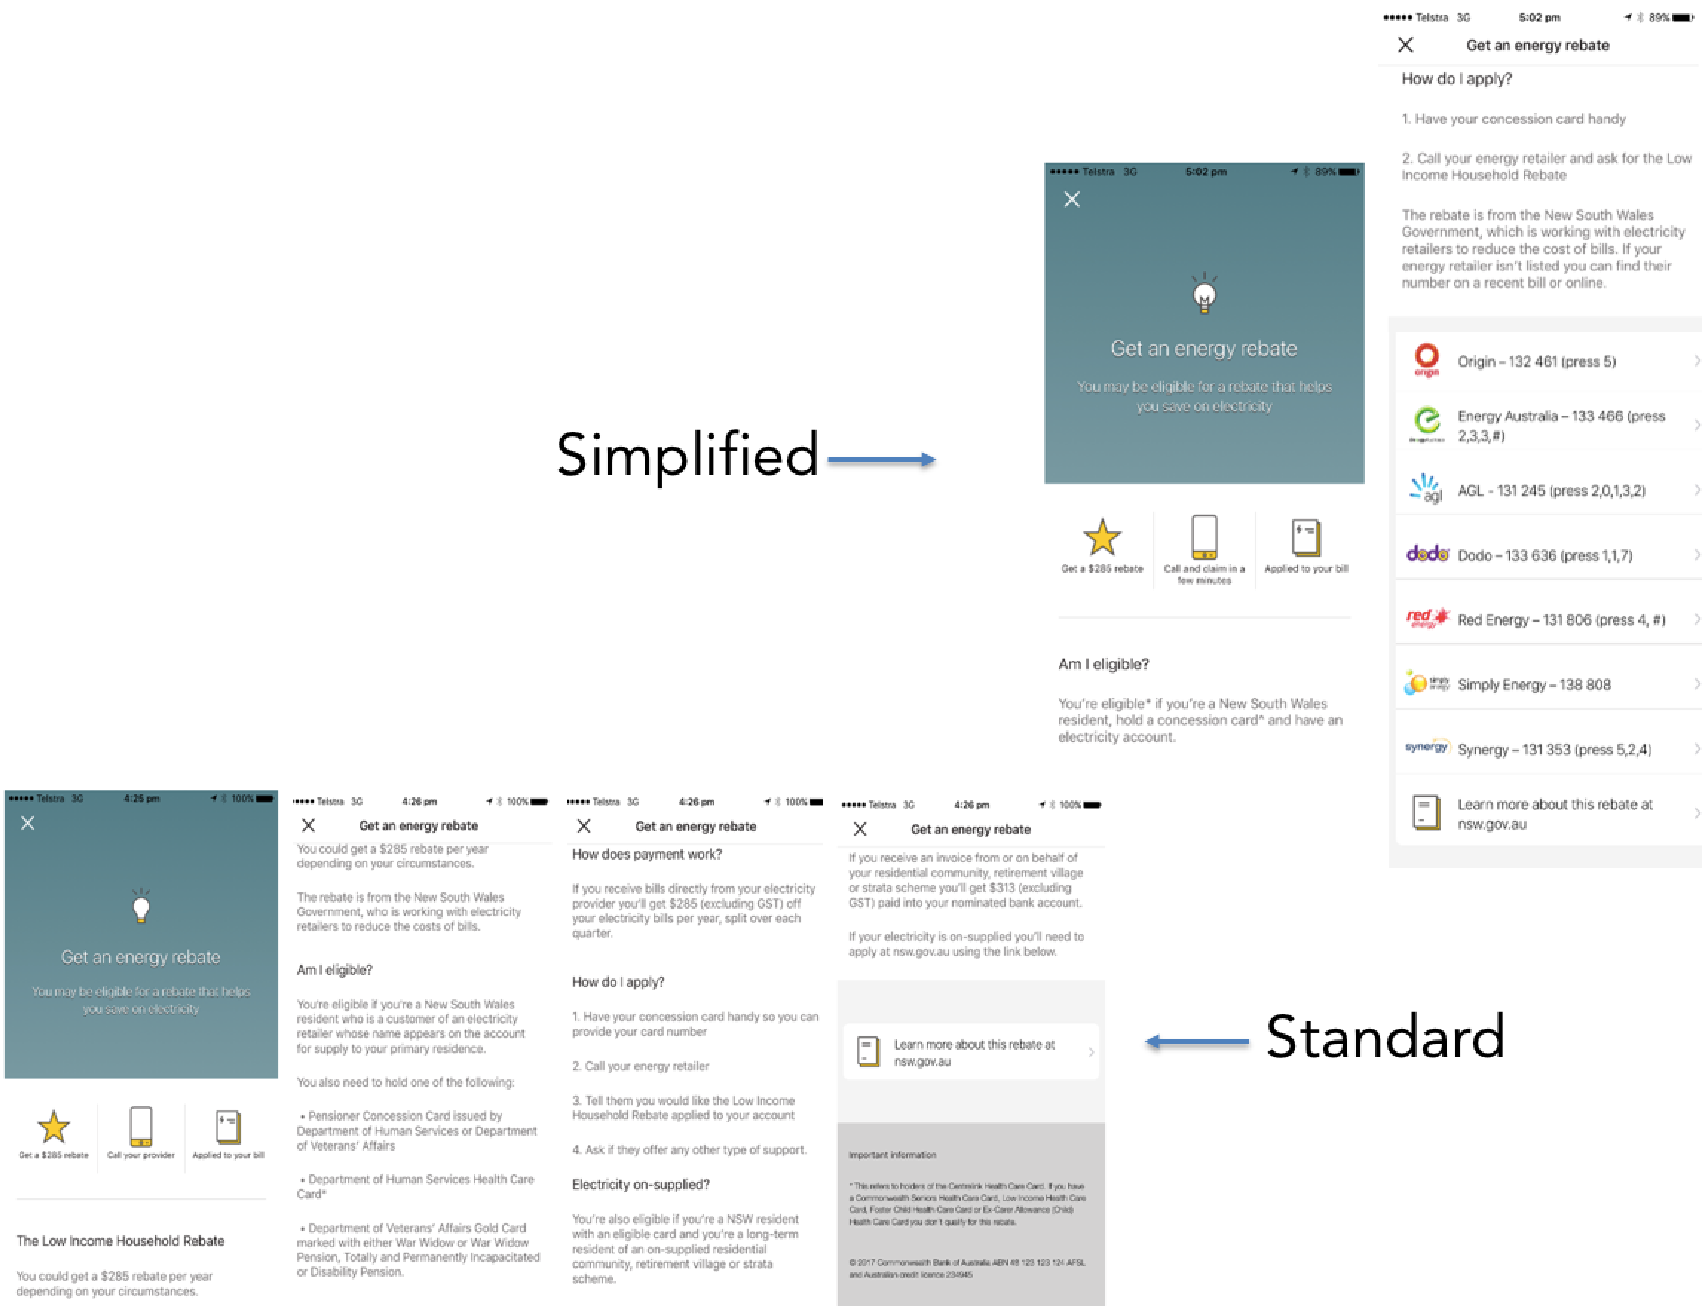

Figure A7: Days between Intervention and Survey Response

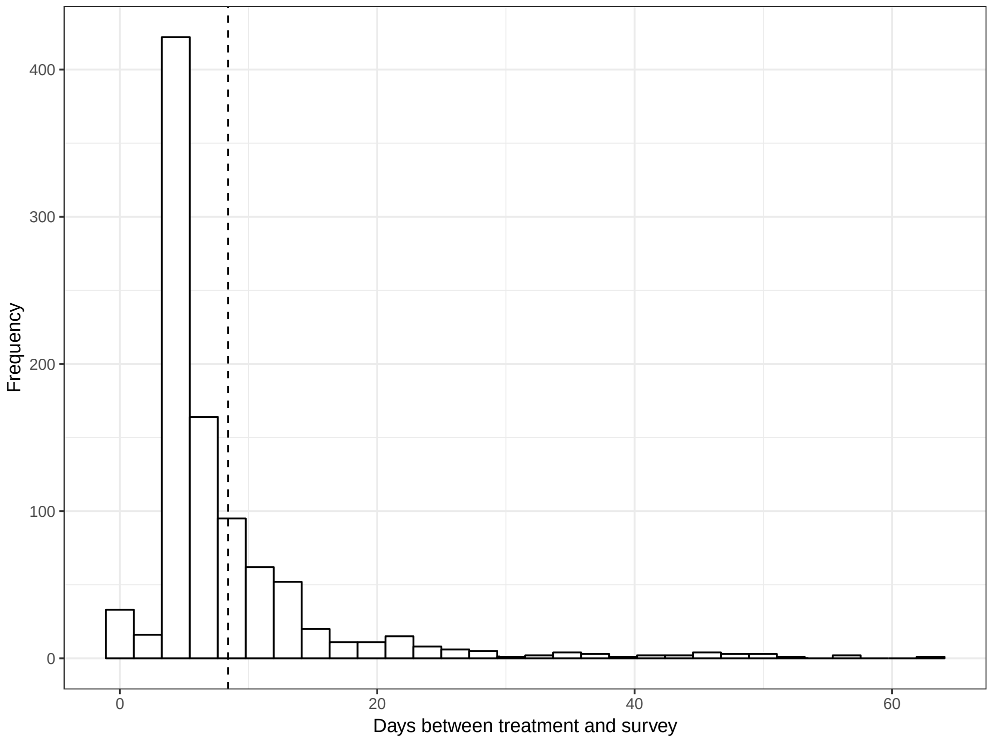

The mean number of days between intervention and submitting a survey response is 8.4 days (median=6) for the whole survey sample. Distributions are similar when broken down by treatment: mean for the Standard Page is 8.3 days, and 8.5 days for the Simplified Page.

A.2 Survey Design

Figure A8: Treatment Survey

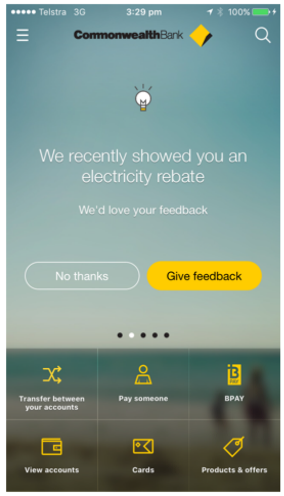

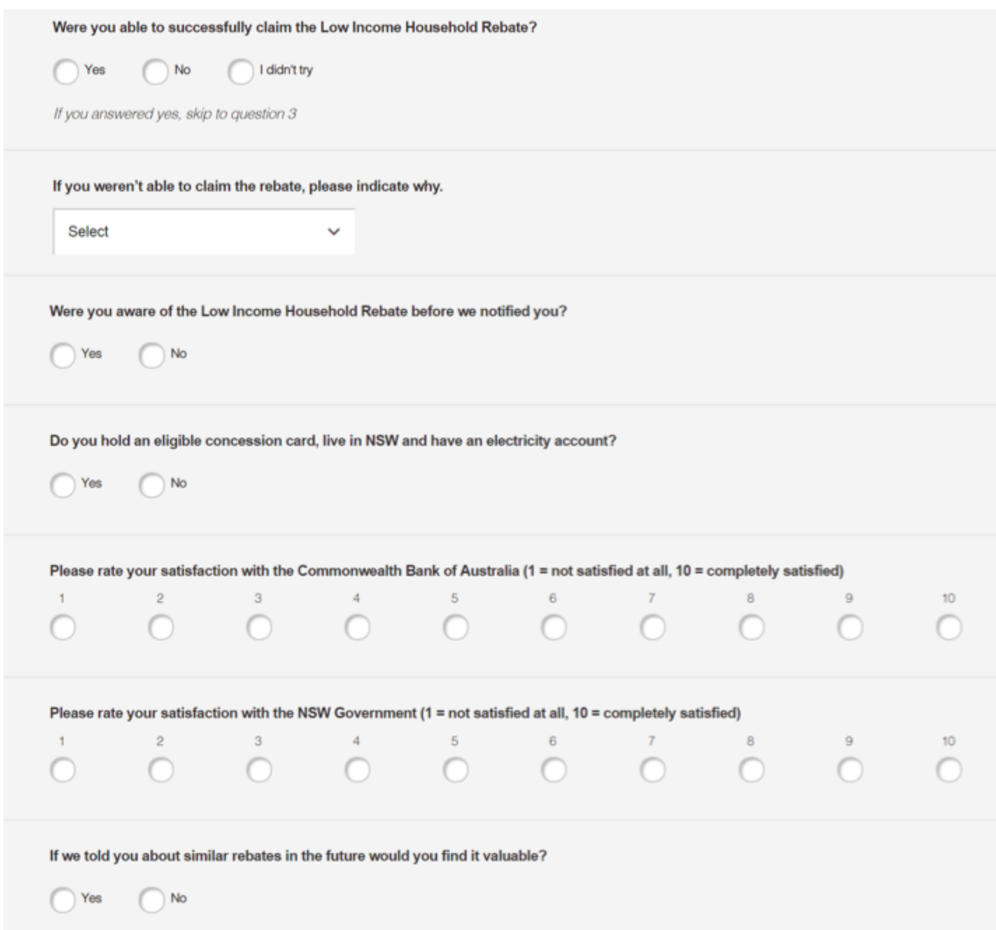

Were you able to successfully claim the Low Income Household Rebate?

Yes No I didn't try

If you answered yes, skip to question 3

If you weren't able to claim the rebate, please indicate why.

Select

Were you aware of the Low Income Household Rebate before we notified you?

Yes No

Do you hold an eligible concession card, live in NSW and have an electricity account?

Yes No

Please rate your satisfaction with the Commonwealth Bank of Australia (1 = not satisfied at all, 10 = completely satisfied)

1 2 3 4 5 6 7 8 9 10

Please rate your satisfaction with the NSW Government (1 = not satisfied at all, 10 = completely satisfied)

1 2 3 4 5 6 7 8 9 10

If we told you about similar rebates in the future would you find it valuable?

Yes No

Figure A9: Control Survey

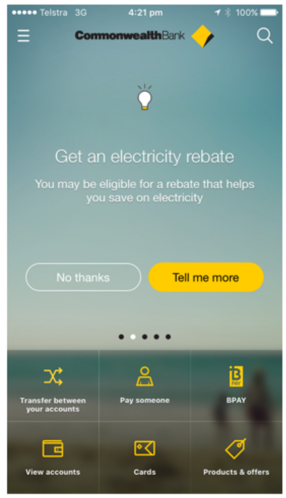

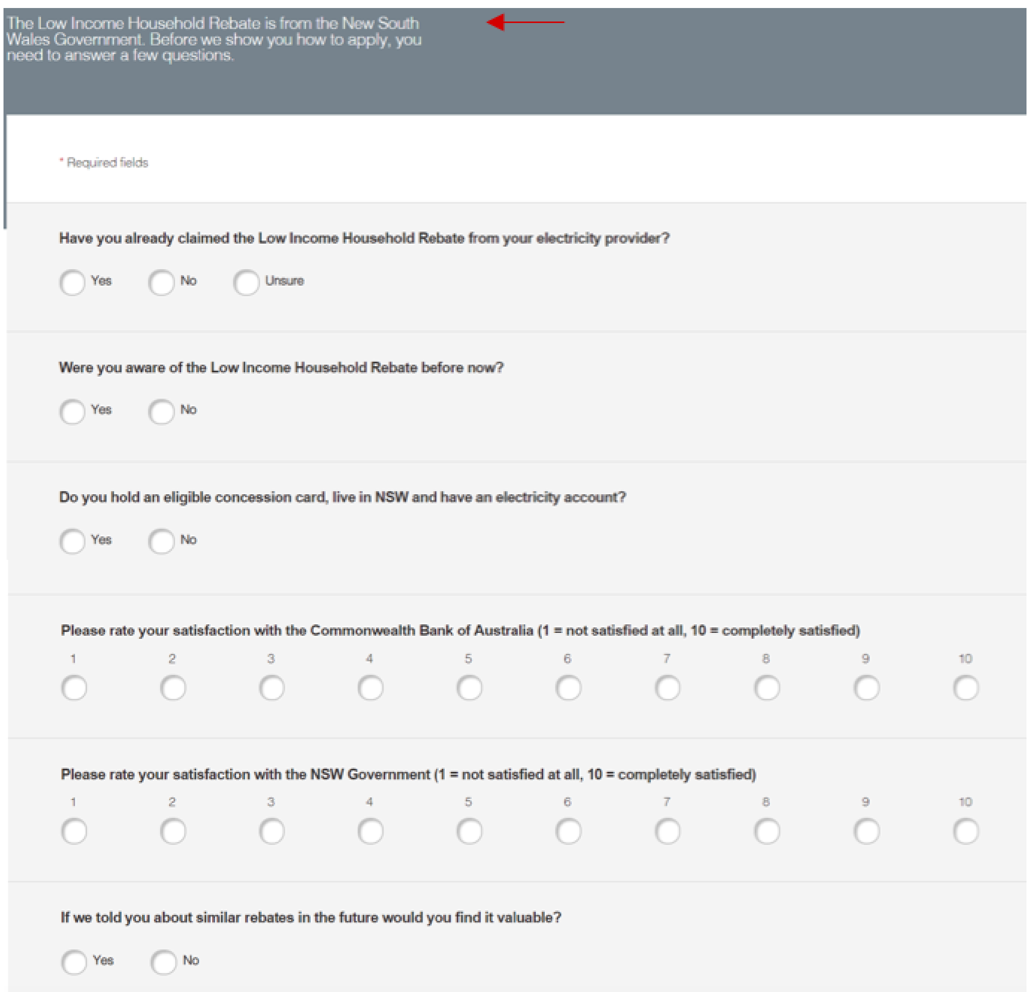

The Low Income Household Rebate is from the New South Wales Government. Before we show you how to apply, you need to answer a few questions.

* Required fields

Have you already claimed the Low Income Household Rebate from your electricity provider?

Yes No Unsure

Were you aware of the Low Income Household Rebate before now?

Yes No

Do you hold an eligible concession card, live in NSW and have an electricity account?

Yes No

Please rate your satisfaction with the Commonwealth Bank of Australia (1 = not satisfied at all, 10 = completely satisfied)

1 2 3 4 5 6 7 8 9 10

Please rate your satisfaction with the NSW Government (1 = not satisfied at all, 10 = completely satisfied)

1 2 3 4 5 6 7 8 9 10

If we told you about similar rebates in the future would you find it valuable?

Yes No

A.2.1 Control Survey Details

The control survey was sent to individuals assigned to the control group on a roll-out schedule matching the timing and volume of the treatment survey invitations. Unfortunately, the control survey responses cannot be used for comparisons in the main analyses, due to decisions taken during implementation.[81] The fact the treatment survey and control surveys are different is reflected in the vastly different response rates. Even though similar proportions of the entire treatment sample (90.4 percent) and control sample (91.1 percent) saw their survey invite, the response rates differ by a factor of 6: the response rate is 0.6 percent for the treatment survey and 3.6 percent for the control survey.[82]

Figure A9 demonstrates how both the control invite and the control survey itself resemble a light treatment condition. The invite looks more like the 'Generic' treatment invitation than the treatment survey invite, and there is no hint that this is a request for feedback. The added text in the top header, which tells individuals that "before we show you how to apply, you need to answer a few questions", is likely to induce response biases such as satisficing or just answering what you think the bank wants to hear.

[81]It is bank policy that surveys cannot be sent to customers too frequently, or to ask for general feedback—this is why the treatment survey invite includes "We recently showed you an electricity rebate", to indicate why the survey is being sent. Since control group indviduals had not previously been given any context about the electricity benefit, the wording had to be slightly different. But the wording that got final legal approval ended up sufficiently different from the treatment survey that it cannot be used as a control.

[82]Only those who login to the app again after the intervention period but before July 21 see the survey invite. The response rates are calculated based on the number of individuals who actually saw their survey invite; for the control survey this means 788/21,721.

Appendix B Electricity Benefit: Summary Statistics and Balance

Table B1: Characteristics of the Study Population in Experiment 1

	Median	Mean	SD	Min	Max	N
Demographics+						
Age	35	39.11	16.46	18	85	195,412
Years with bank	17	16.60	10.78	0	67	195,414
Female	1	0.59	0.49	0	1	195,414
Primary bank	1	0.85	0.36	0	1	195,380
History of arrears	0	0.13	0.34	0	1	195,414
History of hardship	0	0.00	0.02	0	1	195,414
Bank Products						
Credit card	0	0.29	0.52	0	6	195,197
Deposit & trans.	2	1.95	1.12	0	23	195,197
Lending & loan	0	0.07	0.27	0	5	195,197
Term deposit	0	0.04	0.25	0	12	195,197
Home loan	0	0.12	0.50	0	24	195,197
Other product	0	0.44	0.71	0	12	195,197
Account Balances ($)						
Monthly combined	395	-15,154	81,724	-400,155	87,928	186,729

This table shows pre-treatment characteristics. Panel 1 shows means (continuous variables) and proportions (binary variables); Panel 2 mean number of products; Panel 3 combined monthly balance for savings, transaction and credit card accounts (from Jan 2017–April 2018; winsorized at 2.5th and 97.5th percentiles).

Table B2: Characteristics of Study Population by Sub Treatment

	Standard				Simplified				Control
	Low Eff. $	Transp. $	Generic $	Generic	Low Eff. $	Transp. $	Generic $	Generic	Control
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Demographics+									
Age	39.3	39.3	39.0	39.1	39.1	39.2	39.2	39.1	39.1
Years with bank	16.7	16.8	16.6	16.6	16.8	16.6	16.7	16.6	16.5
Female	0.59	0.59	0.59	0.59	0.59	0.59	0.59	0.60	0.59
Primary bank	0.85	0.85	0.85	0.85	0.85	0.85	0.85	0.85	0.85
History of arrears	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13	0.13
History of hardship	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Bank Products									
Credit card	0.28	0.28	0.29	0.29	0.29	0.29	0.29	0.28	0.28
Deposit & trans.	1.94	1.95	1.95	1.96	1.96	1.94	1.95	1.95	1.94
Lending & loan	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07
Term deposit	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
Home loan	0.12	0.12	0.13	0.12	0.13	0.12	0.12	0.12	0.13
Other product	0.43	0.44	0.44	0.44	0.45	0.44	0.44	0.43	0.43
Account Balances ($)									
Monthly combined	-15,211	-14,438	-15,737	-14,610	-15,373	-14,925	-14,542	-15,339	-16,095
Observations (N)	21,636	21,369	21,484	21,293	21,581	21,271	21,366	21,582	23,832

Observations correspond to the sample of 195,414 bank customers included in the study. Each column represents one of the 9 experimental groups. Each row is a pre-treatment characteristic Each cell is the mean (or share) of the variable for that experimental group. Panel 1 shows means or shares, Panel 2 shows mean nr. of bank products, Panel 3 shows account balances, winsorized at the 0.5th and 99.5th percentiles. Variable dates: pre-treatment variables based on available data up to 6 months prior to intervention, which started on May 4, 2018. Primary bank, gender and postcode as of April 30, 2018. Bank products (Panel 2) as of May 4, 2018. Arrears and hardship recorded as 1 if any occurrence between Nov 1, 2017 and April 30, 2018. Age and years with bank as of Oct 2nd, 2018. Account balances (Panel 3) based on available monthly data between Jan 1, 2017 and April 30, 2018. Postcode is excluded but balanced.

Table B3: Balance: Yes Info vs. Not Interested

Pre-treatment Covariates:	**Want more info** (1)	**No thanks/Ignore** (2)	p-value (1) vs. (2)
Age	42.67	38.33	0.00
Years with bank	18.93	16.14	0.00
Female	0.61	0.59	0.00
Primary bank	0.87	0.84	0.00
History of arrears	0.14	0.13	0.00
History of hardship	0.00	0.00	0.36
Credit card	0.42	0.26	0.00
Deposit & transactions	2.12	1.91	0.00
Lending & loan	0.09	0.07	0.00
Term deposit	0.05	0.04	0.00
Home loan	0.18	0.11	0.00
Other product	0.55	0.41	0.00
Monthtly median	-22,877	-13,536	0.00
Monthly mean	-22,609	-13,332	0.00
Observations	31,299	140,283	

Table B4: Balance: Clicking Transparency Message vs. Generic $

Pre-treatment Covariates:	**Click Transparency** (1)	**Click Generic $** (2)	p-value (1) vs (2)
Age	39.21	39.08	0.27
Years with bank	16.64	16.64	1.00
Female	0.59	0.59	0.70
Primary bank	0.85	0.85	0.53
History of arrears	0.13	0.13	0.09
History of hardship	0.00	0.00	0.06
Credit card	0.29	0.29	0.50
Deposit & transactions	1.94	1.95	0.50
Lending & loan	0.07	0.07	0.54
Term deposit	0.04	0.04	0.84
Home loan	0.12	0.13	0.14
Other product	0.44	0.44	0.62
Monthtly median	-14,893	-15,355	0.41
Monthtly mean	-14,681	-15,141	0.41
Observations	42,640	42,850	

Table B5: Balance: Taking Steps Toward Claiming vs. Not

	Steps to Claim	**No Observable Action**	*p*-value (1) vs (2)
Pre-treatment Covariates:	(1)	(2)	
Age	45.04	42.42	0.00
Years with bank	19.75	18.85	0.00
Female	0.67	0.60	0.00
Primary bank	0.87	0.87	0.96
History of arrears	0.16	0.14	0.00
History of hardship	0.00	0.00	0.00
Credit card	0.40	0.42	0.15
Deposit & transactions	2.01	2.13	0.00
Lending & loan	0.09	0.09	0.69
Term deposit	0.05	0.05	1.00
Home loan	0.12	0.19	0.00
Other product	0.52	0.55	0.04
Monthtly median	-11,858	-24,054	0.00
Monthtly mean	-11,618	-23,784	0.00
Observations	3,041	28,258	

Table B6: Balance: Treatment Survey Respondents vs. Non-Respondents

	Submit Survey	**Rest of Sample**	*p*-value (1) vs (2)
Pre-treatment Covariates:	(1)	(2)	
Age	52.38	42.47	0.00
Years with bank	22.59	18.86	0.00
Female	0.64	0.61	0.10
Primary bank	0.87	0.87	0.61
History of arrears	0.12	0.14	0.18
History of hardship	0.00	0.00	0.51
Credit card	0.50	0.42	0.00
Deposit & transactions	2.10	2.12	0.67
Lending & loan	0.09	0.09	0.73
Term deposit	0.08	0.05	0.02
Home loan	0.17	0.18	0.63
Other product	0.63	0.55	0.01
Monthtly median	-15,952	-23,017	0.05
Monthtly mean	-15,631	-22,750	0.04
Observations	950	30,349	

This table shows the covariate balance between treatment survey respondents and non-respondents, among those who saw their randomly assigned application page (n=31,299).

Table B7: Treatment Survey Respondents vs. Control Survey Respondents

Pre-treatment Covariates:	**Treatment Survey** (1)	**Control Survey** (2)	p-value (1) vs (2)
Age	51.87	44.94	0.00
Years with bank	22.20	19.91	0.00
Female	0.63	0.62	0.62
Primary bank	0.88	0.87	0.75
History of arrears	0.13	0.13	0.96
History of hardship	0.00	0.00	0.90
Credit card	0.49	0.44	0.12
Deposit & transactions	2.11	2.14	0.61
Lending & loan	0.09	0.10	0.60
Term deposit	0.08	0.03	0.00
Home loan	0.16	0.18	0.42
Other product	0.60	0.56	0.37
Monthtly median	-16,524	-21,279	0.25
Monthtly mean	-16,197	-21,029	0.24
Observations	950	788	

Table B8: Simplified vs. Standard Survey Respondents

Pre-treatment Covariates:	**Simplified Survey** (1)	**Standard Survey** (2)	p-value (1) vs (2)
Age	51.86	51.87	0.99
Years with bank	22.16	22.26	0.89
Female	0.63	0.63	0.98
Primary bank	0.89	0.87	0.37
History of arrears	0.13	0.14	0.46
History of hardship	0.00	0.00	0.32
Credit card	0.50	0.48	0.63
Deposit & transactions	2.12	2.09	0.64
Lending & loan	0.09	0.08	0.44
Term deposit	0.08	0.08	0.95
Home loan	0.17	0.15	0.72
Other product	0.59	0.61	0.83
Monthtly median	-17,746	-15,120	0.65
Monthtly mean	-17,388	-14,829	0.65
Observations	508	442	

Appendix C Electricity Benefit: Robustness Checks and Extensions

C.1 Application Page Robustness

Table C1: Scrolling Down to the Read the Entire Application Page

	Dependent variable:
	Scrolling Down to Bottom of Page
Simplified	0.147***
	(0.006)
Standard (intercept)	0.339***
	(0.004)
Observations	31,299
R^2	0.022
Adjusted R^2	0.022
Residual Std. Error	0.487 (df = 31297)
F Statistic	709.449*** (df = 1; 31297)

This table shows results from a linear probability model of a binary outcome variable for scrolling down to the bottom of the appication page, on a treatment indicator for the simplified page, with the standard page as the omitted category. The sample includes all people who visited their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

Table C2: Time Spent on Application Page

	Dependent variable:
	Seconds on Page
Simplified	−6.207
	(5.141)
Standard (intercept)	86.517***
	(4.461)
Observations	3,023
R^2	0.0005
Adjusted R^2	0.0002
Residual Std. Error	121.921 (df = 3021)
F Statistic	1.457 (df = 1; 3021)

This table shows results from a linear probability model of a continuous outcome variable for the number of seconds spent on the appication page, on a treatment indicator for the simplified page, with the standard page as the omitted category. We can only measure time by the interval between clicking yes on the invitation and taking action on the application page. Therefore, the sample only includes the people who took any observable step toward claiming on their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

C.1.1 Does the Invitation Framing Matter for Claiming Behavior?

Does the invitation framing matter for behavior at the application page stage? There are several reasons to hypothesize that it might: the invitations contain different information, which could, for instance, attract different types of people, affect motivation to try to apply, or affect beliefs about eligibility. We test this with the following specification:

$$Action_i \quad = \quad \alpha + \beta Invitation_i + \epsilon_i \tag{6}$$

where $Action_i$ is the binary outcome, $Invitation_i$ indicates assignment to a treatment message for individual i, and α and ϵ_i represent the constant and error terms, respectively.

Table C3: Steps toward Claiming, by Invitation

	Dependent variable:
	Clicking 'Call' or 'Government Link'
Low Effort $	0.005
	(0.005)
Transparency $	0.012**
	(0.005)
Generic $	0.003
	(0.005)
Generic (intercept)	0.091***
	(0.004)
Observations	31,299
R^2	0.0002
Adjusted R^2	0.0001
Residual Std. Error	0.295 (df = 31295)
F Statistic	2.137* (df = 3; 31295)

This table shows results from a linear probability model of a binary outcome variable for taking any step toward claiming that we can observe in our application page environment—clicking either the government link (available on both pages) or the call button (simplified page only)—on treatment indicators for the four different invitations. The sample includes all people who clicked 'Yes' in response to the invitation and actually experienced their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

We find evidence that the 'Transparency $' invitation affects subsequent claiming behavior: as shown in Table C3, those who visit an application page as a result of showing interest in the Transparency invitation specifically, go on to take steps toward claiming at a higher rate than those who come to the application page via any of the other three invitations (about 13 percent, or 1.2 percentage points, higher than the Generic invitation, p=0.024).

One plausible interpretation is that the Transparency invitation works better for attracting indviduals who will actually go on to try to claim, because it provides more information pertinent to eligibility. Not only does the Transparency invitation give clues that the eligibility criteria will include being a resident in NSW and having a concession card—thus potentially filtering out non-eligibles up-front; getting this information might also make it easier for people who are already claiming the electricity benefit to identify it as such and realize that further information will not be useful for them. If already-claimers and ineligibles select out more when seeing the Transparency invitation, that would be consistent with the finding that a larger proportion of those who do want more information after seeing the Transparency invitation take steps toward claiming on the application page.

C.2 Bank Account Balances

We test whether the intervention had an effect on bank account balances, using the following specification:

$$AccountBalance_{it} = \alpha_i + \gamma_t + \beta Treat_{it} + \epsilon_{it} \tag{7}$$

where $AccountBalance_{it}$ is the account balance and $Treat_{it}$ the binary treatment status for customer i at month t, respectively. α_i and γ_t are customer and month fixed effects.

In our main analysis, we focus on the sub-sample of customers who experience their randomly assigned application page, since any effect of the simplified treatment should be driven by actually seeing the page. Below, we show that results are robust to different specifications.[83] We also restrict the sample to customers with less extreme levels of debt, to increase precision.[84] Table C4 shows the results of examining account balances, for this restricted sample, for a period of 4 months before and 4 months after the intervention (but different time windows lead to similar conclusions, as seen in Table C5).

Table C4: Effect of Simplified Treatment on Account Balances

	Dependent variable:
	Monthly Balance
Simplified	0.989
	(71.267)
Customer Fixed Effects	Y
Month Fixed Effects	Y
Standard Page Mean	10,595.639
SE	(145.202)
Observations	194,079
Customers	24,278
Time Window	8 months
R^2	0.000
F Statistic	0.0002

This table shows results from a fixed effects regression model of a continuous variable for the monthly combined balance (all savings, transactions and credit card accounts), on a treatment status indicator for the simplified page (vs. standard), and customer and month fixed effects. The sample includes balances for 4 months pre and 4 months post intervention, and is restricted to customers who experienced their randomly assigned treatment page, and have less than $10,000 in pre-treatment debt. *p<0.1; **p<0.05; ***p<0.01

[83] In Table C6, we show that results are statistically unchanged, and still noisy, when comparing pure control group balances with those of any treatment group (standard or simplified).

[84] To limit the influence of customers with large amounts of debt, and make the estimates more precise, we restrict the sample to the 84 percent of customers who have less than $10,000 in pre-treatment debt.

We find no significant difference between simplified and standard treatment group balances. Critically, Table C4 reveals that the account balance data for this population were much noisier than expected.[85] Even in the most extreme scenario—where the intervention causes full take-up in the simplified group (an extra $24 per month for the modal beneficiary) and no take-up in the standard group—we would likely be unable to detect the effect.[86]

Simplified Versus Standard Treatment

Table C5: Effect of Simplified Treatment on Account Balances (Time Windows)

	Dependent variable:		
	Monthly Balance		
	(1)	(2)	(3)
Simplified	69.273	0.989	−15.963
	(81.322)	(71.267)	(66.319)
Customer Fixed Effects	Y	Y	Y
Month Fixed Effects	Y	Y	Y
Standard Page Mean	10,598.710	10,595.639	10,605.400
SE	(205.68)	(145.202)	(118.340)
Observations	97,075	194,079	290,877
Customers	24,277	24,278	24,279
Time Window	4 months	8 months	12 months
R^2	0.00001	0.000	0.00000
F Statistic	0.726	0.0002	0.058

This table shows results from a fixed effects regression model of a continuous variable for the monthly combined balance (all savings, transactions and credit card accounts), on a treatment status indicator for the simplified page (vs. standard), and customer and month fixed effects. Columns differ by the time window around intervention: (1) includes balances for 2 months before and 2 months after, (2) for 4 pre and 4 post, and (3) for 6 pre and 6 post intervention. The sample is restricted to customers who experienced their randomly assigned treatment page, and have less than $10,000 in pre-treatment debt. *p<0.1; **p<0.05; ***p<0.01

Any Treatment Versus Control

[85] No previous trial had been run on these bank data in this kind of population.

[86] Even after restricting the sample to customers with less extreme balances and including customer and month fixed effects, as seen in Table C4 the standard error (71) is larger than twice the maximum expected average monthly effect size (24×2).

Table C6: Effect of Treatment on Account Balances (Time Windows)

	Dependent variable:		
	Monthly Balance		
	(1)	(2)	(3)
Any Treatment (vs. control)	30.902	9.095	14.675
	(42.902)	(36.606)	(33.820)
Customer Fixed Effects	Y	Y	Y
Month Fixed Effects	Y	Y	Y
Control Mean	7,498.36	7,492.36	7,519.92
SE	(134.65)	(95.21)	(77.83)
Observations	653,521	1,306,612	1,958,252
Customers	163,431	163,448	163,450
Time Window	4 months	8 months	12 months
R^2	0.00000	0.00000	0.00000
F Statistic	0.519	0.062	0.188

This table shows results from a fixed effects regression model of a continuous outcome variable for the monthly combined balance (all savings, transactions and credit card accounts), on a treatment status indicator (i.e. any invitation or page but not pure control), and customer and month fixed effects. Columns differ by the time window: (1) includes balances for 2 months before and 2 months after intervention, (2) for 4 pre and 4 post, and (3) for 6 pre and 6 post intervention. Sample is restricted to customers with less than $10,000 in pre-treatment debt. *p<0.1; **p<0.05; ***p<0.01

C.3 Electricity Bill Robustness

Table C7: Effect of Simplified Treatment on Electricity Bills (Restricted Sample)

	Dependent variable:			
	Bill Amount			
	$\pm 4mo$	$-4, +6mo$	$-4, +8mo$	$-4, +18mo$
	(1)	(2)	(3)	(4)
Simplified	−0.181	−0.032	−0.580	−0.643
	(1.824)	(1.702)	(1.608)	(1.465)
Customer Fixed Effects	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y
Standard Page Mean	180.599	180.599	180.599	180.599
SE	(1.070)	(1.085)	(1.078)	(1.075)
Observations	47,175	57,027	67,594	123,899
Customers	17,504	18,781	20,061	25,886
Time Window	8 months	10 months	12 months	22 months
R^2	0.00000	0.000	0.00000	0.00000
Adjusted R^2	−0.590	−0.491	−0.422	−0.264
F Statistic	0.010	0.0004	0.130	0.193

This table shows results from a fixed effects regression model of a continuous variable for electricity bill amounts on a simplified treatment status indicator, and customer and month fixed effects. Columns differ by the time window around intervention: (1) includes bills for 4 months pre and 4 months post, (2) for 4 months pre and 6 post, (3) for 4 months pre and 8 post, and (4) for 4 months pre and 18 months post intervention. **The sample is restricted to customers with bills from the 7 electricity retailers listed on the simplified page, and bill amounts up to $500 per month.** *p<0.1; **p<0.05; ***p<0.01

Table C8: Effect of Simplified Treatment on Electricity Bills (Unrestricted Sample)

	Dependent variable:			
	Bill Amount			
	$\pm 4mo$	$-4, +6mo$	$-4, +8mo$	$-4, +18mo$
	(1)	(2)	(3)	(4)
Simplified	1.697	−0.311	−0.769	−0.101
	(2.678)	(2.537)	(2.381)	(2.133)
Customer Fixed Effects	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y
Standard Page Mean	272.847	272.847	272.847	272.847
SE	(1.645)	(1.707)	(1.689)	(1.674)
Observations	87,648	106,841	127,215	236,175
Customers	30,718	32,683	34,506	42,554
Time Window	8 months	10 months	12 months	22 months
R^2	0.00001	0.00000	0.00000	0.000
Adjusted R^2	−0.540	−0.441	−0.372	−0.220
F Statistic	0.402	0.015	0.104	0.002
DF	(1; 56922)	(1; 74148)	(1; 92697)	(1; 193599)

This table shows results from a fixed effects regression model of a continuous variable for electricity bill amounts on a simplified treatment status indicator, and customer and month fixed effects. Columns differ by the time window around intervention: (1) includes bills for 4 months pre and 4 months post, (2) for 4 months pre and 6 post, (3) for 4 months pre and 8 post, and (4) for 4 months pre and 18 months post intervention. The sample includes customers with bills from all electricity retailers, and bill amounts up to $2,000. *p<0.1; **p<0.05; ***p<0.01

Table C9: Effect of Treatment on Electricity Bills (Restricted Sample)

	Dependent variable:			
	Bill Amount			
	$\pm 4mo$	$-4, +6mo$	$-4, +8mo$	$-4, +18mo$
	(1)	(2)	(3)	(4)
Any Treatment (vs. control)	0.258	0.032	−1.778	−0.998
	(2.565)	(2.396)	(2.265)	(2.068)
Customer Fixed Effects	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y
Control Mean	178.779	178.778	178.778	178.779
SE	(1.997)	(2.025)	(2.013)	(2.006)
Observations	53,890	65,129	77,200	141,584
Customers	19,959	21,385	22,842	29,537
Time Window	8 months	10 months	12 months	22 months
R^2	0.00000	0.000	0.00001	0.00000
Adjusted R^2	−0.589	−0.489	−0.420	−0.264
F Statistic	0.010	0.0002	0.616	0.233

This table shows results from a fixed effects regression model of a continuous outcome variable for electricity bill amounts, on a treatment status indicator (i.e. any invitation or page but not Pure Control), and customer and month fixed effects. Columns differ by the time window around intervention: (1) includes bills for 4 months pre and 4 months post, (2) for 4 months pre and 6 post, (3) for 4 months pre and 8 post, and (4) for 4 months pre and 18 months post intervention. The sample is restricted to customers with bills from the 7 electricity retailers listed on the simplified page, and bill amounts up to $500. *p<0.1; **p<0.05; ***p<0.01

Note 1: the analysis of the effect of treatment (any treatment vs. control) in an unrestricted sample of electricity bills yields more imprecise estimates, and is not shown here.

Note 2: as explained above, we cannot perfectly classify transactions as electricity charges, but create a sample of bills for 64,036 customers based on transaction details such as the name of the company, amounts and frequency. In the main electricity bill amount analyses, to reduce noise, we first restrict the full bills sample to as sample of 48,513 customers with bills smaller than $2,000 and no more than 20 bills in this period. This serves to further hone in on actual electricity bills, which tend to be smaller and quarterly.

C.4 Satisfaction Robustness

Table C10: Satisfaction by Application Page (Full Sample)

	Dependent variable:			
	Positive Rating:		Negative Rating:	
	Bank	**Government**	Bank	Government
	(1)	(2)	(3)	(4)
Simplified	0.001***	0.001***	0.00005	0.0003
	(0.0003)	(0.0002)	(0.0002)	(0.0002)
Standard (intercept)	0.003***	0.002***	0.001***	0.001***
	(0.0002)	(0.0001)	(0.0001)	(0.0001)
Observations	171,582	171,582	171,582	171,582
R^2	0.0001	0.0001	0.00000	0.00001
Adjusted R^2	0.0001	0.0001	−0.00001	0.00001
Residual Std. Error (df = 171580)	0.054	0.044	0.032	0.038
F Statistic (df = 1; 171580)	9.880***	10.525***	0.092	2.488

This table shows results from a linear probability model of binary outcome variables for submitting a survey that rates the bank (1)(3) or government (2)(4) positively or negatively—on a treatment indicator for the simplified page, with the standard page as the omitted category. Here, positive (negative) means a rating above (below) the mean satisfaction score, but results are robust to other specifications. Sample includes all assigned to a treatment group. *p<0.1; **p<0.05; ***p<0.01

C.5 Other Survey Evidence

Figure C1 shows—among the 950 survey respondents—the distribution of self-reported reasons for not managing to claim the rebate. ('NA' means the survey respondent either did not select a reason, or said they did manage to successfully claim the rebate). It is important to keep in mind that we cannot assume that these survey respondents are perfectly representative of the study population. See Table B6 on how the survey respondents differ from non-respondents on observables.

Figure C1: Self-Reported Reasons for Not Managing to Claim

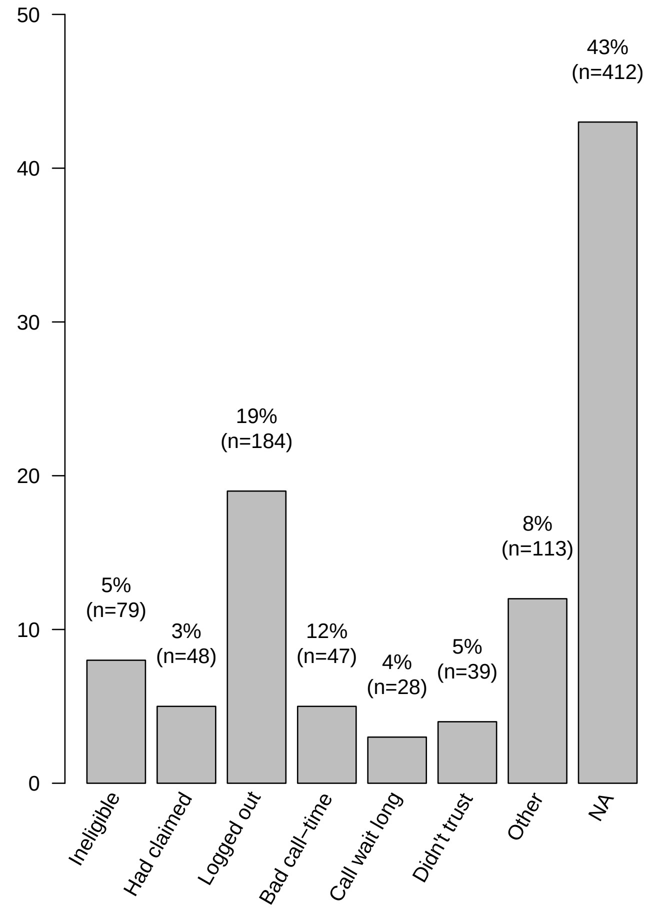

Beyond the subjective measures of satisfaction already discussed, the survey asked three more factual questions, originally included to try to validate self-reports of eligibility, awareness, and take-up. Unfortunately, we do not have objective take-up data at the individual level from government or electricity retail partners, and have no way of verifying how well answers to these questions represent reality. Figure C2 shows the proportion of all respondents who answered these questions affirmatively.

Figure C2: Other Survey Outcomes

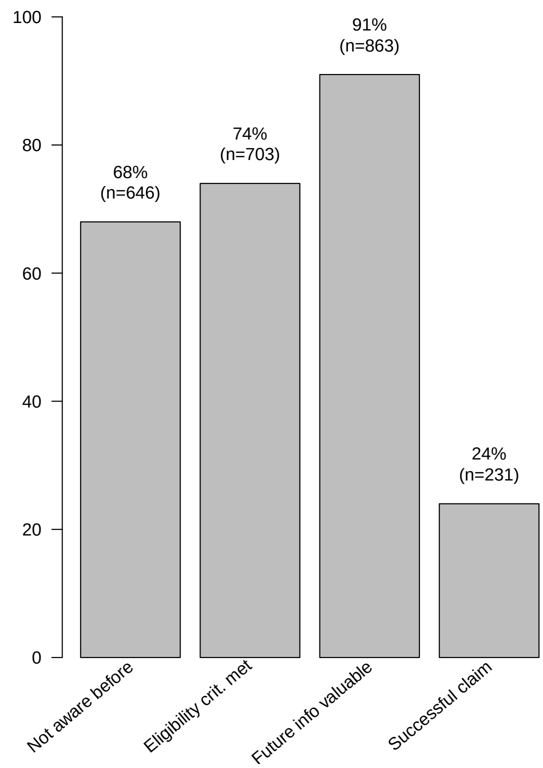

C.5.1 Do People Want Similar Information in the Future?

The survey question: "If we told you about similar rebates in the future, would you find it valuable?", is another measure of satisfaction. As Table C11 shows, the answers provide more evidence that experiencing the simplified application page increased satisfaction. The simplified application page made individuals more likely to submit a survey saying they would find future invitations valuable (but *not* more likely to submit a survey saying they would not find future information valuable).

Table C11: Effect on Wanting Similar Information in the Future

	Dependent variable:	
	Submit Survey Saying:	
	Future Info Valuable	Not Valuable
	(1)	(2)
Simplified	0.006***	0.00000
	(0.002)	(0.0002)
Standard (intercept)	0.016***	0.0004***
	(0.001)	(0.0002)
Observations	31,299	31,299
R^2	0.0005	0.000
Adjusted R^2	0.0005	−0.00003
Residual Std. Error (df = 31297)	0.135	0.021
F Statistic (df = 1; 31297)	15.436***	0.0002

This table shows results from a linear probability model of binary outcome variables for submitting a survey (1) saying one would find similar info valuable in the future, or (2) saying one would not find it valuable—on a treatment indicator for the simplified page, with the standard page as the omitted category. The sample includes all people who visited their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

C.5.2 Self-Reported Claiming

Effect of Simplified Treatment on Self-Reported Claiming

This section includes analyses of survey questions about objective outcomes, such as whether a respondent successfully claimed the electricity benefit. We cannot validate these answers, so they are harder to interpret and should be done so with caution. (In contrast, outcomes such as satisfaction, or whether someone wants information in the future, are inherently subjective questions where we are interested in people's choice to provide a specific answer).

First, respondents assigned to the simplified page are not significantly more likely to say that they managed to successfully claim the electricity benefit. This is true both when looking among respondents only (Table C12) and when comparing—among all those who experienced their randomly assigned application page—the choice to submit a survey self-reporting claiming (Table C13). As seen in Table C13, simplified respondents are both more likely to submit a survey saying they claimed *and* more likely to submit a survey saying they did not claim. This merely reflects the fact that those assigned to the simplified experience are more likely to submit a survey.

Table C12: Effect of Simplified Treatment on Self-Reported Claiming

	Dependent variable:
	Successfully Claimed (self-report)
Simplified	0.036
	(0.028)
Standard (intercept)	0.224^{***}
	(0.020)
Observations	950
R^2	0.002
Adjusted R^2	0.001
Residual Std. Error	0.429 (df = 948)
F Statistic	1.651 (df = 1; 948)

This table shows results from a linear probability model of a binary outcome variable for saying one successfully managed to claim the electricity benefit, on a treatment indicator for the simplified page, with the standard page as the omitted category. The sample includes all survey respondents. *p<0.1; **p<0.05; ***p<0.01

Table C13: Choosing to Self-Report Take-Up, After Experiencing an Application Page

	Dependent variable:	
	Submit Survey Saying:	
	Yes, Successfully Claimed	**No, Did Not Claim**
	(1)	(2)
Simplified	0.002**	0.003***
	(0.001)	(0.001)
Standard (intercept)	0.005***	0.003***
	(0.001)	(0.001)
Observations	31,299	31,299
R^2	0.0002	0.0004
Adjusted R^2	0.0001	0.0004
Residual Std. Error (df = 31297)	0.077	0.064
F Statistic (df = 1; 31297)	5.390**	13.584***

This table shows results from a linear probability model of binary outcome variables for submitting a survey (1) saying one successfully managed to claim the benefit, or (2) saying one did not manage to claim the benefit—on a treatment indicator for the simplified page, with the standard page as the omitted category. The sample includes all people who visited their randomly assigned application page. *p<0.1; **p<0.05; ***p<0.01

Satisfaction and Reported Claiming in Survey Respondents, by Page

The relationship between self-reported claiming and satisfaction is consistently stronger among the survey respondents assigned to the standard page (Table C14). This suggests that in order to be satisfied as a result of interacting with the standard page, actually being successful in getting the money plays a larger role than it does for those experiencing the simplified page.

Table C14: Relationship between Claiming and Satisfaction, by Page

	Dependent variable: positive rating			
	Bank Satisfaction		**Government Satisfaction**	
	Simplified	Standard	Simplified	Standard
	(1)	(2)	(3)	(4)
Claimed (self-report)	0.195***	0.321***	0.167***	0.293***
	(0.050)	(0.054)	(0.049)	(0.051)
Intercept	0.441***	0.315***	0.340***	0.242***
	(0.025)	(0.025)	(0.025)	(0.024)
Observations	508	442	508	442
R^2	0.029	0.076	0.023	0.070
Adjusted R^2	0.027	0.074	0.021	0.068
Degrees of Freedom	506	440	506	440
F Statistic	15.231***	36.055***	11.763***	33.234***

This table shows results from a linear probability model of binary outcome variables for rating the bank (1)(2) or government (3)(4) positively—on an indicator for saying one successfully managed to claim the benefit. Columns (1) and (3) show results for respondents assigned to the simplified page, and columns (2) and (4) show results for respondents assigned to the standard page. As before, positive means a rating above the mean satisfaction score, but results are robust to other specifications. The samples include all survey respondents. *p<0.1; **p<0.05; ***p<0.01

Appendix D National Scale-Up: Design and Implementation

D.1 Intervention Design

Figure D1: Timeline and Outcome Measurements Experiment 2

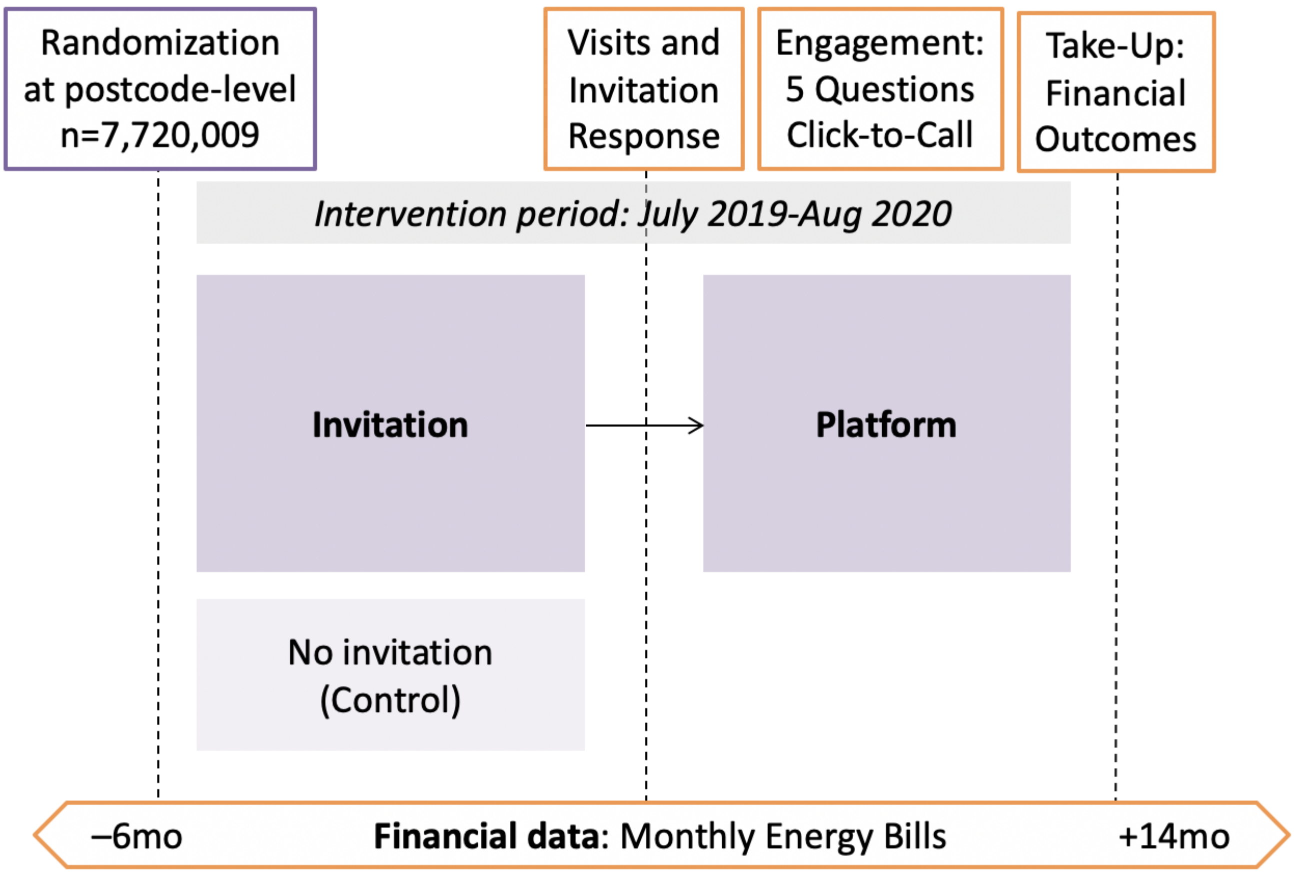

Appendix E National Scale-Up: Robustness Checks and Extensions

E.1 Financial Results

Table E1: Effect of Platform Invitation on Energy Bill Amounts (Full Controls)

	Dependent variable:			
	Energy Bills			
	(1)	(2)	(3)	(4)
Treatment	−6.329***	−2.311***	−2.887***	−1.316***
	(2.345)	(0.816)	(1.323)	(0.542)
Pre-Intervention Bill Average		0.821***		0.794***
		(0.033)		(0.028)
NSW	−73.320***	−39.268***	−68.486***	−38.683***
	(0.844)	(0.579)	(0.811)	(0.574)
NT	−128.603***	−73.249***	−116.275***	−70.613***
	(1.716)	(1.178)	(1.650)	(1.166)
QLD	−142.501***	−73.787***	−136.464***	−73.843***
	(0.863)	(0.594)	(0.830)	(0.588)
SA	−91.534***	−31.164***	−80.757***	−29.438***
	(0.979)	(0.673)	(0.942)	(0.667)
TAS	−72.530***	−25.335***	−61.311***	−22.499***
	(1.076)	(0.739)	(1.034)	(0.732)
VIC	−36.109***	−3.084***	−30.775***	−2.465***
	(0.846)	(0.581)	(0.813)	(0.575)
WA	−186.296***	−60.263***	−179.259***	−62.342***
	(0.907)	(0.627)	(0.872)	(0.621)
Female			−1.349***	−1.450***
			(0.214)	(0.151)
Age			1.184***	0.344***
			(0.009)	(0.006)
Years with Bank			0.935***	0.211***
			(0.011)	(0.008)
Small Business			45.668***	32.394***
			(1.451)	(1.025)
Credit Card			27.448***	13.528***
			(0.147)	(0.104)
Savings Account			−0.622***	0.692***
			(0.096)	(0.068)
Transaction Account			17.874***	5.213***
			(0.091)	(0.065)
Concession Card			−52.158***	−16.063***
			(0.294)	(0.209)
Constant	291.646***	94.590***	170.589***	62.376***
	(0.854)	(0.599)	(0.922)	(0.656)
Observations	2,220,834	2,220,834	2,220,834	2,220,834
R^2	0.081	0.567	0.151	0.576
Adjusted R^2	0.081	0.567	0.151	0.576

This table shows results from a regression model of a continuous variable for average monthly energy bill amount post-intervention on a treatment status indicator (invited to the platform). Standard errors are robust and clustered at the state-level. Columns (2) to (4) show additional controls. Sample is restricted to customers with energy bills observable in bank data; outcomes are based on bill amounts for 6 months pre and 14 months post intervention. *p<0.1; **p<0.05; ***p<0.01

Effect of Visiting the Platform (Local Average Treatment Effect)

Customers are randomly invited (or not) to participate in treatment, but not everyone who gets invited visits the platform, and some in control who do not get invited nevertheless find their way to the platform on their own. We can estimate the effect of actually visiting the platform, the local average treatment effect (LATE), by dividing the intention-to-treat estimates (ITT) in Table 9 by the estimated difference in the share of 'compliers' in treatment versus control (0.206).[87] We find that actually visiting the platform lowers average monthly energy bills by $6.39, in our preferred specification, as seen in Table E2.[88]

Table E2: Effect of Visiting The Platform on Energy Bill Amounts

	Dependent variable:			
	Energy Bill Amount			
	(1)	(2)	(3)	(4)
Treatment (Platform Visit)	−30.723	−11.218	−14.0146	−6.388
Pre-Intervention Bill Control	Y	Y	Y	Y
State Controls	Y	Y	Y	Y
Other Covariate Controls	N	N	Y	Y

This table shows the local average treatment effect (LATE) of visiting the platform on energy bill amounts. The table is produced by scaling the intention-to-treat point estimates from Table 9 by 0.206, the estimated difference in the share of 'compliers' in treatment versus control.

[87] The difference in 'compliance' rates in visiting the platform is 20.6 percent (31.8 − 11.2). From the treatment group, a total of 31.8 percent (n=2,005,546) visit the platform: 18.9 percent (n=1,190,275) via the invitation link, and 12.9 percent (n=815,271) via the bank's website or app menus without clicking on the invitation link. From the control group, 11.2 percent (n=157,222) find their way to the platform.

[88] $1.316 is the intention-to-treat (ITT) effect; which scaled by the difference in compliance rates yields a local average treatment effect (LATE) of $6.39 ($1.316 \div 0.206 = 6.388$).

Acknowledgements

We gratefully acknowledge and thank the phenomenal team at the Commonwealth Bank of Australia for their partnership in designing and executing these field experiments start to finish: Alex Perry, Alister Hearnshaw, Andrea Nicastro, Andrew McMullan, Anika Patel, Ashley Tucker, Ben Grauer, Bernadette Galea, Bruno Rotunno, Carla Zuniga-Navarro, Carlos Vazquez, Chris Hill, Courtney Grigor, Dan Jermyn, Darren Broughton, David Rowe, Emily Daniels, Fred Tran, Grant Burrow, Gulley Shimeld, Ian Calbert, Jackson Davey, Jamie Koay, Jeroen Nieboer, Jessica Gleeson, Kenneth Uro, Madhurima Biswas, Madhu Somineni, Matt Comyn, Melissa Claire, Mohamed Khalil, Muhsin Karim, Nathan Damaj, Nikhil Ravichandar, Pete Steel, Prisca Han, Rafael Batista, Sree Gogineni, Stephen Bush, Sylvia Sara, Thomas Crawley, and Will Mailer. We thank Bekzod Abdullaev, Liam Ryan, Neena Narayanan, and Rizwan Bashir at and the New South Wales Department of Planning, Industry, and Environment for sharing additional data. Thanks to Alexandra Geddes, Filiz Ketenci, Shabana Ali and Vineet Gounder at Service New South Wales; Ciara Sterling and Gabby Sundstrom at the Thriving Communities Partnership; and Catherine Nehring and Ryan Davis. We also thank Harvard University's Sustainability, Transparency, and Accountability Research Lab for support. Any errors are our own.