American Economic Journal: Economic Policy 2025, 17(4): 1–29
https://doi.org/10.1257/pol.20220786

Take-up of Social Benefits: Experimental Evidence from France†

By Laura Castell, Marc Gurgand, Clément Imbert, and Todor Tochev*

We report on two nationwide experiments with job seekers in France. We first show that a meeting with social services to assess eligibility and help with applications to social benefits increased new benefit take-up by 29 percent. By contrast, an online simulator that gave personalized information on benefit eligibility did not increase take-up. Marginal treatment effects show that individuals who benefit the most from the meetings are the least likely to attend. Overall, without ruling out information frictions, our results suggest that transaction costs represent the main obstacle to applying for benefits or accessing government's assistance in applying. (*JEL* D83, I38, J22, J65, R38)

The efficacy of social policies throughout the world is limited by low benefit take-up (Ko and Moffitt 2022). An estimated 25 percent of people eligible for the earned income tax credit (EITC), the largest federal means-tested cash transfer in the United States, do not claim it (Plueger 2005). In France, the setting of our study, 36 percent of eligibles do not claim the out-of-work minimum income (Revenu de solidarité active), a figure that goes up to 68 percent for the in-work component (Domingo and Pucci 2012).

A key question is why poor households would forgo additional resources offered to them. Currie (2004) identifies three possible reasons. The first is the lack of information: Potential claimants may be unaware of their eligibility to social benefits. The second is transaction costs, which are due to complex and lengthy application procedures. The third is welfare stigma, due to the conflict between eligibles' self-image and their negative perception of benefit claimants (Moffitt 1983). These different explanations for non-take-up imply different policy responses: information or anti-stigma campaigns, or application assistance.

* Castell: INSEE (email: laura.castell@insee.fr); Gurgand: Paris School of Economics–CNRS and ENS-PSL (email: marc.gurgand@psemail.eu); Imbert: Sciences Po (email: clement.imbert@sciencepo.fr); Tochev: Institut des Politiques Publiques (email: todor.tochev@psemail.eu). Kory Kroft was coeditor for this article. We are grateful to Johannes Abeler, Manasi Deshpande, Michael Keane, Matthew Notowidigdo, and Stefano Caria, as well as seminar participants at Bocconi University, Institute for Employment Research (IAB), Paris School of Economics, Santa Barbara, Oxford, Sussex, and Warwick for useful comments. This project builds on the collaboration between the Caisse Nationale des Allocations Familiales, the French Ministry of Social Affairs and Health, and Pôle Emploi. We would like to thank these institutions for their support and, in particular, Kim Antunez, François Aventur, Magali Beffy, Anthony Dumser, Céline Grislain-Letremy, Etienne Perron-Bailly, and Augustin Vicard. We thank Marie Borel and Anaëlle Solnon for exceptional research assistance. The research was funded by the European Commission and the French Ministry of Social Affairs and Health. The Paris School of Economics IRB gave ethical approval (CE/2015-004). The project is registered in the AEA Registry as AEARCTR-0001882.

† Go to https://doi.org/10.1257/pol.20220786 to visit the article page for additional materials and author disclosure statement(s).

A natural follow-up question is whether alleviating barriers to benefit take-up will bring in poorer or richer benefit claimants—that is, improve or worsen targeting. For example, complex application procedures may screen out households that need the benefits the least (Nichols and Zeckhauser 1982; Kleven and Kopczuk 2011; Alatas et al. 2016). Or they may discourage less educated, financially stressed households that need the most support (Deshpande and Li 2019; Finkelstein and Notowidigdo 2019; Linos et al. 2020). The specific design of interventions that facilitate benefit take-up—for example, online or in person—may also change their impact on targeting.

In this paper, we use two nationwide experiments with French job seekers to investigate the barriers to take-up of 15 different social benefits, including family- and housing-related transfers, in- and out-of-work income supplements, and subsidized health insurance. In the first experiment, we show that job seekers who attended a meeting with social workers to inform them about their eligibility and help them apply to social benefits were 29 percent more likely to take up any new benefit. The effects were driven by family, housing, and income benefits, for which social workers could directly help them with their application, and, among these, by benefits that had more complex application forms. By contrast, in a second experiment, job seekers learned about their eligibility from an online simulator: This treatment had a much smaller and statistically insignificant effect on take-up. These results suggest that the lack of information about one's own eligibility to benefits, although likely present, is not the main barrier to the take-up of social benefits. Instead, the evidence points toward the importance of transaction costs. Finally, we estimate the marginal treatment effect (MTE) and find that the effectiveness of the meeting is higher for job seekers with a low propensity to attend it. This suggests that transaction costs also prevent job seekers from receiving assistance with applying to benefits.

In the first experiment, we evaluate the impact of a meeting with social workers (*Rendez-vous des droits élargis*, RDVDE). Within a target population of 60,000 job seekers with small or no claims to unemployment benefits, a randomly selected half was invited to a meeting with social services to discuss their potential eligibility for 15 different benefits and to support them in starting the application process. Using a combination of administrative data sources, we evaluate the effect of the intervention on benefit applications and take-up. To shed some light on barriers to take-up, we leverage additional random variation in the type of meeting (face-to-face or phone) and in the content of the letter (anti-stigma or information flyer).

The policy was well implemented: 21 percent of job seekers who were invited had a meeting with social services, thus reaching the official target of 6,000 RDVDE. It was also successful at inducing benefit take-up: The treatment group was 7 percent more likely to have received any new benefit 6 months after the intervention (the probability of taking up new benefits was 27 percent in the control). These effects were primarily driven by (in- and out-of-work) income support benefits, for which take-up increased by 10 percent, and by family and housing benefits, whose take-up rose by 5 percent (the probabilities of taking up these benefits were 15 and 11 percent in the control, respectively). When we estimate the effect of attending the meeting (instrumented by the invitation), we find it to be very large: +29 percent for any benefit, +44 percent for income benefits, and +21 percent for family and housing benefits after 6 months.

The RDVDE is a bundled intervention, which provides both information and application assistance, and may have even alleviated welfare stigma, although this was not its explicit aim. The evaluation findings, however, point to the importance of application costs. Indeed, the intervention was effective for benefits administered by the social services that offered the meeting and that could follow up immediately with the application: family, housing, and income support benefits. By contrast, we find no effect on health benefits, which required an application with other social services. This suggests that the RDVDE's main role is to lower application costs: If the main effect of the RDVDE was through better information or lower stigma, one would expect it to increase health benefit take-up as well.[1] Within the group of family, housing, and income benefits, we further show that the meetings had a stronger positive effect on take-up of benefits whose application forms were more complex. Using random variation in the invitation letters, we also show that the information and the stigma fliers had no additional positive effect on meeting attendance or benefit take-up. However, this last piece of evidence is not very strong, since the qualitative evaluation suggests that job seekers did not remember the flyers.

In order to study the role of pure information provision on benefit take-up, we ran another nationwide experiment one year prior to the RDVDE experiment on a sample of 40,000 job seekers targeted in the same way. We sent job seekers an email with an information leaflet similar to the one in the main experiment, with a direct link to the simulator but without an invitation to any meeting. This email generated interest among recipients, as 67 percent of the treatment group clicked on the link and 18 percent completed a simulation immediately. However, we find that the email only had small and insignificant effects on benefit take-up. To the extent that the information provided by the simulator was credible and clear, this suggests that the lack of information may not be the main barrier to benefit take-up. An alternative interpretation could be that the information provided by the simulator was not as effective as the one provided by social workers during the meetings.

Even if job seekers hold on average accurate beliefs, they could still be misinformed—some underestimating, some overestimating their eligibility. Within the RDVDE experiment, we combine survey data on applications made and administrative data on benefit receipt to show that the success rate of applications, which is about 90 percent in the control, fell slightly in the treatment group. This suggests that the intervention did not improve the quality of applications, as one would expect from better-informed applicants. Instead, our results are consistent with a lowering of application costs, which may have induced lower-quality applications.

In the last part of the paper, we ask two questions: whether the meeting induced poorer or richer claimants to take up social benefits and whether the meeting attracted the job seekers who needed assistance the most. To answer the first question, we

[1] The qualitative evaluation of the RDVDE experiment (Alberola, Muller, and Maes 2018) emphasizes the importance of meeting social workers, who help benefit claimants to fill out the form and to identify the required documents. Personal contact is also useful to focus the attention of benefit claimants and to answer their questions directly, which reduces their cognitive load.

use data on the amounts received by benefit claimants, which we obtained on a predefined 58 percent subsample of the experimental sample. We find that while the meeting increased coverage of income benefits in this subsample, it had a negative effect on the value of benefits received by the marginal claimant. Together with our finding that the success of application declined slightly, this result suggest that the meeting did not improve the targeting of social benefits—that is, it did not help poorer people deterred by transaction costs access social benefits.

All the experimental results we document are effects on compliers—i.e., job seekers who come to the meeting when they are invited. One may wonder whether the treatment effects would be lower or higher for never takers—i.e., people who do not come to the meeting. For this, we use the MTE framework developed by Heckman and Vytlacil (2005). The MTE is the effect of the meeting as a function of the (observed and unobserved) cost of going to the meeting. For identification, we combine experimental variation in the invitation to the meeting (face-to-face or online) and individual characteristics at baseline (e.g., distance to social services). We find that the MTE on income benefit take-up is increasing in the unobserved cost of attending the meeting, which means that people with a lower unobserved probability to come to the meeting have a higher benefit from attending. We also find that individuals with observed characteristics related to higher effects of the treatment have a lower predicted probability to attend the meeting. These results suggest that people who face high transaction costs in applying to benefits also face high costs in joining meetings that could help them apply. The differences between never takers and compliers are large: We estimate that attending the meeting increases the probability of taking up the income benefit by 19 points among never takers, which represents about €20 per month of additional benefits, against 7 percentage points (€6 per month) for compliers.

Our paper contributes to the literature that evaluates interventions aimed at increasing benefit take-up (see Ko and Moffitt 2022 for a more extensive review). Some authors have found positive effects of information provision: e.g., Duflo and Saez (2003) on retirement plan take-up and Kling et al. (2012) on Medicare plan choice by senior citizens in the United States. Bhargava and Manoli (2015) test several barriers on take-up of the EITC by varying the design and presentation of the letter sent to each treatment group. They find that take-up is most affected by the salience and simplicity with which information is presented, while efforts to attenuate stigma or reduce perceived transaction costs do not have any effect. However, Linos et al. (2020) find that similar nudges targeted at first-time tax filers have no effect on take-up, implying that "outreach to hard-to-reach populations will often need to include higher-touch interventions that simplify the underlying processes" (439).

Several (mostly nonexperimental) papers evaluate interventions that changed application costs. For instance, Kopczuk and Pop-Eleches (2007) find that the introduction of electronic filing programs increased take-up of the EITC because it increased the presence of tax preparers who are paid by the poor to help them apply. Similarly, Aizer (2007) finds that a toll-free number to assist claimants of Medicaid increases enrollment, and Herd et al. (2013) show that a plan to reduce administrative burden to claim Medicaid in Wisconsin increased take-up by about 10 percent. More recently, Ganong and Liebman (2018) show that several state-level policies that made it easier to apply to the Supplemental Nutrition Assistance Program (SNAP) did improve take-up, and Deshpande and Li (2019) find that social security office closures in the United States reduce applications to disability benefits.

Our paper belongs to the few papers that test the effect of information provision and application assistance in the same context using an experimental method. Bettinger et al. (2012) show that coupling improved information with personalized assistance in filling out applications leads to increased financial aid applications and receipt, while information provision alone has no effect on applications. Targeting a population of low-income elderly who were not currently enrolled in the SNAP income support program, Finkelstein and Notowidigdo (2019) find that an information letter doubled the probability of take-up, while application assistance tripled it. One obstacle to generalizing findings from the literature is that each study focuses on one particular type of benefit within an often narrowly defined target population. One contribution of our paper is to exploit two nationwide experiments that test an information and an application assistance intervention on take-up of 15 different social benefits (including family, health, and income benefits), implying a wide range of eligibility rules and potential beneficiaries. Our findings suggest that in our context, information provision alone is not enough to increase take-up and that application assistance is essential.

We also contribute to the ongoing discussion about the effects of application costs on targeting. On the one hand, Kleven and Kopczuk (2011) argue that the complexity of the application process to social benefits may be a necessary by-product of optimal targeting. Consistent with this view, the targeting of social benefits in Indonesia improves when social services that provide application support are located farther away (Alatas et al. 2016). On the other hand, Tempelman and Houkes-Hommes (2016) show that take-up of a Dutch health care allowance by eligible individuals decreases with income but is very low for the lowest incomes. Similarly, Deshpande and Li (2019) show that social security office closures in the United States reduce applications among people with better claims to disability benefits, and Gupta (2017) argues that application costs for old-age pensions in India screen out the most vulnerable women.

A key issue in empirical studies to date is that their conclusions regarding targeting are only true for people who comply with the treatment—that is, who access application support. But the same barriers that prevent poor people from applying to benefits may prevent them from accessing application support. Finkelstein and Notowidigdo (2019) find that the marginal applicant to SNAP induced by an information and assistance intervention has lower claims than average (suggesting good targeting), but they also argue that nonapplicants have lower education and may have higher claims (suggesting poor targeting). Our contribution is to formalize and test this idea empirically using MTEs. On the one hand, the intervention induces higher take-up but not higher benefits per recipient. On the other, job seekers who receive assistance are not the ones who would benefit the most from it.

The psychological literature interprets these types of findings from a cognitive perspective as the result of low executive functioning (Christensen et al. 2020). People who suffer from financial distress or poor health have a lower ability to plan and prioritize, lose the overview of the benefits of their actions, and are less resilient to obstacles. These types of mechanisms have also been highlighted by the behavioral economics literature (Mani et al. 2013; Mullainathan and Shafir 2013). This would explain why the most deprived people may also be those for whom administrative burden is felt more strongly, thus accessing benefits less often in spite of higher needs.[2]

Finally, our paper contributes to the literature devoted to the use of digital solutions to improve the functioning of social programs. Recent studies show that the use of IT in beneficiary identification and payments can reduce corruption and improve targeting but may increase exclusions errors or payment delays (Muralidharan, Niehaus, and Sukhtankar 2016, 2025; Banerjee et al. 2020). In a similar vein, we show that, in contrast to phone or in-person meetings with social workers, an online simulator that provides personalized information on benefit eligibility had no effect on benefit take-up. Our results suggest that digital solutions may not improve access to social programs if they are unable to significantly reduce transaction costs for vulnerable populations (job seekers in our case).

The remainder of the paper is structured as follows. Section I presents the context and the interventions. Section II describes the data and the empirical strategy for the evaluation. The main evaluation results are presented in Section III. Section IV discusses the reasons for low take-up, and Section V explores the implications for targeting. Section VI concludes.

I. Context and Intervention

A. *Institutional Setting*

Social security in France is organized in agencies (*Caisses*), each responsible for a specific type of benefit. Income, family, and housing benefits are managed by the Caisse nationale d'allocations familiales (CNAF), and health benefits are managed by the Caisse nationale d'assurance maladie (CNAM). Unemployment benefits for all job seekers are managed by Pôle Emploi, a separate state-controlled agency. These multiple agencies were created at different points in history and kept their autonomy (Direction de la Sécurité sociale 2017). Procedures, amounts, and eligibility rules differ substantially across agencies, which makes benefit applications more complex.

In this paper, we will consider in total 15 benefits, which for simplicity we pool into three groups. The first group, which we call *family benefits*, includes benefits that are given to parents or people with disabilities, but also housing benefits, which are managed by the CNAF. The second group, which we call *health benefits*, includes benefits that subsidize access to complete health care insurance coverage and are managed by the CNAM. The third group, which we call *income benefits*, covers benefits that

[2] Dubois (2003), in his sociological analysis of social workers in France, offers a different interpretation: He argues that social workers provide less assistance to benefit claimants who are more socially distant from them (and, hence, poorer). However, this interpretation does not seem in line with the observation made in the qualitative evaluation of the RDVDE experiment (Alberola, Muller, and Maes 2018), which reports that social workers were eager to help the most deprived individuals access benefits and were disappointed by the fact that the people most in need did not attend the meetings.

directly support low-income individuals, whether in or out of work, and are managed by the CNAF.[3]

B. *Intervention*

The issue of low benefit take-up has been discussed in France since at least the late 1990s (Warin 2012). It received more public attention in the early 2010s, when studies suggested that less than 50 percent of individuals eligible for the *Revenu de solidarité active* (the main income benefit managed by the CNAF) actually claimed it (Domingo and Pucci 2012). In 2013, increasing take-up became part of the mission of the CNAF, which committed to implement the "*Rendez-vous des droits*" (RDVDD) (Decobecq 2013). Launched in 2014, the RDVDD is a one-hour meeting with a social worker offered to potential claimants to assess their eligibility for social benefits. The aim is to cover *all* nationally available social benefits, not only those administered by the CNAF. It focuses on individuals who are known to the CNAF, either because they are themselves benefit recipients or because they are part of a recipient's household. It targets among them those who were flagged because they have likely experienced a recent negative income shock. CNAF social workers are trained for the RDVDD to deliver comprehensive information on the 15 social benefits, including non-CNAF benefits.

This paper evaluates a new version of the RDVDD, known officially as the *Rendez-vous des droits élargi* (RDVDE), that was experimented with during the last trimester of 2017. As compared to the RDVDD, the RDVDE had a broader targeting scope. The trigger-based strategy was carried forward but was no longer limited to CNAF beneficiaries. Instead, the target population was low-income individuals who had registered as unemployed at Pôle Emploi and many of whom were not present in the CNAF's databases. Two types of individuals were considered. *Population 1* had registered as job seekers between one and three months before the randomization date and didn't have a recent previous unemployment spell.[4] *Population 2* included job seekers whose rights to unemployment benefits (*Allocation de retour à l'emploi*) were due to expire less than three months before the randomization date. We removed from the sample the most affluent individuals by excluding those who received an unemployment benefit amount above a certain threshold (see Supplemental Appendix E). Because the two populations just experienced an income shock, they may have become eligible to benefits and could benefit from the RDVDD.[5]

The structure of the RDVDE was the following. The meeting lasted about an hour. At the start of the RDVDE, the agent went through a series of basic questions on the household composition, occupational status, and income of the potential claimant. A

[3] See Supplemental Appendix Table B1 for more detail about each benefit.

[4] Specifically, a previous spell, if any, had to be closed for more than three months before randomization. This condition removes individuals from industries (e.g., the arts) characterized by frequent short unemployment spells who are likely to be well aware of their benefit eligibility.

[5] The targeting strategy was chosen based on exploratory small-scale trials conducted by the SGMAP (Secrétariat général pour la modernisation de l'action publique) and members of the research team, in partnership with two CNAF agencies and Pôle Emploi in 2013.

decision tree algorithm helped social workers determine which social benefits should be discussed, but they were free to overrule the algorithm's recommendations. The qualitative evaluation of the RDVDE experiment (Alberola, Muller, and Maes 2018) highlights that the algorithm was a central and distinguishing feature of the intervention, which social workers found especially useful to assess eligibility to benefits that were not provided by CNAF, although it was not used systematically in practice. During the meeting, the worker could consult information on benefit criteria stored on the CNAF intranet, as well as on the internet, and perform simulations using online benefit calculators. Once a potential for new benefits had been identified, the worker instructed the person on the process to claim those benefits, which was often cumbersome, and usually involved a different requirement and paperwork for every single benefit. When the benefit was provided by CNAF (i.e., all family and income benefits), the worker could start the application process on the spot.[6] This was not possible, however, for health benefits, which are provided by another agency. Finally, social workers were prompted by the intranet tool to fill out a questionnaire about the meeting and to report which benefit they had discussed with the potential claimant.

C. *Experimental Design*

This study uses an encouragement design. Half of the job seekers from the two populations received letters by mail inviting them to the RDVDE. The allocation of the meetings was managed by the CNAF. There were two ways for potential claimants to arrange a meeting. The first was to call the landline number provided on the invitation letter, in which case they were transferred to the CNAF's call center department, which then allocated meeting slots matching the local CNAF agency's availabilities with the individual's time and location preferences. Less than 10 percent of the meetings in our sample were allocated in this way. Most meetings were agreed to when the call center directly called potential claimants, which it did systematically in the treatment group: Social workers were instructed to call each individual up to three times.

The content of the invitation letter was randomized: There were six treatment groups following a 3×2 crosscutting design. Supplemental Appendix D presents the invitation letters. In the first dimension, there were three types of letters.[7] The *neutral* letter (see Supplemental Appendix Figure D5) explained the reason for the invitation and the purpose of the meeting, invited recipients to call to fix a meeting, and informed them that they may also be called by a CNAF employee. The *information* letter (see Supplemental Appendix Figure D6) consisted of the same invitation, with a second page listing several common life events that led to negative income shocks and the corresponding social benefits. It also included a flyer (see Supplemental Appendix Figure D7) that described four fictitious households, with the types and amount of benefits they were eligible for, and invited readers to visit the online simulator mes-aides.gouv.fr to obtain a personalized estimate of their

[6] The qualitative evidence gathered from interviews with both potential applicants and social workers suggests that both groups consider this to have been one of the most useful aspects of the meeting.

[7] All three types were printed on a joint CNAF/Pôle Emploi letterhead and posted using the same envelopes as used by the CNAF in its day-to-day contacts with the public.

eligibility. The *anti-stigma* letter contained the same invitation as the neutral letter but included a flyer that told three real-life stories of people who experienced negative income shocks: birth of a child and job loss, bereavement, and family separation (Supplemental Appendix Figure D8). Each story presented social benefits as a safety net and benefit receipt as a legal right.

The second dimension of the design aimed to create variation in the costs incurred by individuals in attending the RDVDE. For a randomly selected half of the treatment group, each of the earlier three types of invitation letters offered individuals a choice between participating in a face-to-face meeting with a social worker or having the meeting over the phone. In the other half, a face-to-face meeting was the only option offered in the letter. The phone group thus faced potentially lower meeting costs.

D. *Sampling*

Randomization and invitations happened in two separate waves. In wave 1, a sample of job seekers registered as unemployed on May 31, 2017, received invitations from September 18, 2017, onward. In wave 2, a sample of job seekers registered as unemployed on July 31, 2017, received invitations from November 2, 2017, onward. In total, we drew 90,000 individuals from administrative unemployment registers, 45,000 for each target population. They were informed of the fact that their data would be used for research purposes: About 2,000 individuals refused and were withdrawn from the sample. Due to logistical constraints from the CNAF, we had to select a smaller subsample of 60,000 individuals to participate in the study (24,000 for wave 1 and 36,000 for wave 2) with a probability proportional to the number of job seekers registered in each *département* within each population (1 and 2).

We allocated job seekers to six treatment and one control groups using a stratified random sampling procedure carried out by computer. The strata correspond to cells formed by all the possible combinations of three variables: CNAF branch to which the individual is affiliated, type of target population (1 or 2), and a variable indicating whether an individual's distance to the nearest CNAF center is smaller than the median in their département (calculated over all the individuals in the base sample). Half of the sample (30,000) was allocated to the control group, who received no invitation at all, with the other half distributed equally across all the six treatment groups (5,000 per subgroup).

The research team set the minimum number of meetings needed to detect the effect of the intervention at 6,000 (20 percent of the treatment group of 30,000). In practice, the CNAF network used 6,000 as a target and stopped arranging meetings once the quantitative objective set to each département was achieved. We do not dispose of the list of job seekers who were called; hence, we cannot use this information in our analysis. However, in a few départements, no meetings were planned for second-wave individuals, because all available slots had been used for the first wave: We drop these wave $\times$ département samples altogether. These are balanced by construction since the randomization was stratified by wave $\times$ département. In the end, the baseline sample is composed of 54,418 individuals. The baseline characteristics of the sample are presented in Table 1.

Table 1—Baseline Characteristics and Balance Tests

Variable	Control (1)	Phone (2)	In person (3)	*F*-test *p*-value (4)
Panel A. Sociodemographic variables				
Married	0.3263	0.0012	0.0011	0.9602
Number of children (Pole Emploi)	0.7345	0.0094	−0.0156	0.2205
Age in years	35.2695	−0.0345	−0.0460	0.9222
Female	0.4715	0.0109	0.0062	0.1129
Education: High school or less	0.1794	−0.0045	−0.0032	0.4888
Education: Missing	0.0055	−0.0007	0.0007	0.2965
Education: CAP/BEP	0.3501	0.0050	0.0013	0.6153
Education: Bac général	0.2487	−0.0044	0.0006	0.5878
Education: Higher education	0.2218	0.0039	0.0013	0.6899
French national	0.8505	−0.0004	0.0006	0.9744
EU national	0.0357	0.0036	0.0019	0.1950
Non-EU European national	0.0061	−0.0001	−0.0007	0.6732
Rest of the world	0.1070	−0.0032	−0.0023	0.5587
Years of experience in target job	5.6770	0.0359	−0.0100	0.8644
Unemployment duration in months	11.9919	−0.1709	−0.3205	0.1879
Monthly unemployment benefits (€)	435.5438	6.8806	5.0500	0.3505
CNAF beneficiary	0.5793	−0.0012	−0.0022	0.9122
log daily reference salary (€)	3.8480	0.0067	0.0029	0.5691
Panel B. Baseline benefit usage				
Any benefit at baseline	0.4763	0.0057	−0.0029	0.3737
Any family benefit at baseline	0.3639	0.0054	−0.0004	0.5380
Any health benefit at baseline	0.1636	0.0025	−0.0037	0.4014
Any income benefit at baseline	0.1668	0.0061	0.0015	0.3291
AF at baseline	0.1553	0.0019	−0.0036	0.4629
ASF at baseline	0.0361	−0.0002	−0.0003	0.9861
Any AL at baseline	0.3172	0.0045	−0.0008	0.6140
AAH at baseline	0.0126	−0.0016	0.0001	0.3478
AEEH at baseline	0.0084	0.0002	0.0009	0.6632
PAJE: Prime à la naissance	0.0157	−0.0006	−0.0009	0.7528
PAJE: Allocation de base	0.0655	0.0015	−0.0019	0.5510
PAJE: Prestation partagée d'éducation de l'enfant	0.0042	−0.0002	0.0001	0.9406
PAJE: Complément de libre choix du mode de garde	0.0016	−0.0001	−0.0004	0.5381
ACS at baseline	0.0017	−0.0001	0.0001	0.9639
CMUC at baseline	0.1626	0.0025	−0.0039	0.3698
ASI at baseline	0.0003	0.0001	0.0002	0.7369
RSA at baseline	0.0718	0.0040	0.0007	0.3747
PA at baseline	0.1080	0.0033	0.0004	0.6102
Observations	29,563	11,911	12,944	

Notes: The first column presents the corresponding control group mean for the row variable. The second and third columns show the difference between the corresponding treatment group mean and the control group mean. The final column shows the *p*-value associated with a joint test of equality between the treatment group that received an invitation for a phone meeting, the treatment group that received an invitation for an in-person meeting, and the control group. See Supplemental Appendix Table B1 for more detail about each benefit.

E. *Additional Experiment: Online Simulator*

We carried out a second experiment, which evaluates the online social benefits simulator mes-aides.gouv.fr. This online simulator was created by the Secrétariat Général pour la Modernisation de l'Action Publique (SGMAP), a government agency responsible for modernizing government services. It covers the same range of benefits as the RDVDE and offers personalized information on benefit eligibility, including the amounts for which potential beneficiaries might be eligible.

The sampling for the additional experiment followed the same structure as in the main experiment: We selected job seekers from the two target populations (Populations 1 and 2) in June 2016 using Pôle Emploi data and randomly allocated each individual to one of two experimental groups. The treatment group received an email containing generic information on the social benefits covered by the simulator, as well as a personalized link inviting them to visit the mes-aides.gouv.fr website. The control group did not receive any email. To reduce the risk of information spillovers, we first randomized Pôle Emploi agencies and allocated each agency to the control or the treatment group. In every agency, we then randomly drew 50 individuals from each of the two target populations. Job seekers registered to the same agency received the same treatment. The total sample size is 40,000 individuals, with 20,000 individuals in each treatment group (10,000 for each target population).

II. Data and Empirical Strategy

A. *Data Sources*

The data used in this paper are drawn from four administrative sources and one household survey. The first data source was jointly provided by Pôle Emploi and CNAF and compiled by the Direction de la Recherche, des Études, de l'Évaluation et des Statistiques (Drees) of the Ministère des Affaires Sociales (Drees 2016, 2018a). Pôle Emploi collects information on job seekers' characteristics, including age, sex, marital status, number of children, nationality, educational achievement, type of work sought, reason for registering as a job seeker, length of the unemployment spell, and unemployment benefit amount and duration. The sociodemographic information is collected at the time of registration and updated for every unemployment spell. The unemployment information is updated every month by job seekers; it is a condition for receiving unemployment benefits. Pôle Emploi transmitted this dataset to the CNAF, which added an indicator for individuals who had previously received a CNAF benefit either themselves or in their household.

The second dataset is extracted from the Répertoire National Commun de la Protection Sociale (RNCPS), a social security database that contains information on the benefit receipt status of every individual holding a French social security number. This database is hosted by the CNAV, which maintains an interface accessible to other French social security agencies (Friconneau 2014). Data for each wave were queried via this interface by the CNAF at baseline and three months and six months after the end of the intervention: February and May for the first wave and April and July for the second wave (CNAF 2018). The RNCPS provides cross-sectional individual-level social benefit receipt data but does not contain information on benefit amounts. Information on benefits that are no longer received is only kept if the individual does not receive the same benefit again—and only up to one year. If a new claim is approved, all information about the old claim is overwritten.

We use RNCPS data to construct three different measures of benefit take-up. The first is a *stock* measure, defined as a dummy variable equal to 1 if an individual is registered as receiving benefits at month t and 0 otherwise. The second, a *flow* measure, is a dummy variable equal to 1 if the individual is registered at month t

as receiving a benefit that they did not receive at baseline and 0 otherwise. Our *preferred* measure is a dummy variable equal to 1 if the individual is registered as having received a benefit that they did not have at baseline, at any point between baseline and month t, even if they do not receive it anymore at time t.[8] While we focus on the preferred measure in most of our analysis, we also provide results using the other two measures for completeness. As per the pre-analysis plan, we construct these measures for each benefit group (family, health, and income; see Section IIA), and our baseline outcome is having received at least one benefit in the group.[9]

We supplement the RNCPS data using a third source, the CNAF's own internal database, which provides information on family and income benefit amounts (CNAF 2021). These data are only available for individuals in the RDVDE sample who were known to the CNAF at baseline (58 percent of the full RDVDE sample), and they do not cover health benefits.

The fourth data source, provided by the CNAF and compiled by the Drees, details information about the RDVDE meeting (Drees 2018b). Social workers reported in an intranet tool developed for the experiment whether they held meeting over the phone or face-to-face and which social benefits they discussed.[10] In addition to these data, a team of sociologists from the Centre de Recherche pour l'Étude et l'Observation des Conditions de Vie (CREDOC) conducted qualitative interviews with social workers and meeting beneficiaries in four départements.

Finally, we conducted a phone survey with 10,000 individuals selected from the experimental sample (Castell et al. 2024).[11] The response rate was about 61 percent (16,489 individuals were called).[12] The survey collected information about benefit receipt by the household at the time of the survey and asked respondents about benefit applications they or someone from their family had made since September (first-wave sample) or November (second-wave sample) that had not (yet) led to benefit receipt because they were either unsuccessful, pending, incomplete, or abandoned. We use these data to compute success rates of the applications.

For the additional experiment, which evaluates the online simulator, we measure compliance outcomes, such as the email open rates and the simulator usage rate through electronic trackers (Drees 2016). Because the randomization unit is geographical in that design, we also collect connections by postal codes to assess differential compliance. The main outcome is benefits take-up, which we measure in the RNCPS at 8 and 12 months after the start of the intervention but not at baseline, 3, and 6 months as in the main experiment (CNAF 2017).

[8] Hence, by construction, both the preferred and the flow measures are normalized to 0 at baseline.

[9] We also construct three measures of take-up (stock, flow, and preferred) based on the *number* of benefits received by group, and report the treatment effect on those outcomes too.

[10] The tool had a checklist for popular benefits and free text entry for the others. Many social workers reported all benefits discussed as a free text entry. We ran a regular expression search of the free text portions to search for benefit names, abbreviations, and possible variant spellings, and completed the checklist data.

[11] The survey sample was drawn by stratified random sampling using the same strata as the randomization process. The sampling probabilities for individuals called were $1/9$ for each treatment group and $1/3$ for the control group. Therefore, the treatment group was oversampled compared to the control group.

[12] Supplemental Appendix Table B3 compares the characteristics of survey respondents with those of the whole sample. Survey respondents are significantly more likely to be female and are younger, less experienced, more educated, and less likely to be married. These differences are, however, very small (less than 1 percent of the mean).

B. *Randomization Checks*

We perform randomization checks by estimating the following equation:

$$X_i = \alpha + \sum_j \delta_j T_i^j + \pi_s + \varepsilon_i, \tag{1}$$

where X_i is a characteristic of individual i at baseline, including take-up of benefits covered by the RDVDE, T_i^j are dummies for the different treatment groups, and π_s are randomization strata fixed effects. A stratum is defined by the interaction of four dimensions: implementation wave, target population, distance to the nearest CNAF branch, and département. As randomization occurred at the individual level, standard errors are only adjusted for heteroskedasticity. Table 1 presents balance tests for sociodemographic variables (panel A) and baseline benefit take-up (panel B): There is good balance between the control and the two types of invitation (phone and in-person).

Supplemental Appendix Table B2 presents the same balance tests with the full set of six treatments (interacting the type of invitation with the flyer received with the invitation).[13]

C. *Empirical Strategy*

In most of the paper, we pool all treatment groups into being invited to a meeting or not. Let T_i be a dummy variable equal to 1 if a job seeker received any type of invitation to the RDVDE ($T_i = \sum_j T_i^j$). Let Y_i denote the outcome of interest. We estimate by OLS the intention-to-treat (ITT) effect of being invited to the RDVDE:

$$Y_i = \beta_{OLS} T_i + \mathbf{X}_i' \boldsymbol{\delta} + \pi_s + \varepsilon_i, \tag{2}$$

where π_s are strata fixed effects. Given the large number of variables available, we use the robust double-lasso procedure (Belloni, Chernozhukov, and Hansen 2014b) to select controls, $\mathbf{X}_i$. Standard errors are robust to heteroskedasticity.[14]

In order to estimate the effect of attending the meeting itself, we also implement an IV strategy in which we regress outcomes on meeting attendance and instrument attendance by the invitation. Formally, we use two-stage least squares (2SLS) with the following two equations to estimate the local average treatment effect (LATE):

$$M_i = \gamma T_i + \mathbf{X}_i' \boldsymbol{\alpha} + \pi_s + u_i, \tag{3}$$

$$Y_i = \beta_{2SLS} \hat{M}_i + \mathbf{X}_i' \boldsymbol{\delta} + \pi_s + \varepsilon_i, \tag{4}$$

[13] Supplemental Appendix Table B4 reports balance tests for survey respondents: There are few significant differences across treatment groups. Supplemental Appendix Table B19 presents balance tests for the evaluation of the online simulator, which is also balanced.

[14] When we estimate regressions on the online simulator experiment, standard errors are clustered at the employment agency level.

where M_i denotes attendance to a phone or in-person meeting, as measured in the RDVDE extranet module, and $\mathbf{X}_i$ are the same variables, as in the OLS specification. The IV approach is only valid if the invitation is exogenous, has a significant effect on meeting attendance on average, affects meeting attendance in the same direction for all job seekers (no defiers), and does not affect benefit take-up in other ways than through meeting attendance.

In the last part of the paper, we go beyond the LATE and estimate the MTE to study whether job seekers who attended the meetings were the ones who stood the most to gain from it (Heckman and Vytlacil 2005; Heckman 2010). We follow the literature and estimate the MTE in two stages (Brinch, Mogstad, and Wiswall 2017; Cornelissen et al. 2018; Bhuller et al. 2020). The first stage uses a probit to estimate the probability of attending the meeting conditional on receiving an in-person invitation, a phone invitation, observable characteristics, and the interaction between the two treatments and observables. We use the latent variable index model:

$$(5) \qquad M_i = \mathbf{1}_{\{M_i^* > 0\}}$$

with

$$M_i^* = \gamma_1 T_i^1 + \gamma_2 T_i^2 + T_i^1 \times \mathbf{X}_i' \boldsymbol{\alpha}_1 + T_i^2 \times \mathbf{X}_i' \boldsymbol{\alpha}_2 + \mathbf{X}_i' \boldsymbol{\alpha}_3 + \pi_s + u_i,$$

where T_i^1 and T_i^2 are two indicators for the standard invitation and the phone invitation, which are randomized. $\mathbf{X}_i$ are predictors of meeting attendance chosen by robust lasso (Belloni, Chernozhukov, and Hansen 2014a). The second stage estimates a flexible relationship between the predicted propensity to attend and outcomes:

$$(6) \qquad Y_i = a\widehat{p(\mathbf{X}_i, \mathbf{Z}_i)} + b\widehat{p(\mathbf{X}_i, \mathbf{Z}_i)}^2 + \mathbf{X}_i'\mathbf{c} \times \widehat{p(\mathbf{X}_i, \mathbf{Z}_i)} + \mathbf{X}_i' \boldsymbol{\delta}_4 + \pi_s + \varepsilon_i,$$

where $\mathbf{Z}_i$ is a vector including the two treatment indicators T_i^1 and T_i^2, and $\widehat{p(\mathbf{X}_i, \mathbf{Z}_i)}$ is the propensity score computed in the first stage. We apply the local instrumental variable (LIV) estimator to identify how treatment effects vary with the predicted propensity. Standard errors are based on 500 bootstrap replications of both stages. The MTE is the derivative of this equation with respect to the propensity score: It is equal to $a + \mathbf{X}_i'\mathbf{c} + 2b\widehat{p(\mathbf{X}_i, \mathbf{Z}_i)}$. It has an observed and an unobserved component, which are assumed to be separable (as in Brinch, Mogstad, and Wiswall 2017; Cornelissen et al. 2018; Bhuller et al. 2020). If the vector $\mathbf{c}$ is not zero, the treatment effect varies with observables $\mathbf{X}$, which captures the traditional heterogeneity in treatment effects. The term $2b\widehat{p(\mathbf{X}_i, \mathbf{Z}_i)}$ measures treatment effect heterogeneity along the unobserved dimensions that drive the decision to receive the treatment. Specifically, a positive coefficient b indicates that people with a higher unobserved cost to attend the meeting have higher treatment effects.

In our specification, the MTE is linear in $p(\mathbf{X}_i, \mathbf{Z}_i)$. Our instrument takes three values (control, standard, and phone invitation) and is interacted with predictors $\mathbf{X}$ ($T_i^1 \times \mathbf{X}_i$ and $T_i^2 \times \mathbf{X}_i$ in equation (5)), which generates continuous variation. The separability assumption implies that the MTE is identified over the unconditional support of $p(\mathbf{X}_i, \mathbf{Z}_i)$, as opposed to the support of $p(\mathbf{Z}_i)$ conditional on $\mathbf{X}_i = \mathbf{x}_i$. In principle, under this separability assumption, we could identify more flexible forms

(Brinch, Mogstad, and Wiswall 2017). In practice, however, the amount of variation generated by the data is limited, and we find that cubic and quartic terms are never significant in the second stage of the MTE (see Supplemental Appendix Table B18).

III. Results

This section presents the main findings from the RDVDE experiment: the effects of invitations on meeting attendance, the effects of invitations on benefit take-up, and the effects of meeting attendance on benefit take-up.

A. *Compliance*

Supplemental Appendix Table B5 presents the treatment effects on the probability of attending the RDVDE. As column 1 shows, in the control, the probability of attending a RDVDE is null. The average compliance rate for the treatment group is 21.3 percent, which corresponds to the official target of 6,000 meetings. Within the treatment group, participants who were given the choice between meeting in person and on the phone were only slightly more likely to attend than those who were only invited to a face-to-face meeting (column 2). We also test in column 3 whether the content of the invitation letter had any effect on the probability of attending the meeting. We find that beneficiaries who received an information and an anti-stigma flyer were significantly *less* likely to attend the meeting by 1.3 and 1.1 percentage points: The lack of information on eligibility or welfare stigma do not seem to act as barriers to meeting attendance.[15]

Columns 4 to 6 in Supplemental Appendix Table B5 present the estimated treatment effect on the probability of having a meeting on the phone. Nobody in the control group had a phone meeting, and 8.3 percent of the treatment group had a phone meeting—that is, a bit more than a third of all RDVDE were held on the phone (column 4). As column 5 shows, the probability of having a phone meeting is 10.1 percent in the treatment group that was offered a phone meeting in the invitation mail and 6.7 percent in the group that was only offered a face-to-face meeting initially. The deviation from the study protocol is due to the success of phone meetings, which were offered increasingly to participants of all treatment groups, even when they were not supposed to. Once again, the anti-stigma flyer seems to have slightly reduced meeting attendance.

In Supplemental Appendix Table B7, we provide additional information on the content of the RDVDE meeting based on administrative data from the intranet tool used by social workers. Specifically, we estimate that 19 percent of job seekers who received an invitation to the meeting attended and discussed at least one social benefit they may be eligible for. When we break it down by benefit group, we find that 14 percent discussed at least one family benefit, 11 percent a health benefit, and 16 percent an income benefit. These results suggest that over 90 percent of job seekers who attended the RDVDE meeting were potentially eligible

[15] As Supplemental Appendix Table B6 shows, these findings are virtually identical if we correct for contamination bias using the interacted regression solution proposed by Goldsmith-Pinkham, Hull, and Kolesár (2024).

Table 2—Effect of Invitation and Attendance on Benefit Take-up

	Any benefit	Family benefit	Health benefit	Income benefit
Panel A. Effect of receiving an invitation to the RDVDE (*ITT*)				
Any invitation	0.0174	0.0054	0.0017	0.0150
	(0.0037)	(0.0027)	(0.0023)	(0.0031)
Control group mean	0.2669	0.1099	0.0761	0.1544
Observations	54,418	54,418	54,418	54,418
Panel B. Effect of attending the RDVDE (*LATE*)				
Attended meeting	0.0815	0.0255	0.0079	0.0704
	(0.0172)	(0.0126)	(0.0106)	(0.0145)
Complier control mean	0.2859	0.1195	0.0897	0.1618
Observations	54,418	54,418	54,418	54,418

Note: Panel A presents the ITT estimate of being invited to the RDVDE on benefit take-up. Panel B presents the 2SLS estimate of the LATE of the RDVDE meeting. In both panels, the outcome variable is the preferred measure of benefit take-up at six months after the start of the intervention. The preferred measure is a dummy variable equal to 1 if the individual was registered as receiving any benefit from the corresponding benefit group at any point since the start of the intervention and that benefit was not received at baseline, and 0 otherwise. In panel B, we instrument meeting attendance with being assigned to any group receiving an invitation letter. Controls are selected using the robust double-lasso method, and coefficients are omitted for brevity. All specifications contain strata fixed effects. Standard errors are robust to heteroskedasticity.

for social benefits, including family, health, and income benefits. We will next test whether this potential eligibility translated into actual take-up.

B. *Effect of Invitations on Benefit Take-up*

We now use equation (2) to estimate the effect of the RDVDE invitation on our main outcome, the take-up of social benefits. Panel A in Table 2 presents the treatment effects on new benefit receipt during a period of six months since the intervention. As column 1 shows, 26.7 percent of job seekers in the control group received a new benefit in these 6 months. As compared to the control, the treatment group was 1.74 percentage points more likely to have received any new benefit. The effects are driven by family and housing benefits (0.54 percentage point increase; see column 2) and income support benefits (1.5 percentage point increase; see column 4), with no treatment effect for health benefits (column 3). Supplemental Appendix Table B9 presents the results for each benefit. The rise in family benefits receipt is entirely due to increased housing benefit (AL) take-up. The positive effect on income benefits is present both for in- (PA) and out-of-work (RSA) benefits but is stronger for the former.

We provide more evidence on the effects of the RDVDE in the Supplemental Appendix. First, Supplemental Appendix Table B8 shows that the effects increased between three and six months after the intervention. Second, Supplemental Appendix Table B10 presents treatment effects on alternative outcome measures. In panel A, we use the *stock* measure—that is, all benefits received at endline (three or six months to the meeting), including those being already received at baseline. In panel B, we use the *flow* measure—that is, only benefits still open at endline, not including those that may have been closed between the intervention and the endline. The results in panel B of that table are qualitatively similar to our preferred measure, but the treatment effects are smaller. Taken together, these results suggests that part of the benefit take-up generated by the RDVDE may have been temporary—for example, to help transitions from unemployment to a paying job.[16] When we use the *stock* measure (panel A), we find even smaller effects, which is expected since the meeting would have a low effect on benefits at endline for households who already had one at baseline. Finally, we estimate the treatment effects on the number of benefits instead of dummies for any benefit receipt: Supplemental Appendix Table B11 shows that the intervention increased the number of benefits received, with stronger effects for our preferred measure than for the *stock* or *flow* measures.

We next turn to the effect of the different types of invitations to the RDVDE. Supplemental Appendix Table B12 presents the results. In panel A, we test whether being offered a phone meeting made a difference to the treatment effect on benefit take-up. The treatment group who was invited to a phone meeting had a slightly higher chance of opening new benefits, but the difference is not significant. In panel B, we test whether the content of the invitation letters had any differential effect on benefits. We find that, as compared to the neutral invitation, the group who was sent the information flyer was more likely to receive new benefits after six months, but the difference is not significant. There is no evidence that the anti-stigma letter had any differential effect on take-up either. Overall, it does not seem that the type of invitation changed the effectiveness of the meeting in inducing benefit take-up.[17]

C. *Effect of Meeting Attendance*

The results presented so far have documented the effect of receiving an invitation to the RDVDE. We next estimate the effect of attending the meeting itself and estimate the IV specification of equation (3). The IV estimates are presented in panel B of Table 2. As expected given the relatively low compliance rate, the LATE is much larger than the ITT. Attending the meeting increases the probability of receiving any new benefit by 8.15 percentage points, a 29 percent increase as compared to the complier control mean (column 1). As column 4 shows, the meeting was particularly successful at inducing income benefit take-up: The meeting increased the probability of receiving a new income benefit over six months by 44 percent (7.04 percentage point increase). The meeting also increased the take-up of any new family benefit over six months by 21 percent (2.55 percentage point increase; see column 2). By contrast, there is no significant effect on health benefit take-up, and the point estimate is 9 percent of the complier control mean. The very large effects of the meetings suggest that the take-up of social benefits is far from 100 percent in the population of job seekers we study: Among compliers, *at least* 8 percent do

[16]The qualitative evaluation of the experiment (Alberola, Muller, and Maes 2018) notes that the RDVDE attracted workers that alternate frequently between employment and unemployment. This population, which is often unknown to social services, is not very poor but is entitled to in-work (PA) benefits, often temporarily, thus with limited incentive to claim it.

[17]Naturally, the attendance rates to the meeting are different for each of those incentives, but scaling the effects to that rate (i.e., estimating LATEs) does not change this general finding.

not receive benefits they are entitled to. We will show in Section VI that job seekers who are the least likely to attend the RDVDE are also the ones for which the benefit of the treatment is highest; this suggests that benefit take-up could be even lower in the noncomplier population.

IV. Barriers to Take-up

A. *Evidence from the RDVDE Experiment*

The RDVDE is a bundled intervention: It provides information on benefit eligibility, as well as assistance for applications for family and income support benefits. It could also alleviate welfare stigma, although this was not the explicit aim of the meeting between social workers and beneficiaries. A natural question is which of these components were the most important in generating the increase in benefit take-up following the RDVDE.

In the previous section, we showed that, in contrast to family and income benefits, the RDVDE had very little effect on the take-up of health benefits (Supplemental Appendix Table B8). This is despite the fact that 11 percent of the treatment group (and half of the job seekers who went to meetings) discussed their eligibility for health benefits during the meeting (Supplemental Appendix Table B7). The main difference between health benefits and the other types of benefits is that they are provided by a different agency and the CNAF social workers in charge of the RDVDE could not offer application assistance for them. This result suggests that the reduction of transaction costs is an essential part of the RDVDE—that is, that providing information or alleviating welfare stigma alone are not enough to induce take-up in this context.

To provide more evidence on the role of the RDVDE in reducing transaction costs, we can exploit the varying degrees of complexity of the application to the different benefits covered in the meeting. Specifically, we focus on family and income benefits, which are delivered by the CNAF, and count the number of fields that need to be filled in the application forms of each benefit. We then plot the treatment effect of the invitation to the meeting on the take-up of each benefit against this measure of complexity of the application process. As Figure 1 shows, the RDVDE was more effective in raising benefit take-up for benefits that are harder to apply to. This confirms that the meetings contribute to the reduction of application costs.

The only direct evidence that we have on the role of welfare stigma comes from the anti-stigma invitation letter, which we showed had no effect on meeting attendance or benefit take-up. This finding suggests that welfare stigma may not be important in our context but that stigma could also be so deeply entrenched that a simple letter may not be sufficient to alleviate it, so the evidence is inconclusive. In a similar vein, we can test the role of information through the information leaflet, which provided general information about the existence of benefits and amounts received by households in different situations. We showed that it did not have any impact on meeting attendance or benefit take-up. This test is, however, a weak one because the information in the leaflet may have been too general to induce take-up.

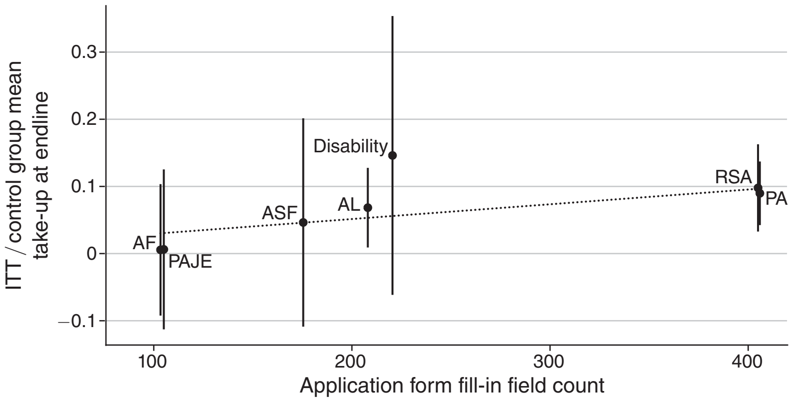

Figure 1. Treatment Effect on Take-up and Complexity of Application for Each Benefit

Notes: The figure presents the relationship between the effect of the invitation to the meeting with social services and the number of fields to fill in the application form for each benefit. We exclude health benefits and focus on benefits that are directly managed by the CNAF social services that held the meetings. AF, PAJE, and ASF are family benefits; AL is housing benefits; and RSA and PA are income support benefits. We plotted 90 percent confidence intervals.

B. *Evidence from the Online Simulator*

To test the role of information more precisely, we leverage the second nationwide experiment, which invited a random sample of job seekers to evaluate their personal eligibility to social benefits using an online simulator provided by the government. The simulator provided customized and detailed information at a small time and effort cost but did not provide any help in benefit application. We see this as a good test of the effect of "pure" information content delivered in a credible and neutral manner. From the 20,000 job seekers in the treatment group, 67 percent opened the invitation email, 39 percent clicked on the link in the email to start a simulation immediately, and 18 percent completed a simulation in the same session. By contrast, simulator usage in the control was negligible: Since randomization took place at the level of unemployment agencies, we can use geolocalized connection data at the zip code level to check that there were almost no connections in control areas. Indeed, at the time of the experiment, only social workers used the simulator, which was not publicized until nine months later. The net compliance rate to this treatment is thus comparable to that of the main experiment. Out of the 18 percent completed simulations, the logs of the computations that were run indicate that most job seekers were eligible for some benefit (16 percent of the total sample) and, in particular, for health benefits (14 percent). These results suggest that the information was well received and well targeted; in fact, these numbers likely underestimate simulator usage since we can only track the first connections.

As Table 3 shows, however, the information only had a very small effect on benefit take-up: The estimate for any benefit is a very small and insignificant 0.29 percentage points (0.6 percent) 8 months after the emails were sent; the point estimate is even negative for family benefits. Supplemental Appendix Table B20 confirms this result using the number of benefits as an alternative outcome. Taken together, these results

Table 3—Effect of the Invitation to the Online Simulator

	Any benefit received over eight months			
	Any benefit	Family	Health	Income
Any invitation	0.0029	−0.0041	0.0058	0.0039
	(0.0064)	(0.0059)	(0.0048)	(0.0045)
Control group mean	0.4944	0.3898	0.1286	0.1635
Observations	40,000	40,000	40,000	40,000

Notes: This table presents the ITT estimates of being invited to use the online simulator on receiving any benefit measured eight months after the start of the intervention. Controls are selected using the double-lasso method and omitted for brevity. Standard errors are robust to heteroskedasticity and clustered at the unemployment agency level. All specifications contain strata fixed effects.

suggest that providing personalized information—and only information—to job seekers who are likely eligible for benefits does not have a large effect on take-up overall.[18]

This evidence does not completely rule out that low benefit take-up is due to a lack of information on potential benefits, however. First, the only outcome we have for the online simulator is identical to the *stock* measure defined for the main experiment rather than the *preferred* one. The corresponding effects of the RDVDE are 0.68 percentage points (1.4 percent) for any benefits and 0.82 percentage points (2.3 percent) for family benefits (Supplemental Appendix Table B10). Given the standard errors, we cannot statistically reject that the two interventions had the same effects. Second, it could be that the meetings are more effective because social workers provide information more credibly or more clearly to job seekers than the simulator, but we cannot test this in our context. In Supplemental Appendix C we use a simple model to clarify that a small effect of information provision could reflect (i) that job seekers are not well informed individually but on average hold accurate beliefs about their eligibility or (ii) that job seekers systematically underestimate their eligibility but the bias is small relative to transaction costs, so correcting them does not visibly change the decision to apply. We test hypothesis (i) in the next section.

C. *Success Rate of Applications*

If job seekers are misinformed about their eligibility but do not on average overestimate or underestimate them, then providing information about potential benefit eligibility would not increase the probability of applying to benefits but might make them more successful. The model in Supplemental Appendix C makes this argument explicit: It shows that better information on eligibility improves self-selection into applying, which should improve the success rate of applications. Based on the phone survey in the RDVDE experiment, we can test this

[18] We can assume that there are job seekers eligible for benefits because (i) the log of simulations indicates it and (ii) this is a similar population as in the main experiment, where we find large effects on take-up.

TABLE 4—EFFECT OF THE INVITATION ON APPLICATION SUCCESS RATES

	Application success rate			
	All benefits	Family benefits	Health benefits	Income benefits
Any invitation	−0.0218	−0.0797	0.0099	−0.0305
	(0.0137)	(0.0283)	(0.0374)	(0.0200)
Control group mean	0.8950	0.9060	0.9230	0.8790
Observations	2,552	1,015	543	1,606

Notes: This table presents the ITT estimates of being invited to the RDVDE meeting on the success rate of benefit applications. The success rate is computed as the number of new benefits received based on the preferred measure from the RNCPS divided by the sum of the number of new benefits received and the number of unsuccessful applications measured using the phone survey. The preferred measure is a dummy variable equal to 1 if the individual was registered as receiving any benefit from the corresponding benefit group at any point since the start of the intervention and that benefit was not received at baseline, and 0 otherwise. The total number of new benefits is the sum, within the benefit group, of the preferred measures for each benefit in that group. Controls are selected using the double-lasso method. Standard errors are robust to heteroskedasticity. All specifications contain strata fixed effects. Control coefficients are omitted for brevity.

using information on benefit applications that did not lead to the receipt of benefits, because they were either rejected, pending, incomplete, or abandoned. We construct the application success rate by dividing the number of benefits opened according to the administrative data after three months (i.e., around the time of the survey) by the sum of the number of benefits opened and the number of rejected applications according to the survey.[19]

Table 4 presents the treatment effect on the success rate of applications. The baseline rate is high, about 90 percent, which implies that control individuals do not strongly overestimate their eligibility.[20] Column 1 in Table 4 shows that the success rate of applications did not improve overall: The coefficient is negative and insignificant. The success rate of applications for family benefits declined significantly (column 2). These results suggest that meetings with social workers did not mostly increase take-up by providing more accurate information on benefit eligibility. Instead, the small decrease in the success rate of application could be explained if the meetings encouraged job seekers with lower or weaker claims to apply by lowering transaction costs. The next section provides evidence on the negative selection of compliers.

[19] By construction, this measure is missing for people who did not apply; thus, the number of nonmissing observations in Table 4 is 2,225 (for all benefits), although survey respondents are 9,969.

[20] By contrast, Finkelstein and Notowidigdo (2019) observe high rejection rates (23 percent compared to 10 percent here) and find strong effects of information provision on take-up.

V. Targeting

A. *Amounts*

Our findings so far suggest that many eligible job seekers do not take up social benefits to a large extent because of transaction costs and that these costs can be lowered by meetings with social services with large positive effects on take-up. The key question that remains is whether the intervention improved the targeting of social benefits—that is, whether it helped people with higher claims to apply. The model in Supplemental Appendix C illustrates that, by decreasing transaction costs, the meeting lowers the threshold of benefit amounts for which one is ready to apply. Thus, in general, the average value of claims should decrease with treatment. We now take this prediction to the data.

We were able to collect information on benefit amounts for only 58 percent of our sample, composed of the job seekers who were already known to social services because they had taken up family or income benefits at some point prior to the baseline. We estimate the ITT effect of being invited to the RDVDE on this population on three outcomes: the probability of receiving any new benefit since the invitation, the amount received conditional on receiving a benefit, and the total amount received, giving zero to people who did not receive any benefit. Panel A in Table 5 shows that in this subsample, the effect of the invitation on the take-up of benefits is slightly smaller but quite comparable to that in the full sample (0.0155 versus 0.0174; see Table 2). In panel B, we report the effect of the invitation on average amounts received conditional on claiming them, which is negative, as expected if the treatment reduces transaction costs. In other words, the sample in panel B is made of inframarginal claimants who receive benefits in the absence of the intervention and marginal claimants who take up benefits only if they get invited (and subsequently attend the meeting). The control group mean in panel B provides the average value of amounts received by the inframarginal claimants (€158 monthly). Those who receive benefits in the treatment group receive €5.7 less on average; because this group is composed of inframarginal and marginal individuals, we can recover the average amounts received by the marginal claimants: It is €91 less than the inframarginal individuals (bootstrapped 95 percent confidence interval $[-317, 28]$).[21] Thus, the intervention attracted potential beneficiaries with substantially lower claims, which they probably wouldn't think worth claiming if they did not receive some assistance. This finding is consistent with Finkelstein and Notowidigdo's (2019) findings on SNAP recipients.

Finally, in panel C we estimate the effect of being invited to the meeting on benefits received attributing zero values to nonclaimants, which combines the two previous effects. We find that each invitation resulted in an unprecisely estimated €1.26 of new income benefits per month, which reflects a higher take-up of lower

[21] Let B denote the value of benefits, and I inframarginal and M marginal recipients, with $p(k)$, $k = I, M$ their proportions in the population. Then the control group mean in panel B is $E[B \mid I, B > 0]$ and the coefficient is $p(I)E[B \mid I, B > 0] + p(M)E[B \mid M, B > 0] - E[B \mid I, B > 0]$, from which we recover $E[B \mid M, B > 0]$ using the estimates of $p(I)$ and $p(M)$ in panel A.

Table 5—Effect of Receiving an Invitation to the RDVDE on Monthly Benefit Amounts

	Any benefit	Family benefit	Income benefit
Panel A. Benefits take-up			
Any invitation	0.0155	0.0012	0.0154
	(0.0048)	(0.0031)	(0.0043)
Control group mean	0.2310	0.0797	0.1721
Observations	31,444	31,444	31,444
Panel B. Amounts received conditional on take-up			
Any invitation	−5.7192	−10.5496	−1.9391
	(3.9593)	(6.6148)	(4.0261)
Control group mean	158.4604	171.3916	133.3288
Observations	7,556	2,526	5,713
Panel C. Benefits amounts received			
Any invitation	1.2589	−0.6397	1.8737
	(1.1963)	(0.7109)	(0.9113)
Control group mean	36.6055	13.6613	22.9442
Observations	31,444	31,444	31,444

Notes: This table presents the ITT estimates of being invited to the RDVDE on three different measures of benefit take-up. The measures are the preferred dummy variable take-up measure (panel A), the monthly amount of benefits received conditional on the benefit being newly received as per the preferred measure (panel B), and the total amounts of benefits received attributing a zero value to nonclaimants (panel C). Regressions are estimated on the sample known to services at baseline for which benefit amount data are available. All outcome variables are measured six months after the start of the intervention. Controls for the specifications are selected using the robust double-lasso method; coefficients are omitted for brevity. All specifications contain strata fixed effects. Standard errors are robust to heteroskedasticity.

amounts. Rescaled by the share who attended the meeting, this ITT effect corresponds to a gain of €6 *per meeting*, which is equal to about 15 percent of the control group mean. More precise figures are estimated for income benefits, where the gain per meeting would be about 39 percent compared to the control group. Hence, overall, the analysis of benefit amounts confirms some positive average effects on benefits, but they are limited by the strong negative targeting of marginal beneficiaries.

B. *Marginal Treatment Effect*

Our results so far have focused on the effects on compliers—that is, people who decided to attend the RDVDE. Since compliers are self-selected, the effects of the meeting could be different for them and for the rest of the population. As an illustration, Supplemental Appendix Table B13 and Supplemental Appendix Table B14 present heterogeneous effects of the invitation on attendance and benefit take-up, respectively, along four dimensions that we prespecified in the pre-analysis plan. Job seekers known to social services (CNAF beneficiaries, column 3) had both smaller effects on benefit take-up from the invitation than the rest of the population and a higher propensity to attend the meeting. Therefore, the effect of the meeting for them (i.e., the Wald estimator) is smaller. This is a situation where those most likely to attend the meeting had a lower benefit from it. We find a similar pattern for job seekers who were at the end of unemployment benefits (Population 2, column 2), who were less likely to attend the meeting but had higher treatment effect on benefit take-up (the latter difference is not statistically significant).

To explore this issue more systematically, we apply the MTE method to estimate how the effect of the meeting varies with the propensity to attend. We first estimate the propensity score using specification (5) (Supplemental Appendix Table B17). Variables selected by the lasso to be part of the propensity score are the main sociodemographic characteristics (citizenship, gender, age, and marital situation) and variables that capture the reasons for and the duration of unemployment, as well as unemployment benefits received, salary in previous jobs, and the type of work sought. We interact these variables with two treatment dummies: any invitation and phone invitation. Several interactions are significant: For instance, invitations increased attendance more for women, older individuals, foreigners, those with lower unemployment benefits and lower pre-unemployment salaries, those who had reentered the labor market after an illness or maternity leave, and those who were previously known to social services. All of them thus have a lower net cost of attending the meetings and are more likely to belong to the complier population. None of the interaction terms with the phone invitation dummy are significant on their own, but the p-value of the F-test of joint significance is 0.05. As Figure A1 in the Supplemental Appendix shows, the combination of the treatments and observable characteristics induces variations in the predicted propensity to attend the meeting from 0 to about 50 percent. We can thus estimate the MTE on a support of individuals with very low to medium meeting attendance cost.

We then estimate the second stage, following specification (6), which regresses measures of benefit receipt on the propensity score, controls, and the propensity score interacted with controls. Table 6 presents the coefficients on the propensity score (coefficients a and b of equation (6)). We find a positive and significant coefficient on the square of the propensity score for income benefits only; it is small and insignificant for family and health benefits, and for all benefits together. We will hence focus on income benefits for the rest of the analysis.

We start with the unobserved component of the MTE, which is the derivative of the second-stage equation with respect to the propensity score: They represent $a + \bar{\mathbf{X}}'\mathbf{c} + 2b\widehat{p(\mathbf{X}_i, \mathbf{Z}_i)}$, where the intercept is estimated at the average sample value of the $\mathbf{X}$s. This derivative identifies how the treatment effect varies as the unobserved cost of attending the meeting increases (or the probability to be a complier decreases). Figure 2, panel A presents the unobserved component of the MTE for income benefits as a function of the propensity score (see Supplemental Appendix Figures 2, panels A–C, for other benefits). The MTE is positively sloped; it increases from a treatment effect of about zero for those with the lower attendance costs to an effect of about 0.25 for those with the highest costs. In comparison, in Section III we estimated a LATE of 0.07 on the take-up of income benefit, which suggests that there is strong heterogeneity in the treatment effect along the attendance cost dimension.

Table 6—Marginal Treatment Effect on Benefit Take-up

	Any new benefit take-up over six months			
	Any benefit	Any family benefit	Any health benefit	Any income benefit
Propensity score	0.0682	0.0822	0.0429	−0.0673
	(0.0655)	(0.0468)	(0.0405)	0.0531
Propensity score^2	0.0917	−0.0253	−0.0305	0.2156
	(0.1381)	(0.1040)	(0.0909)	(0.1168)
Observations	53,795	53,795	53,795	53,795

Notes: This table presents the MTE estimates on benefit take-up, where take-up is measured at six months after baseline using the preferred measure. The preferred measure is a dummy variable, which is 1 if the individual was registered as receiving any benefit from the corresponding benefit group at any point since the start of the intervention and that benefit was not received at baseline, and 0 otherwise. The propensity score is obtained by a linear probability model of meeting attendance based on treatment assignment and controls selected by double lasso. The estimates of the first stage are presented in Supplemental Appendix Table B17. Standard errors are obtained through 500 bootstrap replications of both stages. All specifications include strata fixed effects and controls used in the propensity score interacted with treatment probability.

To get a sense of the range of monetary gains from the meeting, we can use the subsample known to social services for which we observe amounts received. We reproduce the MTE, using as outcome the monthly average of the cumulative amount received from benefits not present at baseline, with zeros when there was no new benefit received. Supplemental Appendix Figure A3, panel C shows that attending the meeting can increase income benefits received by more than €20 per month for job seekers with the highest unobserved cost of attending the meeting, which would represent a doubling of benefits received as compared to the control group mean of €22 and is more than three times larger than the LATE of €6 derived in the previous section.[22]

We next turn to the observed component of the MTE: Its identification is more direct than the unobserved component. We first compute the average treatment effect (ATE) for each individual with observed characteristics $\mathbf{X}_i$ as $\left(\mathbf{X}_i'\mathbf{c}\right)$ from equation (6). We then compute the propensity score for every individual, and correlate the two. To present this graphically, we have formed 50 bins of the propensity score (of 0.01 points width) and computed the average $\left(\mathbf{X}_i'\mathbf{c}\right)$ in each bin. Figure 2, panel B shows a mostly negative pattern: Individuals who are most likely to attend the meeting based on their observed characteristics have smaller predicted treatment effects.[23] Hence, the observed component of the MTE tells the same story as the unobserved component (see Supplemental Appendix Figure 2, panels E–G, for other benefits).

[22] Supplemental Appendix Figure A4 shows that the MTE results on take-up for the subsample known to social services are similar to the whole sample (Figure 2).

[23] Values of the MTE in that figure are relative effects, not accounting for the constant and the term in $2b\overline{p(\mathbf{X}_i, \mathbf{Z}_i)}$.

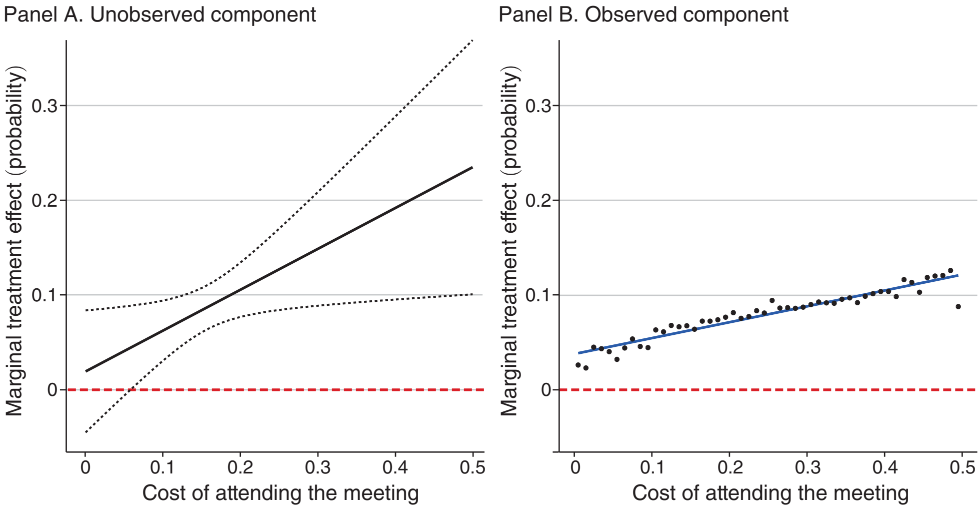

Figure 2. Marginal Treatment Effect on Income Benefit Take-up

Notes: Panel A presents the relationship between the MTE on any new income benefit take-up and the unobserved cost from participating to the meeting with social services. The x-axis measures the cost of attending the meeting. Take-up is measured using the preferred measure, a dummy variable that is 1 if the individual received any income benefit at any point since the start of the intervention and that benefit was not received at baseline, and 0 otherwise. It is measured over six months since the start of the intervention. The linear relationship is obtained from the estimates in Table 6, evaluated at the mean of the relevant sample covariates. Ninety percent confidence intervals are plotted. Standard errors are obtained through 500 bootstrap replications of the first and second stages. See Section II for more details. Panel B presents the relationship between the observable component of the MTE (i.e., the treatment heterogeneity based on observables) and the cost of attending the meeting (x-axis). The observed component of the MTE is computed as $\mathbf{X}_i'\mathbf{c}$ from equation (6). The cost of attending the meeting is the propensity score transformed by $0.5 - \widehat{p(\mathbf{X}_i, \mathbf{Z}_i)}$, in order to facilitate comparison with panel A. We have formed 50 bins of the propensity score (of 1-point width) and computed the averages $\mathbf{X}_i'\mathbf{c}$ in each bin, which are represented by the black points. The blue line is a linear adjustment.

The MTE can further be used to compute weights to estimate policy parameters of interest such as the ATE, the average treatment on the treated (ATT), and the average treatment on the untreated (TUT) (Heckman and Vytlacil 2005; Cornelissen et al. 2016). Table 7 shows the estimates of those parameters. First, notice that the LATE reconstructed from those weights is equal to the ATT (which is expected given that only invited job seekers can attend the meeting) and close to our 2SLS estimate.[24] Most interestingly, the ATT is much lower than the TUT, which confirms that individuals attending the meeting are not those who would benefit most from it. Orders of magnitude are large: While the meeting increases the probability to take a benefit by 7 percentage points on average among those who attended, it would increase it by 19 points among those who did not. The ATE is also much larger than our initial LATE estimate, at about 18 percentage points.

Overall, our findings suggest that claimants with higher costs of attending the meeting would gain the most from it. Through the lens of our representation of compliance and take-up, job seekers who have a high cost of attending the meetings

[24] LATE and 2SLS may not be identical, because they are based on somewhat different specifications, and the weights computation introduces measurement error.

Table 7—Treatment Effect Parameters for Income Benefit Take-up

	Any new income benefit take-up over six months				
	ATE	ATT	TUT	LATE	2SLS
Attended meeting	0.1776	0.1254	0.1888	0.125	0.0704

Notes: This table presents the average treatment effect (ATE), the average treatment on the treated (ATT), the average treatment on the untreated (TUT), and the LATE, computed using our estimate of the linear MTE. Column 5 also recalls the 2SLS estimator from Table 2.

only attend if they anticipate large benefits relative to the transaction costs of application. In that case, they would tend to claim the benefit even without support from social workers, and the treatment impact would be low for them. Our results point to another likely feature: that job seekers with high costs of attending also have high transaction costs to claim benefits on their own and, thus, tend not to apply without the treatment. When they do attend the meeting, application assistance is particularly effective for them. Our results for job seekers known to social services support this interpretation. They have lower costs of going to the CNAF agency, which explains their higher compliance rate; they also face lower transaction costs in applying to benefits, which explains that the meetings are less useful for them. Our MTE results suggest that this holds more generally: The distribution of compliance and transaction costs is such that self-selection tends to bring the wrong people to the meeting.

VI. Conclusion

This paper reports on two nationwide randomized experiments aimed at increasing social benefit take-up in France. The first invited job seekers to a meeting with social services to evaluate their eligibility for a wide range of social benefits and to help them apply. The impact of the meetings on compliers was large: Benefits opened within six months of the meeting increased by +44 percent for income benefits and +21 percent for family and housing benefits. By contrast, a companion experiment that provided personalized eligibility information via an online benefits simulator had no impact on take-up. Based on this second experiment and on additional results from the first, we argue that application costs rather than lack of information seem to be the main driver of non take-up in our context. In the final part of the paper, we show that the marginal claimant induced by the lowering of application costs had slightly lower claims, which suggests that among compliers, application costs may worsen targeting. However, when we estimate the MTE, we find that the people who stood the most to gain from the meeting were less likely to attend it, which implies that the same costs that deter eligible individuals from applying to social benefits also prevent them from accessing assistance to help them apply.

REFERENCES

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

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

Alberola, E., J. Muller, and C. Maes. 2018. *Evaluation qualitative de l'expérimentation du rendez-vous des droits élargi*. CREDOC.

Banerjee, Abhijit, Esther Duflo, Clément Imbert, Santhosh Mathew, and Rohini Pande. 2020. "E-governance, Accountability, and Leakage in Public Programs: Experimental Evidence from a Financial Management Reform in India." *American Economic Journal: Applied Economics* 12 (4): 39–72.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014a. "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives* 28 (2): 29–50.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014b. "Inference on Treatment Effects after Selection among High-Dimensional Controls." *Review of Economic Studies* 81 (2): 608–50.

Bettinger, Eric P., Bridget T. Long, Philip Oreopoulos, and Lisa Sanbonmatsu. 2012. "The Role of Application Assistance and Information in College Decisions: Results from the HR Block FAFSA Experiment." *Quarterly Journal of Economics* 127 (3): 1205–42.

Bhargava, Saurabh, and Dayanand Manoli. 2015. "Psychological Frictions and the Incomplete Take-up of Social Benefits: Evidence from an IRS Field Experiment." *American Economic Review* 105 (11): 3489–529.

Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Løken, and Magne Mogstad. 2020. "Incarceration, Recidivism, and Employment." *Journal of Political Economy* 128 (4): 1269–324.

Brinch, Christian N., Magne Mogstad, and Matthew Wiswall. 2017. "Beyond LATE with a Discrete Instrument." *Journal of Political Economy* 125 (4): 985–1039.

Castell, Laura, Marc Gurgand, Clément Imbert, and Todor Tochev. 2018. "Information and Take-up of Social Benefits in France." AEA RCT Registry. https://doi.org/10.1257/rct.1882-3.0.

Castell, Laura, Marc Gurgand, Clément Imbert, and Todor Tochev. 2025. *Data and Code for:* "*Take-up of Social Benefits: Experimental Evidence from France*." Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E209001V1.

Christensen, Julian, Lene Aarøe, Martin Baekgaard, Pamela Herd, and Donald P. Moynihan. 2020. "Human Capital and Administrative Burden: The Role of Cognitive Resources in Citizen-State Interactions." *Public Administration Review* 80 (1): 127–36.

CNAF. 2017. *Mes-aides participants benefit usage*. Caisse nationale des allocations familiales.

CNAF. 2018. *Rendez-vous des droits élargi participant benefit usage*. Caisse nationale des allocations familiales.

CNAF. 2021. *Rendez-vous des droits élargi benefits amounts*. Caisse nationale des allocations familiales.

Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schöenberg. 2016. "From LATE to MTE: Alternative Methods for the Evaluation of Policy Interventions." *Labour Economics* 41 (1): 47–60.

Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schöenberg. 2018. "Who Benefits from Universal Child Care? Estimating Marginal Returns to Early Child Care Attendance." *Journal of Political Economy* 126 (6): 2356–409.

Currie, Janet. 2004. "The Take-up of Social Benefits." NBER Working Paper 10488.

Decobecq, Laurence. 2013. *Convention d'Objectifs et de Gestion entre l'Etat et la Cnaf 2013–2017*. Ministère de L'Économie et des Finances. La Gazette Des Communes.

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

Direction de la Sécurité Sociale. 2017. "Historique du Système Français de Sécurité Sociale." http://www.securite-sociale.fr/Historique-du-systeme-francais-de-Securite-sociale (accessed March 31, 2017).

Domingo, Pauline, and Muriel Pucci. 2012. "Les Non-recourants au RSA." *CNAF l'e-ssentiel* 11: 124.

Drees. 2016. *Mes-aides participants.* Directorate of Research, Studies, Evaluation and Statistics. https://drees.solidarites-sante.gouv.fr/ (accessed 2016).

Drees. 2018a. *Rendez-vous des droits élargi participant*. Directorate of Research, Studies, Evaluation and Statistics. https://drees.solidarites-sante.gouv.fr/ (accessed 2018).

Drees. 2018b. *Rendez-vous des droits élargi post-meetings*. Directorate of Research, Studies, Evaluation and Statistics. https://drees.solidarites-sante.gouv.fr/ (accessed 2018).

Dubois, Vincent. 2003. *La vie au guichet*. Économica.

Duflo, Esther, and Emmanuel Saez. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics* 118 (3): 815–42.

Finkelstein, Amy, and Matthew J. Notowidigdo. 2019. "Take-up and Targeting: Experimental Evidence from SNAP." *Quarterly Journal of Economics* 134 (3): 1505–56.

Friconneau, Claude. 2014. "La Contribution du Répertoire National de la Protection Sociale à la Détection des Droits Potentiels et des Droits Ouverts Indûment." *Regards* 46 (2): 115–22.

Ganong, Peter, and Jeffery B. Liebman. 2018. "The Decline, Rebound, and Further Rise in SNAP Enrollment: Disentangling Business Cycle Fluctuations and Policy Changes." *American Economic Journal: Economic Policy* 10 (4): 153–76.

Goldsmith-Pinkham, Paul, Peter Hull, and Michal Kolesár. 2024. "Contamination Bias in Linear Regressions." Preprint, arXiv. https://doi.org/10.1257/aer.20221116.

Gupta, Sarika. 2017. "Perils of Paperwork: The Impact of Information and Application Assistance on Welfare Program Take-up in India." Unpublished.

Heckman, James J. 2010. "Building Bridges Between Structural and Program Evaluation Approaches to Evaluating Policy." *Journal of Economic Literature* 48 (2): 356–98.

Heckman, James J., and Edward Vytlacil. 2005. "Structural Equations, Treatment Effects and Econometric Policy Evaluation." *Econometrica* 73 (3): 669–738.

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

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

Kling, Jeffery R., Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. 2012. "Comparison Friction: Experimental Evidence from Medicare Drug Plans." *Quarterly Journal of Economics* 127 (1): 199–235.

Ko, Wonsik, and Robert A. Mofftt. 2022. "Take-up of Social Benefits." NBER Working Paper 30148.

Kopczuk, Wojciech, and Cristian Pop-Eleches. 2007. "Electronic Filing, Tax Preparers and Participation in the Earned Income Tax Credit." *Journal of Public Economics* 91 (7–8): 1351–67.

Linos, Elizabeth, Allen Prohofsky, Aparna Ramesh, Jesse Rothstein, and Matthew Unrath. 2022. "Can Nudges Increase Take-up of the EITC? Evidence from Multiple Field Experiments." *American Economic Journal: Economic Policy* 14 (4): 432–52.

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." *Science* 341 (6149): 976–80.

Mofftt, Robert. 1983. "An Economic Model of Welfare Stigma." *American Economic Review* 73 (5): 1023–35.

Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity: Why Having Too Little Means So Much*. Macmillan.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. "Building State Capacity: Evidence from Biometric Smartcards in India." *American Economic Review* 106 (10): 2895–929.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2025. "Identity Verification Standards in Welfare Programs: Experimental Evidence from India." *Review of Economics and Statistics* 107 (2): 372–92.

Nichols, Albert L., and Richard J. Zeckhauser. 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review: Papers and Proceedings* 72 (2): 372–77.

Plueger, Dean. 2005. *Earned Income Tax Credit Participation Rate for Tax Year 2005*. Internal Revenue Service.

Tempelman, Caren, and Aenneli Houkes-Hommes. 2016. "What Stops Dutch Households from Taking Up Much Needed Benefits?" *Review of Income and Wealth* 62 (4): 685–705.

Warin, Philippe. 2012. "Le Non-recours aux Droits. Question en Expansion, Catégorieen Construction, Possible Changement de Paradigme dans la Construction des Politiques Publiques." *SociologieS*. https://doi.org/10.4000/sociologies.4103.