# At-Risk or Responsive?
# Targeting in Unemployment Policy

Sebastian Y. D. Ernst — LSE
Andreas I. Mueller — University of Zurich
Johannes Spinnewijn — LSE
Joost Bollens — VDAB*

March 26, 2026

**Abstract**

This paper evaluates the use of algorithmic risk profiling to direct job search support toward individuals most at risk of long-term unemployment. Using three regression discontinuity designs and rich administrative data from the Flemish Public Employment Service (VDAB) in Belgium, we address two key dimensions: who and when to target. First, job search counseling raises employment on average, but treatment effects decline sharply with predicted risk. This reveals a fundamental trade-off between targeting those most at risk and those most responsive to intervention. Second, we find no evidence of depreciation in treatment effects at the individual level, but strong dynamic selection: as spells lengthen, surviving job seekers have both higher baseline risk and lower treatment responsiveness, generating a trade-off between the timing of treatment and the composition of the remaining pool. We develop and calibrate a conceptual framework for optimal targeting whose key sufficient statistic is the ratio of predicted treatment effects to predicted job-finding probabilities. Targeting rules combining both objects substantially outperform the status quo, yielding welfare gains nearly three times larger than random assignment, while targeting on risk scores alone performs worse than random targeting.

**JEL Codes:** J64, J65, H43, C21
**Keywords:** unemployment, targeting, treatment effect heterogeneity, regression discontinuity, public employment services, machine learning

*Addresses: Ernst: London School of Economics, Department of Economics, Houghton Street, London WC2A 2AE (email: s.y.ernst@lse.ac.uk). Mueller: University of Zurich, Department of Economics, Schönberggasse 1, CH-8001 Zurich (email: andreas.mueller@econ.uzh.ch); Spinnewijn: London School of Economics, Department of Economics, Houghton Street, London WC2A 2AE (email: j.spinnewijn@lse.ac.uk); Bollens: VDAB, Keizerslaan 11, BE-1000 Brussels (email: joost.bollens@vdab.be).

# 1 Introduction

A central challenge for public institutions operating under resource constraints is how to allocate interventions where they will do the most good. Public employment services (PES) across many countries have turned to algorithmic risk profiling to address this challenge: by predicting which job seekers are most at risk of long-term unemployment, they aim to direct caseworker support toward those who need it most [Desiere et al., 2019]. Yet the logic of this approach rests on an assumption that rarely goes tested, namely that individuals most in need of assistance are also the ones who benefit most from it. This paper shows that the assumption fails along two distinct dimensions. First, risk and responsiveness to intervention are negatively correlated, so those most at risk benefit the least from job search counseling. Second, dynamic selection over the unemployment spell compounds this problem. As spells lengthen, the surviving pool becomes less responsive, so that delaying intervention reaches an increasingly unresponsive group.

We study the targeting of job search counseling in Flanders, Belgium, where the Flemish Public Employment Service (VDAB) has since September 2018 used a machine-learning model with over 400 input variables to predict job seekers' probability of finding employment within six months. Job seekers whose predicted probability falls below a threshold of 0.65 are prioritized for contact with a caseworker, while those above the threshold are classified as self-reliant and not actively contacted. The VDAB's system exemplifies the international trend toward algorithmic profiling in active labor market policy. We study its consequences using administrative data covering the universe of registered unemployment spells between September 2018 and November 2022 linked to rich information on labor market histories, risk scores, caseworker interactions, and employment outcomes. This combination of institutional detail, data richness, and quasi-experimental variation provides an unusually clean setting in which to evaluate the welfare consequences of risk-based targeting.

We investigate two main components of targeting. The first and most fundamental is *who* should be targeted. Current practice targets those with the lowest predicted probability of job finding and thus most at risk of long-term unemployment. But optimal targeting requires identifying those for whom the intervention yields the greatest return, which depends jointly on responsiveness to treatment and the cost of remaining unemployed. To address this question, we leverage the rule in Flanders that job seekers are automatically reassigned to their previous caseworker only if their prior unemployment spell ended within the past six months, and estimate a regression discontinuity (RD) design around the 6 month threshold. Exploiting this design, we estimate highly heterogeneous treatment effects across the distribution of long-term unemployment risk. The effect of caseworker assignment on days worked is a precise zero for job seekers with the lowest predicted job-finding rates (the most at-risk group), while large and statistically significant effects are estimated for job seekers with higher predicted job-finding probabilities. This finding represents a sharp trade-off, as the individuals currently prioritized for intervention respond the least to it.

The second component we examine is *when* to target. To address this question, we exploit two further discontinuities in the VDAB's caseworker assignment process. After five weeks of unemployment, those with predicted job-finding rates below 0.65 are called by the service line and potentially

assigned to a caseworker. After a further three months, all remaining unassigned job seekers are contacted regardless of their risk score, reversing the original discontinuity for those above the threshold. By comparing treatment effect estimates at these two points in the unemployment spell for similar groups of job seekers (those near the 0.65 threshold), we find no evidence of true duration dependence in treatment effects at the individual level, as the estimated treatment effects are very similar. However, we find evidence for strong dynamic selection, as those who remain unemployed longer are both higher risk and less responsive to treatment. This creates a second trade-off, as delaying treatment may better target high-risk individuals, but it reaches a less responsive pool.

To evaluate these trade-offs jointly, we develop a conceptual framework that formalizes the optimal targeting problem. The framework accommodates observed and unobserved heterogeneity in both job-finding probabilities and treatment effects, as well as duration dependence in job finding and treatment efficacy. The key insight is that the return to treating a given job seeker depends on both the magnitude of their treatment effect and the continuation loss from unemployment they would otherwise face, with the latter being larger for higher-risk individuals. A formal analytical result shows that the optimal targeting index is approximately proportional to the ratio of predicted treatment effects to baseline job-finding probabilities. In other words, the index places positive weight on both responsiveness and risk, and thus captures the tension at the heart of the targeting problem.

We calibrate the framework to the Flemish data and evaluate several alternative targeting policies. The welfare gains of optimal targeting are nearly three times those of random assignment, reflecting substantial heterogeneity in both risk and responsiveness across job seekers. A simple rule based on the ratio of predicted treatment effect to predicted job finding risk performs nearly as well as the fully optimal policy. By contrast, the current practice of targeting exclusively on risk scores performs worse than random targeting. This is not merely a theoretical possibility, as in our calibrated model, risk-score targeting yields only about half the welfare gains of random assignment. The reason is that risk scores and treatment effects are strongly negatively correlated in our setting, so that targeting the lowest scorers systematically selects the individuals with the smallest returns to intervention. Risk alone is therefore not just an insufficient but in our context counterproductive targeting criterion.

We also use the framework to characterize *when* during the unemployment spell it is optimal to intervene. The welfare gains from optimal targeting relative to random assignment decline strongly over the first few months of the spell. This pattern reflects that the scope for selective targeting is greatest at the start of the spell, when the pool of unemployed workers is large and heterogeneous and the optimal policy can concentrate treatment on those with the highest ratio of predicted treatment to predicted job finding. As the spell progresses, dynamic selection erodes this heterogeneity. Fewer workers remain unemployed, and those who do are increasingly similar in both risk and responsiveness, so that the optimal policy converges toward treating the average surviving worker. These results imply that employment agencies should prioritize early intervention and invest in identifying responsive individuals at the outset of the spell, when the returns to selective targeting are highest.

Our paper contributes to several strands of literature. It relates to a growing literature on heterogeneity in job-finding and the predictability of long-term unemployment [Ahn and Hamilton,

2020; Mueller et al., 2021; Ahn, 2023; Alvarez et al., 2024; Mueller and Spinnewijn, 2025]. Our finding that job search assistance has higher treatment effects for low-risk individuals is consistent with the broader finding in the active labor market program literature that treatment effects of job search counseling are larger for the short-term unemployed than the long-term unemployed [Koning, 2009; Liu et al., 2014; Card et al., 2018], while training and information-only programs tend to exhibit the opposite pattern [Altmann et al., 2018; Card et al., 2018; Altmann et al., 2022; Belot et al., 2026]. Our analysis relates to recent work on treatment targeting under heterogeneous effects [Kitagawa and Tetenov, 2018; Athey et al., 2025; Yadlowsky et al., 2025], and most closely to Haushofer et al. [2025], who study the same trade-off between targeting deprivation and targeting impact in the context of cash transfer programs in Kenya. The latter finds that targeting solely on the basis of deprivation is unattractive unless the planner's redistributive preferences are very strong, a conclusion that mirrors our own in the context of unemployment policy. We also build on the growing literature on algorithmic profiling in public employment services [Desiere et al., 2019] by providing causal evidence on the welfare consequences of risk-based targeting and characterizing the optimal alternative. While our empirical findings are specific to the context studied, our conceptual framework characterizes the trade-off between risk and responsiveness more broadly, and can be applied to the targeting of any active labor market program. This includes training programs, where the trade-off likely takes a different form given their tendency to exhibit larger effects for higher-risk individuals.

The remainder of the paper is organized as follows. Section 2 develops the conceptual framework and characterizes optimal targeting. Section 3 describes the institutional setting and data. Section 4 estimates the average effects of job search counseling. Section 5 documents heterogeneity in treatment effects across the risk distribution. Section 6 examines duration dependence in treatment effects. Section 7 presents the quantitative evaluation of alternative targeting policies. Section 8 concludes.

# 2 Conceptual Framework

In this section, we develop a conceptual framework that clarifies both *who* should be targeted and *when* during an unemployment spell targeting is most valuable. The framework provides a simple characterization of the individual returns to treatment as a function of predicted treatment effects and baseline job-finding prospects. By linking these objects, the analysis yields transparent sufficient statistics for optimal targeting and highlights the key trade-offs that govern the allocation of scarce program resources across workers and durations.

## 2.1 Setup

We first describe the setup of our model, where we consider a cohort of unemployed workers who may differ in their response to treatment as well as their baseline job finding risk.

**Job Finding and Survival.** We assume that the job-finding probability at unemployment duration $d$ takes the following form:

$$P_d^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau) = (\pi + \varepsilon^\pi)\,\theta_d^\pi + (\tau + \varepsilon^\tau)\,\theta_d^\tau\,T_d(\pi, \tau), \tag{1}$$

where $\pi$ denotes an individual's baseline observable job-finding type, and $\varepsilon^\pi$ captures unobserved heterogeneity in baseline job-finding prospects. Similarly, $\tau$ denotes the individual's predictable treatment effect, while $\varepsilon^\tau$ represents unobserved heterogeneity in treatment responsiveness. The indicator $T_d(\pi, \tau) \in \{0, 1\}$ specifies whether an individual of observed type $(\pi, \tau)$ is treated at duration $d$. We restrict targeting rules to depend only on observable characteristics, so that treatment cannot be conditioned on the unobserved components. The superscript $T$ on $P$ denotes the job-finding probability is conditional on a specific targeting policy, which we define as a vector of treatment dummies $\{T_d(\pi, \tau)\}_{d=0}^{\bar{D}}$ for each type $(\pi, \tau)$. $\bar{D}$ is a maximum duration of unemployment, after which we do not allow for any further treatment.

To simplify the analysis, we assume that duration dependence is proportional across types. The term $\theta_d^\pi$ governs how baseline job-finding probabilities evolve with unemployment duration, while $\theta_d^\tau$ captures how treatment effects depreciate (or appreciate) over the spell. Finally, we assume that observed and unobserved components are independent, although the unobserved baseline type $\varepsilon^\pi$ and the unobserved treatment type $\varepsilon^\tau$ may be correlated. Formally, types are distributed on the unit interval according to the joint density $f(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau) = g(\pi, \tau)h(\varepsilon^\pi, \varepsilon^\tau)$.

We define the survival probability of individual of type $(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau)$ as:

$$S_d^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau) = \prod_{\delta=0}^{d-1}(1 - P_\delta^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau)),$$

where we adopt the short form $S_d(\pi, \varepsilon^\pi)$ in the absence of treatment since in this case the treatment effects are irrelevant.

**Utility losses from unemployment and cost of treatment.** We assume that unemployed individuals incur a per-period utility loss $u_d$, which may vary with the duration of the unemployment spell, i.e., $u_d = u\,\theta_d^u$. This loss captures the reduction in consumption associated with unemployment. More generally, it may also be interpreted as including unemployment insurance (UI) benefit payments, since these constitute a fiscal cost for the agency. We further assume that providing treatment entails a constant per-period cost $c$, and that there is no discounting over time.

For a given type, this implies that the expected utility cost of a spell of unemployment is:

$$L^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau) = \sum_{\delta=0} u_\delta S_\delta^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau),$$

where the superscript $T$ refers to a specific targeting policy and which is equal to $u$ times the expected

duration of unemployment in the case where $u_d = u$. The expected budgetary cost of treatment is:

$$B^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau) = \sum_{\delta=0} cT_\delta(\pi, \tau) S_\delta^T(\pi, \tau, \varepsilon^\pi, \varepsilon^\tau).$$

**Aggregation across types.** Integrating over unobserved heterogeneity $(\varepsilon^\pi, \varepsilon^\tau)$ with density $h$ gives type-level objects $L^T(\pi, \tau)$ and $B^T(\pi, \tau)$ in the natural way. We can then further aggregate across the population and write the expected utility loss and budgetary cost of an unemployment spell averaged across all job seekers as:

$$\begin{aligned} L^T &= \int\int L^T(\pi, \tau) g(\pi, \tau) d\pi d\tau \\ B^T &= \int\int B^T(\pi, \tau) g(\pi, \tau) d\pi d\tau, \end{aligned}$$

where we write L in the absence of treatment.

## 2.2 Who to target?

We consider the problem of an unemployment agency operating under a fixed budget $\bar{B}$. Given a per-period targeting cost $c$, the agency can either treat $\bar{B}/c$ individuals for one period or provide treatment to fewer individuals over multiple periods. We abstract from issues related to distortionary taxation and the optimal overall size of the program. Instead, we focus on the agency's allocation problem: how to distribute treatment across a subset of unemployed individuals given an exogenously determined budget. The agency's objective is to minimize the expected welfare losses associated with unemployment as follows:

$$\begin{aligned} \min_T \quad & \int\int L^T(\pi, \tau) g(\pi, \tau) d\pi d\tau \\ & s.t.\ \bar{B} \geq \int\int B^T(\pi, \tau) g(\pi, \tau) d\pi d\tau. \end{aligned} \tag{2}$$

Define the return to targeting for observable type $(\pi, \tau)$ as the reduction in expected utility losses over the unemployment spell per unit of treatment cost:

$$R^T(\pi, \tau) \equiv -\frac{L^T(\pi, \tau) - L(\pi)}{B^T(\pi, \tau)}. \tag{3}$$

For targeting policies restricted to individuals at duration $d$, we can prove the following proposition.

**Proposition 1.** ***Constrained optimal targeting.*** *Restrict targeting to individuals at a fixed duration $d$, $T \in T_d$, and assume that the distribution function $g(\pi, \tau)$ is continuous. Then the constrained optimal policy described in equation 2 treats all individuals of type $(\pi, \tau)$ for whom the return to targeting exceeds a threshold, i.e.*

$$T_d(\pi, \tau) = 1 \iff R^{T_d}(\pi, \tau) \geq \bar{R}^{T_d},$$

*where the threshold* $\bar{R}^{T_d}$ *is determined implicitly by*

$$\bar{B} = \int \int c\mathbb{I}\left[R^{T_d}(\pi,\tau) \geq \bar{R}^{T_d}\right] S_d(\pi) g(\pi,\tau)\, d\pi\, d\tau.$$

*Proof.* See Appendix A.1.

Note that we also assume that the effect of treatment is positive for all job seekers and thus it is always optimal to exhaust the budget for the agency. In the quantitative evaluation further below, we also characterize the returns to targeting over the spell of unemployment, but we do not characterize the optimal policy over all types and the spell of unemployment.

To build intuition for the optimal targeting policy, we consider the case where there is no unobserved heterogeneity in treatment effects. In this environment, the individual return to treatment at duration $d$ can be written as

$$R^{T_d}(\pi,\tau) = \frac{\tau_d}{c} L_{d+1}(\pi,\tau),$$

where $L_{d+1}(\pi,\tau) = \sum_{\delta=d+1}^{\infty} u_\delta\, S_{\delta|d+1}(\pi)$ denotes the continuation loss from unemployment from period $d+1$ onward and where $S_{\delta|d+1}(\pi) = \frac{S_\delta(\pi,\tau)}{S_{d+1}(\pi,\tau)}$ is the survival probability to period $\delta \geq d+1$, conditional on remaining unemployed until $d+1$. This expression implies that it is desirable to target individuals and durations with both large treatment effects and large continuation losses. We can prove the following proposition.

**Proposition 2.** ***Approximate returns to targeting under geometric heterogeneity.*** *Suppose baseline job-finding rates and treatment effects evolve geometrically:* $\pi_{d+1} = \theta_\pi \pi_d$ *and* $\tau_{d+1} = \theta_\tau \tau_d$*, with initial values* $(\pi_0, \tau_0) \equiv (\pi, \tau)$*, and assume utility losses are constant over the spell,* $u_d = u$*, and identical across job seekers. One-period targeting returns* $R_d^T$ *(defined in equation 3) are approximately proportional to treatment effects and inversely proportional to baseline job-finding probabilities,*

*(i) across types. For two individuals with initial characteristics* $(\pi,\tau)$ *and* $(\pi',\tau')$ *at the same duration* $d$*:*

$$\frac{R^{T_d}(\pi',\tau') - R^{T_d}(\pi,\tau)}{R^{T_d}(\pi,\tau)} \approx \frac{\tau'-\tau}{\tau} - \frac{\pi'-\pi}{\pi}.$$

*(ii) across durations. For the same individual across consecutive durations:*

$$\frac{R^{T_{d+1}}(\pi,\tau) - R^{T_d}(\pi,\tau)}{R^{T_d}(\pi,\tau)} \approx \theta^\tau - \theta^\pi,$$

*where the approximation is exact in* $\theta_\tau$ *and holds in a neighborhood of* $\theta_\pi = 1$*.*

*Proof.* See Appendix A.1.

For expositional simplicity, we have assumed that the utility loss from unemployment is constant over the spell and identical across individuals. Appendix A.2 shows that the result extends straightforwardly to the case of duration-dependent and heterogeneous unemployment losses.

To provide further intuition, consider the special case of a stationary environment with constant job-finding probability $\pi_d = \pi$ and constant utility loss $u_\delta = u$. Then the continuation loss takes

the simple form $L_{d+1}(\pi, \tau) = \frac{u}{\pi}$, so that the individual return to targeting becomes $R^{T_d}(\pi, \tau) = \frac{\tau_d}{c} \times \frac{u}{\pi}$. Returns are therefore given by the treatment effect relative to the cost times the avoided unemployment loss, which is larger for *at-risk* job seekers with low baseline job-finding probabilities.

Proposition 2 thus implies that optimal targeting prioritizes individuals with high predicted treatment effects *relative* to their baseline job-finding prospects. Moreover, returns depend crucially on how treatment effects depreciate over the spell *relative* to the decline in job finding. If treatment effects fall with duration while job finding remains stable ($\theta^{\pi} = 1$), returns are highest early in the spell. Conversely, if treatment effects are stable but job finding deteriorates with duration ($\theta^{\pi} < 1$), returns may increase later in the spell. While the latter may appear counterintuitive at first, it reflects that continuation losses are much larger late in the spell: early on, many individuals would exit unemployment even without treatment, which reduces the value of intervening at short durations.

# 3 Institutional Setting and Data

## 3.1 Institutional Setting

In Belgium, unemployed individuals must register with the regional Public Employment Service in order to receive unemployment benefits. The VDAB is the Public Employment Service for the Flemish Region, the Dutch-speaking part of Belgium and home to over 6.8 million inhabitants. Core aspects of the VDAB's mission include getting every working age resident into sustainable employment, providing mediation and training to all job seekers and eliminating the mismatch between supply and demand in the labor market [VDAB, 2024].

**Caseworkers.** The primary point of contact between the VDAB and a job seekers is the caseworker. The VDAB employs caseworkers at regional offices throughout the Flemish Region. Job seekers who have been assigned to a caseworker must meet with their caseworker regularly (possibly over the phone). Caseworkers provide job search counseling services to job seekers, monitor their job search activity and oversee their assignment to orientation and training programs. Job seekers who fail to show up for meetings or actively search for work face the threat of unemployment benefits sanctions.

**Assignment to caseworkers.** This section describes the process with which job seekers were assigned to caseworkers from October 2018 until May 2023.[1] Upon registration with the VDAB, job seekers receive an introduction email describing what they need to do in order to receive unemployment benefits. Job seekers who were unemployed in the last 6 months are automatically re-assigned to their previous caseworker, if they had one. For job seekers who have not yet been assigned to a caseworker, communication with VDAB occurs primarily through the 'service line.' The service line is the VDAB's call center. Mediators working at the service line provide advice to job seekers and determine which job seekers should be assigned to a caseworker at a regional office.

[1] Prior to October 2018, the VDAB did not use predicted risk to assist with assignment. From May 2023, each step of the caseworker assignment process was done three week earlier.

Job seekers are not required to call the service line in the first four weeks of unemployment. A portion of job seekers, however, call the service line of their own accord. Some of these callers are then assigned to a caseworker. Although uncommon, job seekers can, in theory, also show up to a regional VDAB office in person and be assigned to a caseworker at any time if capacity permits.

After four weeks of unemployment, job seekers aged 26-59 who have not yet been assigned to a caseworker receive an email instructing them to call the 'service line.' These calls involve an in-depth half-hour interview with a mediator who must then decide whether a job seeker is *self-reliant* or whether they should be assigned to a caseworker. Job seekers classified as 'self-reliant' are left to continue their search for work independently for the time being. Those assigned to a caseworker receive an appointment at a nearby regional office. Time until the first appointment depends on backlog, possibly taking up to 8 weeks at the busiest offices.

After five weeks of unemployment, the VDAB calls those job seekers who have not yet called the VDAB and who are considered to be at risk of long-term unemployment. In order to determine a job seeker's risk of long-term unemployment, the VDAB developed a machine learning prediction model that predicts job seekers' probability of employment within 6 months. Those with a probability below 65% are called and either assigned to a caseworker or determined to be self-reliant.[2] Job seekers who do not respond to calls from the VDAB face unemployment benefit sanctions. If after seven weeks of unemployment a job seeker with a score above 0.65 has not called the VDAB, they are classified as self-reliant.

Eleven weeks after being classified as self-reliant, job seekers who are still unemployed and who are still not assigned to a caseworker receive another email instructing them to call the service line. Those who do not call within the next week are called by the service line. Risk scores are not used for these follow-up interviews. Instead, all job seekers must speak to a mediator. Job seekers with a score above 0.65 who were not called and instead classified as self-reliant after seven weeks of unemployment are therefore called by the service line after 19 weeks of unemployment. In this second round of interviews, mediators again either assign job seekers to a caseworker or classify them as self-reliant. The process is repeated eleven weeks later in a third round of interviews, during which all remaining job seekers are assigned to a caseworker.

**The risk score.** The VDAB's prediction model, called *Kans op Werk* or "Employment Prospects," predicts the probability that a job seeker begins a period of employment lasting at least 28 consecutive days within the next 6 months. The random forest model uses more than 400 input variables, with a particular focus on labor market history and the current unemployment spell. Other variables include socio-demographic characteristics, job search preferences (e.g., occupation and geography) and job search behavior (e.g., interactions with the VDAB website). The model is retrained monthly and predictions are updated daily.

Prior to October 2020, risk scores were typically generated and sent to the service line once per week (usually on a Friday). In October 2020, however, the "INOS" system was introduced. Under this system, the risk scores used by the service line when deciding who to call were updated automatically on a daily basis.

[2]While in principle available to the mediators, risk scores are not used in their assignment decision.

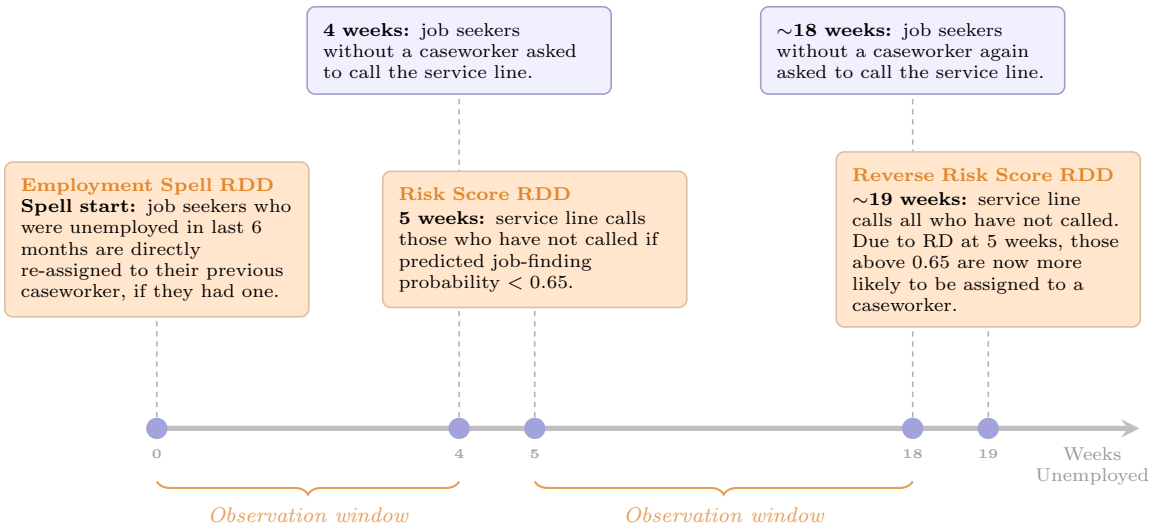


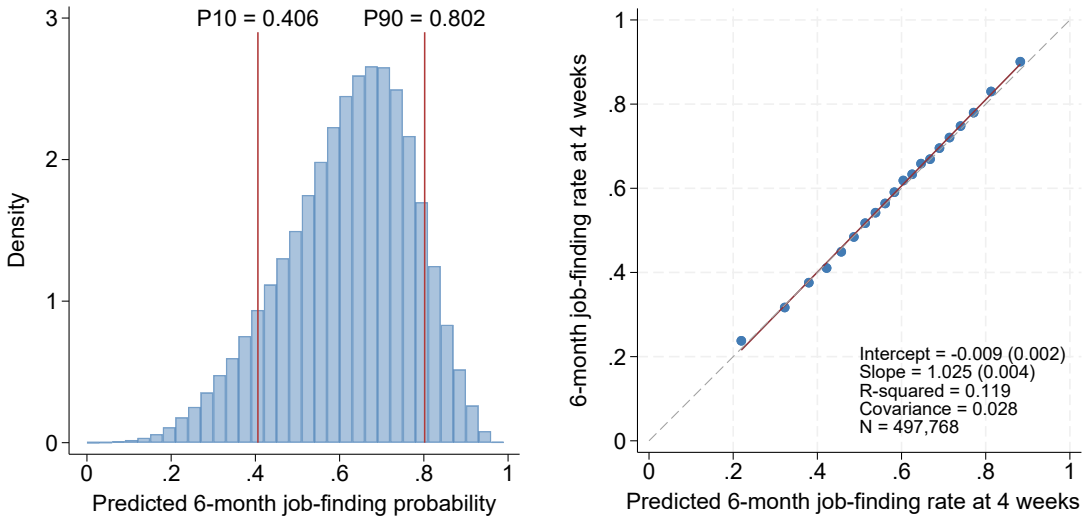

Figure 1: Assignment to Caseworkers

## 3.2 Data

We use administrative data from the Flemish Public Employment Service (VDAB) covering the universe of registered unemployment spells starting between 1 September 2018 and 30 November 2022 in Flanders, Belgium. The data includes a rich set of variables on individuals socio-demographic characteristics, labor market histories and outcomes, the VDAB's risk scores, interactions with the VDAB (e.g., contact with the service line, caseworkers, participation in programs, benefit sanctions), job search behavior (e.g., logins on the VDAB website, self-reported job applications), preferred occupations and occupational experience.

**Risk scores.** We use three datasets containing the risk scores predicted by the VDAB. The first contains all predicted 6-month job-finding probabilities predicted by the VDAB for job seekers in our sample on the 8th, 29th, 92nd, 183rd and 365th days of unemployment. Figure 2a plots the distribution of predicted 6-month job-finding rates on the 8th day, showing considerable dispersion. Figure 2b shows that the predictions made on the 29th day of unemployment are on average accurate. In Appendix B.1, we also show binned scatter plots for predictions made at the other four durations. Those made on the 8th day tend to be slightly less accurate than those made later in the spell.

The other two datasets we use contain the risk scores that were provided to the service line and used to profile job seekers. One contains predictions made from 9 November 2018 to 25 September 2020. These predictions are from before the INOS system was introduced. The other contains predictions from 3 October 2020 to 20 October 2022, which were handled by the INOS system. For the pre-INOS sample, we only have access to one prediction per job seeker, meaning that some observations are missing for job seekers who were called by the service line in two separate spells between November 2018 and September 2020.




(a) Distribution of Predictions

(b) Binned Scatter Plot: Predictions vs. Outcomes

Figure 2: The VDAB Risk Score

*Note:* Panel (a) shows a histogram of predicted 6-month job-finding probabilities generated on the 8th day of unemployment. Panel (b) then shows a binned scatter plot comparing the predicted 6-month job-finding rates generated on the 29th day of unemployment to the empirically observed 6-month job-finding rates for the same sample.

**Employment.** The VDAB defines employment as a period of at least 28 days in work. The VDAB also classifies individuals as employed if they are engaged in temporary work (also called interim work) or *individuele beroepsopleiding* (a subsidized work and training program lasting 1–6 months). The VDAB classifies an individual as engaged in temporary work if they have done at least 10 days of temporary work in the past 28 days. Temporary work involves fixed short-term contracts and is very common in Belgium. It is used by firms to replace workers on leave, handle temporary increases in workloads and trial new employees.

It is not clear whether, for the purposes of our analysis, temporary and subsidized work should be counted as employment, as many job seekers transition from such work back into unemployment. For our baseline analysis, we therefore count these forms of work as employment if they are directly followed by a transition into employment. Otherwise, they are counted as unemployment. We nevertheless check the robustness of our results to always counting temporary and subsidized work as employment as well as never counting them as employment.

**Labor market histories.** While we only have job seekers with unemployment spells beginning between 1 September 2018 and 30 November 2022, we have the employment history for all of these job seekers up until 31 July 2023, providing us with detailed information on their unemployment spells prior to 1 September 2018.

**Interactions with the VDAB.** We have detailed day-by-day data on job seekers contact with the service line and caseworkers, including both in-person contact and phone calls. After the INOS

system was introduced, so from October 2020, we have even more detailed data, including whether the job seeker or the service line made the call, whether a call was made to the job seeker but not answered and whether the service line classified a job seeker as self-reliant or assigned them to a caseworker.

We also have data on all programs that job seekers participate in. This includes, for example, job search orientation, training, job application requirements and external employment services. Programs vary greatly in terms of intensity, lasting from one day to many months.

Our data also includes information on sanctions investigations by the VDAB and the outcome of these investigations. Investigations are initiated for job seekers who do not show up to meetings or fulfill tasks. Outcomes can include partial or full suspensions of unemployment benefits.

# 4 Effects of Job Search Counseling

## 4.1 Empirical Strategy

To estimate the effects of job search counseling on job seekers, we first turn our attention towards the clean discontinuity at a risk score of 0.65 in the probability of being called by the VDAB service line and the resulting probability of being assigned to a caseworker. After five weeks of unemployment, job seekers with a predicted 6-month job-finding rate below 0.65 are considerably more likely to be interviewed by the service line and assigned to a caseworker.

We employ a regression discontinuity (RD) design, estimating the effect of having a score below 0.65 using the following baseline specification:

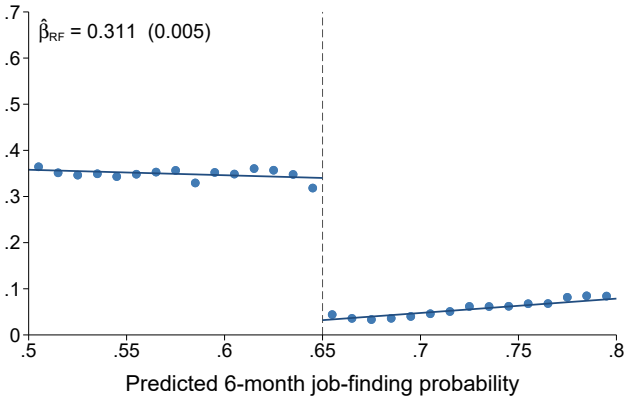
$$y_i = \alpha + \beta \mathbb{1}\{P_i < 0.65\} + \gamma_1(P_i - 0.65) + \gamma_2(P_i - 0.65)\mathbb{1}\{P_i < 0.65\} + \boldsymbol{X}_i'\boldsymbol{\delta} + \epsilon_i, \tag{4}$$

where $y_i$ is the outcome variable for unemployment spell $i$, $P_i$ is the predicted 6-month job-finding probability used by the service line when determining whether to contact job seekers and $\boldsymbol{X}_i$ is a vector of spell characteristics. We include observations in a bandwidth of 0.15 either side of the cut-off and apply a triangular kernel, thereby assigning more weight to observations closer to the cut-off. In Appendix C.2, we show that estimates are robust to applying a uniform kernel and a wide range of bandwidths.

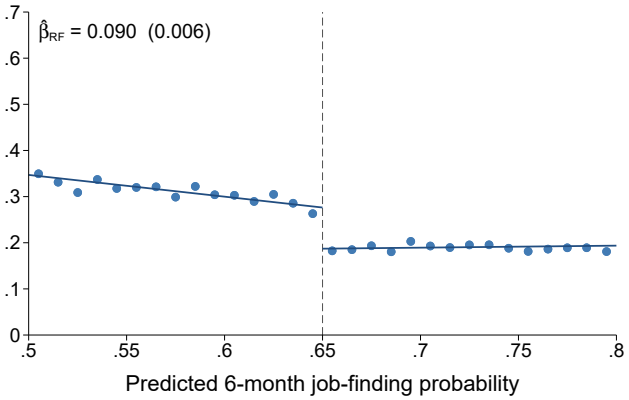
As risk scores are generated automatically by an algorithm, human manipulation around the cut-off is unlikely to be a concern. One may, however, worry that the random forest model, either randomly or due to patterns around the cut-off in the training data, assigns risk scores in a discontinuous manner around the cut-off. We show in Appendix C.1 that there is no bunching in risk scores around the cut-off. Appendix C.3 further shows that a number of predetermined covariates do not vary significantly at the cut-off.

## 4.2 Results

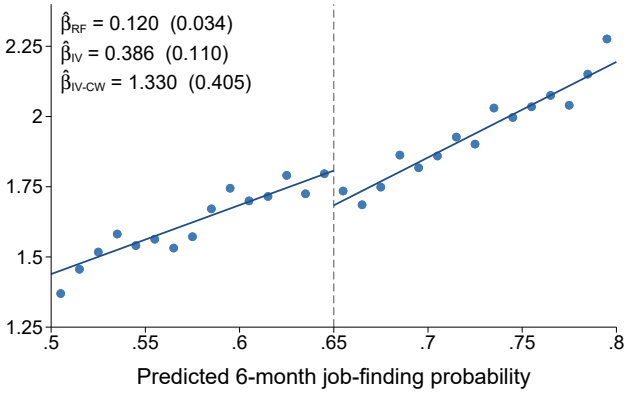
The service line begins calling those below the cut-off after five weeks of unemployment. As follow-up interviews begin three months later, we set a 12-week observation window covering weeks 6-17




(a) First Stage 1: Contact with the Service Line




(b) First Stage 2: Contact with a Caseworker




(c) Days Worked per Week

Figure 3: Risk Score Regression Discontinuity

*Note:* The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment, thereby excluding those who had already exited unemployment or been assigned to a caseworker. Panels (a) and (b) plot the proportion of spells in which, in weeks 6-17 of unemployment, the job seeker had contact with the service line and a caseworker, respectively. Panel (c) plots the mean days worked per week in weeks 6-17 controlling for quarter-year, gender, education, age, arrondissement, average previous unemployment duration, indicators for whether the last job or current unemployment spell began on the first of the month, and the day of the week on which the last job and the current unemployment spell began. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

of unemployment. We define two binary treatments: having had contact with the service line or contact with a caseworker during this period. For weeks 6-17 of unemployment, Figures 3a and 3b plot the proportions of job seekers around the cut-off who had contact with the service line and with a caseworker, respectively.

Our baseline outcome variable is the number of days worked per week throughout the 12-week observation window. Figure 3c plots days worked per week for job seekers around the cut-off. The reduced form interpretation of our estimate is that being below the cut-off, and thereby more likely to interact with the service line and be assigned to a caseworker, increases the number of days worked per week by 0.12.

We also report two fuzzy RD estimates, one using the discontinuity at the cut-off as an instrument for contact with the service line and another using it as an instrument for contact with a caseworker. The former, reported as $\hat{\beta}_{IV}$, can be interpreted as the combined effect of the service line interviewing job seekers and assigning a portion of them to caseworkers. The latter, reported as $\hat{\beta}_{IV-CW}$, is a valid estimate of caseworker effects if the service line interview by itself has no effect, an assumption that may not hold. As individuals can be considered randomly assigned around the cut-off and since being *above* the cut-off should not make anyone more likely to be assigned to a caseworker (local monotonicity), the other assumptions required for a fuzzy RD to be valid are satisfied [Hahn et al., 2001]. A second caveat in interpreting the effect of $\hat{\beta}_{IV-CW}$ is that, if the anticipation of an upcoming appointment with a caseworker affects job finding, the fact that we use contact with a caseworker as an indicator of treatment rather than assignment to a caseworker would lead to $\hat{\beta}_{IV-CW}$ being overestimated. We therefore report reduced form estimates as well as both IV estimates, with the stated caveats applying to any interpretation of $\hat{\beta}_{IV-CW}$.

While we find a positive effect on employment, it is important to consider whether job seekers compromise on the quality of employment in order to find employment sooner. In Appendix C.4, we find that, if anything, those just below the cut-off are likely to find *longer*-lasting employment than those just above the cut-off. One interpretation of this result is that the treatment not only increases the speed at which job seekers find employment, but also the stability of employment thereafter. The result could also, however, be driven by treatment effect heterogeneity, with job seekers who are ex-ante more likely to find stable employment also being more responsive to the treatment. This would cause the treatment group to have, on average, more stable employment conditional on job finding than the control group. In the following section, we explore the extent to which treatment effects are heterogeneous.

# 5 Heterogeneity in Effects

## 5.1 Heterogeneity by Predicted Risk

Given that predicted risk scores are used in the treatment assignment process, they are also one of the key dimensions on which we seek to test for treatment effect heterogeneity. The risk score RD employed in the previous section does not, however, allow us to estimate heterogeneity by risk score.

We therefore introduce another empirical strategy, leveraging the fact that individuals registering

for unemployment who were unemployed in the past six months are immediately reassigned to their previous caseworker if they had one in their previous spell. We exploit this discontinuity to estimate the heterogeneity in caseworker effects by predicted risk. We refer to this empirical strategy as the "employment spell RD."[3] We restrict the baseline observation window to four weeks, as after four weeks those above the cut-off are asked to call the service line and begin to receive the treatment.

In comparing results from the risk score RD and the employment spell RD, it is important to note that the samples are very different. The former involves individuals with above average predicted risk, who do not call the VDAB when asked to. The latter involves job seekers who were recently unemployed. The treatments themselves are also different, with those in the risk score RD being treated later in the spell, talking to service line mediators as well as caseworkers and likely meeting their caseworker for the first time rather than being reassigned to their previous caseworker.

To estimate heterogeneity, we split the sample into three groups based on their 6-month job-finding rate $P_i$ predicted on the 8th day of unemployment.[4] We first define a group that is most comparable with the risk score RD sample, those with $P_i \in [0.6, 0.7)$. We define the two other groups as those with $P_i \in [0, 0.6)$ and $P_i \in [0.7, 1]$, respectively.

We use the following baseline specification for estimating the employment spell RD:

$$y_i = \alpha + \beta \mathbb{1}\{d_i < 183\} + \gamma_1(d_i - 183) + \gamma_2(d_i - 183)\mathbb{1}\{d_i < 183\} + \boldsymbol{X}_i'\boldsymbol{\delta} + \epsilon_i, \tag{5}$$

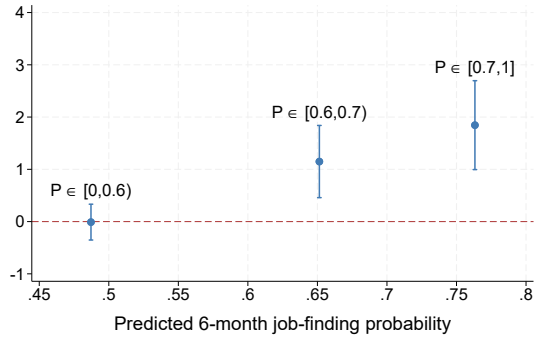
where $y_i$ is the outcome variable for unemployment spell $i$, $d_i$ is the duration in days between the beginning of the current unemployment spell and the end of the previous unemployment spell and $\boldsymbol{X}_i$ is a vector of spell characteristics. We include observations in a bandwidth of 20 weeks either side of the cut-off.

A threat to identification is the presence of bunching in the durations between unemployment spells. In Appendix D.1, we plot the distribution of durations, showing clear bunching at three, six, nine and twelve months. These appear to be reference points used by employers and employees. To avoid bias caused by selection into bunching, we estimate a 'donut' RD, excluding observations in the bunching region (182–194 days between unemployment spells). Results are also robust to including the excluded region. Due to the 'donut' hole, we apply a uniform kernel to baseline estimates, but results are robust to using a triangular kernel as well as a wide range of bandwidths, as shown in Appendix D.2.

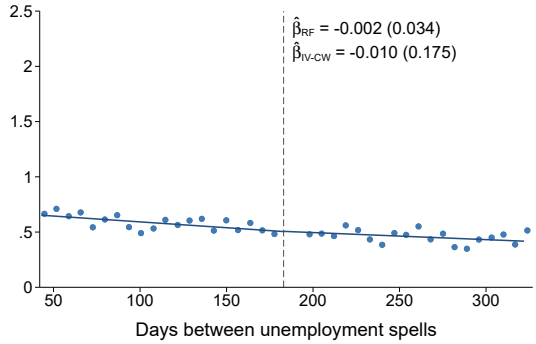
We may also be concerned that job seekers who were employed for less than six months differ on some dimension from those who were employed for more than six months. This could be the case if job seekers either side of the cut-off tend to work in different occupations or under different types of contracts. As a simple first check, we see that, outside the bunching region, the density of spells evolves smoothly around the cut-off. It would be concerning if there were, for example, many more durations below six months than above.

[3] Note, the running variable is not necessarily employment spell duration but rather 'duration not in unemployment.' For the majority of spells, however, the two are equal.

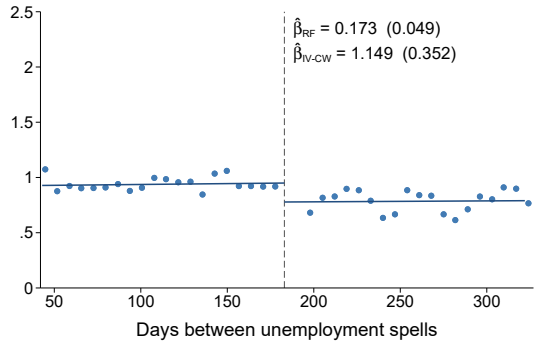
[4] Note, we use the prediction generated on the 8th day of unemployment as earlier predictions are considered less reliable. Predictions on the 8th day are also available for individuals who find work before reaching 8 days of unemployment. The predictions do not take into account that these individuals are already in work, thus not differing greatly for those who find work before and after 8 days.



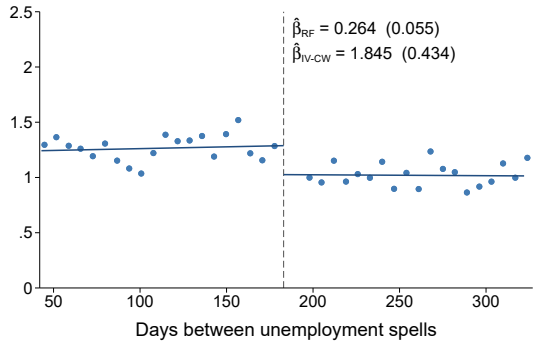

(a) Effect on Days Worked per Week (IV Estimates)




(b) Days Worked per Week ($P \in [0, 0.6)$)




(c) Days Worked per Week ($P \in [0.6, 0.7)$)




(d) Days Worked per Week ($P \in [0.7, 1]$)

Figure 4: Heterogeneity in Treatment Effects

*Note:* The sample is restricted to job seekers who had contact with a caseworker in their previous spell. Job seekers are then split into three groups by the probability of finding a job within 6 months $P$ that was predicted for them after 1 week of unemployment. Panel (a) plots, for each group, the average $P$ and the IV RD estimate of the effect of caseworkers on days employed per week in weeks 1–4 of unemployment. Panels (b) to (d) show the RD plots for each risk group, where the y-axis shows mean days worked per week plus the residual of days worked per week regressed on a set of controls. The controls are: quarter-year, gender, education, age, arrondissement, average previous unemployment duration, indicators for whether the last job or current unemployment spell began on the first of the month, and the day of the week on which the last job and the current unemployment spell began. We reject the null hypothesis that effects on the low and the high risk groups are equal with a p-value of 0.000. We report 95% confidence intervals. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

In Appendix D.3, we show the first stages of the employment spell RD for each risk group. We define the first stage as the proportion of job seekers in contact with a caseworker in the first four weeks of unemployment. We observe strong first stages for each risk group, with it being largest for those with $P_i \in [0, 0.6)$.

Figure A15 presents the estimated effects on employment by risk group. We estimate that the effect of caseworkers on days worked is a precise null for those with the lowest predicted job-finding probabilities—those *most* at-risk of long-term unemployment. Meanwhile, for job seekers in the other two groups, we estimate large treatment effects that are statistically significantly different from those estimated for the most at-risk population. While our baseline specification estimates effects over the first four weeks of unemployment, in Appendix D.4 we extend the analysis to a period of 16 weeks with similar results.

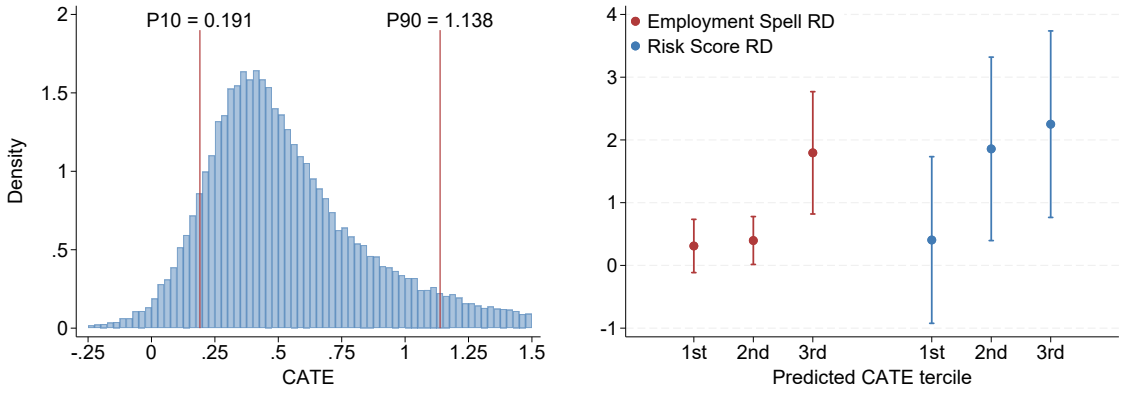
## 5.2 What Drives Heterogeneity by Predicted Risk?

In Appendix D.5, we explore whether the observed heterogeneity in treatment effects is due to heterogeneity in the type of treatment received. This would be the case if, for example, the most at-risk group were more likely to be assigned by their caseworker to long training programs—with lock-in effects thereby reducing effects on employment in the short-term. In the first four weeks of unemployment, we only find one task or program that those below the cut-off are more likely to engage in than those above the cut-off: the requirement to apply to a job. This result applies for each risk group. We therefore find it unlikely that lock-in effects are at play. There are also no large differences in task and program activity by risk group, suggesting that high-risk job seekers are indeed less responsive than low-risk job seekers to the same treatment.

Since the risk score is simply a function of other variables, we now consider which of those variables may drive the heterogeneity. We adopt a data-driven approach, applying an instrumental forest algorithm to estimate conditional average treatment effects (CATEs) and to identify the variables most important for explaining treatment effect heterogeneity [Athey et al., 2019]. An instrumental forest model will then also allow us to predict treatment effects for any unemployment spell.

We use the employment spell RD to estimate the instrumental forest model. One challenge we face with this approach is that we are unable to include a trend in the running variable. Luckily, the treatment status and our outcome do not exhibit strong trends in the running variable. We therefore proceed by reducing the bandwidth to 10 weeks on either side of the cut-off and estimating the instrumental forest using caseworker contact as the treatment assignment indicator and being below the cut-off as a binary instrument. Estimating the model without a trend in the running variable does, however, introduce some bias into our predicted CATEs. Comparing various RD specifications to CATE estimates, we find that differences between the two are primarily explained by the absence of trends in the running variable.

In Figure 5a, we plot the distribution of CATEs, showing that the algorithm predicts a large amount of variation in treatment effects. To assess the internal validity of the predicted CATEs, we estimate caseworker effects using the employment spell RD separately for each CATE tercile.




(a) Distribution of CATEs

(b) Internal and External Validity

Figure 5: Predicted Conditional Average Treatment Effects (CATEs)

*Note:* Panel (a) plots the distribution of estimated CATEs for the sample on which the instrumental forest was trained (note, each CATE is still an out-of-bag prediction). The red estimates in panel (b) are IV estimates of the effect of a caseworker on days worked per week by CATE tercile estimated using the employment spell RD. We reject the null hypothesis that effects on the 1st and 3rd terciles are equal with a p-value of 0.006. In blue is the estimated effect of a caseworker on days worked per week by CATE tercile estimated using the risk score RD. The p-value for the null hypothesis that the effects on the 1st and 3rd terciles are equal is 0.072. The tercile boundaries used are the same for both sets of estimates. We display 95% confidence intervals for the estimated effects.

The results are reported in Figure 5b. We find that the estimated treatment effect is significantly larger in the highest tercile, confirming that the CATE estimates capture meaningful treatment effect heterogeneity. We then also estimate treatment effects using the risk score RD for each CATE tercile and again find CATEs to be predictive of treatment effects. This latter finding is particularly noteworthy. The CATEs are estimated using the employment spell RD and therefore predict heterogeneity in the short-term effects of reassigning job seekers to their previous caseworker for a sample of individuals who were recently unemployed. The risk score RD, on the other hand, estimates effects for a slightly different treatment occurring later in the spell and over a longer window of observation for a sample that has spent less time unemployed, is more educated, older and more likely to speak Dutch well. It is not a given that the attributes that predict treatment effects for one context will also predict treatment effects in the other. Nevertheless, we find that CATEs estimated using the employment spell RD are also able to predict treatment effects for the risk score RD. We view these results as evidence supporting the external validity of findings from the employment spell RD.

The instrumental forest estimation also provides us with a useful measure for determining the most important variables for explaining treatment effect heterogeneity. This measure, "variable importance," is the number of times the forest split on a variable weighted by the depth at which the split occurred. We find unemployment history, risk score, age and occupation to be most important predictors of treatment effects.

A final potential cause of treatment effect heterogeneity by predicted risk comes from the fact that predicted risk is itself an outcome of treatment. Consider that the VDAB risk score is trained on a sample of individuals, most of whom either exit unemployment or are assigned to a caseworker within

six months. The algorithm should assign higher predicted job-finding scores to individuals with higher treatment effects. This alone could cause risk scores and treatment effects to be correlated.

Specifically, suppose

$$P_i = \pi_i + \tau_i T_i, \tag{6}$$

where $P_i$ is individual $i$'s probability of finding a job, $\pi_i$ is $i$'s baseline job-finding probability and $\tau_i$ and $T_i$ are treatment effect and treatment status, respectively. The focus of our analysis has been $Cov(P_i, \tau_i)$. Substituting for $P_i$ using Equation 6,

$$Cov(P_i, \tau_i) = Cov(\pi_i, \tau_i) + Cov(\tau_i T_i, \tau_i).$$

If treatment is randomized,

$$\begin{aligned} Cov(P_i, \tau_i) &= Cov(\pi_i, \tau_i) + Pr(T_i = 1) Var(\tau_i) \\ Cov(P_i, \tau_i) &> Cov(\pi_i, \tau_i) \end{aligned}$$

Hence, even if the job-finding probability in the absence of treatment $\pi_i$ is uncorrelated with the treatment effect $\tau_i$, the job-finding probability $P_i$ will be positively correlated with treatment effects if some random portion of the population is treated. Note that if the likelihood of treatment is declining in $\tau_i$, the correlation will be attenuated and could become negative.

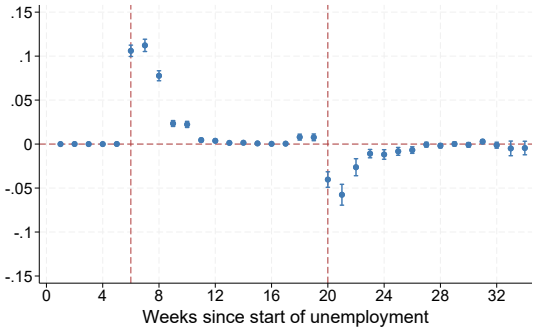
Targeting treatment on an outcome can therefore lead to adverse selection of those with low treatment effect into treatment. Ideally, the solution is to target on $\pi_i$ rather than $P_i$. However, this may be difficult if $\pi_i$ is difficult to predict. Possible solutions could involve training a model only on untreated individuals or using a function of predictions for $P_i$, $\tau_i$ and $T_i$ to predict $\pi_i$.

# 6 Duration Dependence in Effects

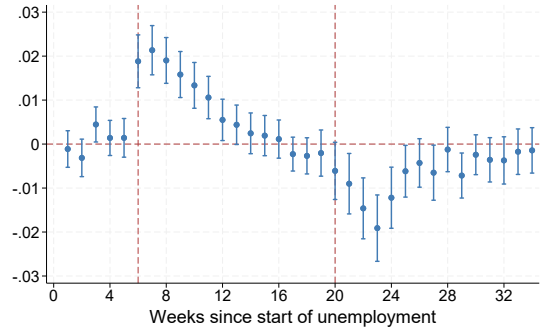
In this section, we investigate the dynamics of the assignment of job seekers to caseworkers. We first introduce another regression discontinuity design that allows us to shed light on the extent of true duration dependence in treatment effects.

In Section 4, we used the risk score RD to estimate caseworker effects for individuals with a predicted 6-month job-finding probability just below 0.65. This first round of treatment occurs after five weeks of unemployment. Since those below the cut-off are more likely to have been treated during the first round, those who were above the cut-off are more likely to be treated in the second round of treatment, when the cut-off is no longer in use. We refer to this empirical strategy as the "reverse risk score RD."

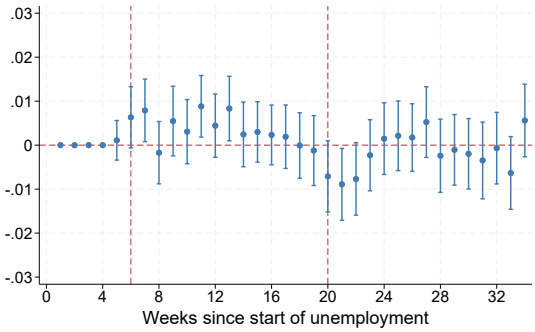
Figure 6a plots the RD estimate of being below the cut-off on contact with the service line at each week of unemployment, conditional on being unemployed at the beginning of the week. Those below the cut-off are much more likely to be in contact with the service line during the first round of treatment and much less likely during the second round of treatment. Figure 6b depicts a similar pattern for initial contact with a caseworker, though caseworker contact lags service line contact by up to a few weeks.



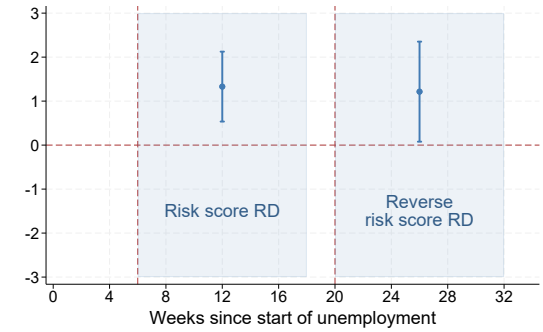

(a) Effect on Contact with Service Line (Weekly)




(b) Effect on First Contact with Caseworker (Weekly)



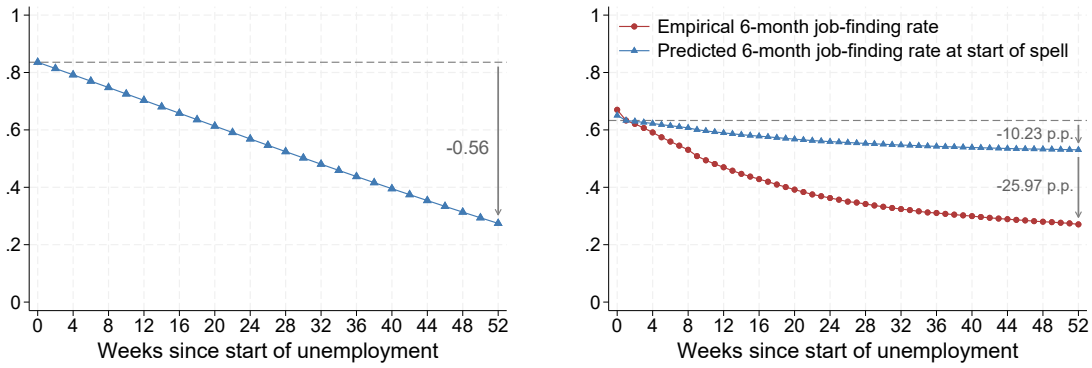

(c) Effect on Entry into Employment (Weekly)




(d) IV Effect on Days Worked per Week (Over 12 Weeks)

Figure 6: Dynamics

*Note:* All estimates use the baseline specification used in the risk score RD. Panel (a) plots RD estimates of the effect of being below $P = 0.65$ on contact with the service line each week, conditional on being unemployed at the beginning of that week. Panel (b) plots the estimated effect on having the first caseworker meeting of the unemployment spell each week. Panel (c) plots the estimated effect on employment each week. Panel (d) plots the IV estimate of the effect of meeting a caseworker for the first time in the 12-week period on days worked per week in the same period. We display 95% confidence intervals for the estimated effects.




(a) Simulated Effect of CW on Days Worked per Week

(b) Empirical vs. Initial Predicted Job Finding

Figure 7: Duration Dependence in Treatment Effects and Job Finding

*Note:* Panel (a) simulates the evolution of the average treatment effect at different unemployment durations using heterogeneity in treatment effects and predicted survival rates by predicted risk quintile. Treatment effects are estimated using the employment spell RD using baseline specifications. Survival rates are estimated using the average quintile risk score and are assumed to be constant over time. Panel (b) plots the empirical 6-month job-finding rate for the surviving population at each duration against 6-month job-finding rate predicted by VDAB on day 8 of unemployment for the surviving population at each duration. All predicted rates are multiplied by a constant so that the predicted and empirical rates are equal at a duration of 1 week.

Figure 6c plots the RD estimate of being below the cut-off on employment at the end of each week, conditional on being unemployed at the beginning of the week. We find that employment responds rapidly to the treatment, with effects concentrated in the short run. Figure 6d plots IV estimates of the effect of caseworkers on days worked per week over a 12-week period. Treatment is defined as having the first caseworker meeting of the spell during the 12-week period. In the first treatment round, we estimate a caseworker effect of 1.33 days worked per week, while in the second, we estimate an effect of 1.22.

The key threat to our identification strategy here is that the first round of treatment affects the survival probabilities at week 20 differently below and above the cut-off. As more individuals are treated below the cut-off, they are less likely to remain in unemployment until week 20. This is not a problem if treatment effects are uncorrelated with unobserved job-finding probability. However, if treatment effects are, for example, positively correlated with unobserved job-finding probability, then by week 20 we would expect the surviving population above the cut-off to have higher job-finding rates than those below the cut-off even in the absence of the second round of treatment. We argue, however, that the precise timing of our estimated employment effects and the zeroes estimated in weeks 18 and 19 shown in Figure 6c suggest that this form of selection is unlikely to drive our results.

Finally, we use the results of the risk score RD and the reverse risk score RD to assess duration dependence in caseworker treatment effects. Determining whether estimates at different durations represent true duration dependence or not can be a challenge due to dynamic selection of the sample. The key strength of our strategy is that both effects are local to job seekers who were assigned a predicted job-finding probability of around 0.65 after five weeks of unemployment. To the extent that the risk score is accurate, our design therefore compares the treatment effects of similar individuals at

different durations. While the confidence intervals are large, the results do not suggest the presence of true duration dependence in treatment effects.

Overall, however, a random sample of job seekers treated later in the spell are likely to have lower predicted job-finding probabilities than a random sample of job seekers treated earlier in the spell. Given the strong heterogeneity in treatment effects by predicted risk we identified in Section 5, we expect the selection of high-risk job seekers into long-term unemployment to cause the average treatment effect of unemployed job seekers to decrease with unemployment duration.

We simulate this in Figure 7a by splitting the sample into predicted risk quintiles and simulating the average treatment effect of the surviving population at different unemployment durations, assuming treatment effects and survival rates are the same within quintile and heterogeneous across quintiles. Quintile treatment effects are estimated using the employment spell RD. Quintile survival rates are estimated using the average predicted risk score within each quintile. The simulation predicts that dynamic selection causes the average treatment effect to fall from 0.84 for job seekers just entering unemployment to 0.27 for those who have been unemployed for a year.

In Figure 7b, we then investigate duration dependence in job-finding rates. Empirically, we find that 6-month job-finding rates fall from 0.63 at one week of unemployment to 0.27 at one year of unemployment, a reduction of 57%. Following Mueller and Spinnewijn [2025], the decline in the job-finding rates that were predicted at the start of the spell can give us a lower bound estimate of the proportion of duration dependence in job-finding explained by dynamic selection rather than true duration dependence. While the decline in job-finding caused by dynamic selection is a bit smaller than that documented by Mueller and Spinnewijn [2025] in Sweden, the overall decline in job-finding is much larger (in Sweden the decline is only 33% over a year). This suggests there may be considerable true duration dependence in job-finding among job seekers in Flanders.

# 7 Quantitative Evaluation

In this section, we quantify the welfare gains from different targeting policies in the context of our empirical analysis in Flanders. Our quantitative evaluation deviates from the approximation above and allows for duration dependence in job finding, both at the individual level and through unobserved heterogeneity, and carefully parametrizes the model to moments in the data. Therefore, it also allows us to evaluate how well the approximation works when we significantly deviate from the stationary world with no depreciation.

**Calibration.** We calibrate the model to key moments in the data, as summarized in Table 1. We begin by using the empirical distribution of job-finding risk scores (RS) from the VDAB.[5] Because the model is calibrated at the monthly frequency, we map each risk score—defined as the predicted 6-month job-finding probability over the first six months of an unemployment spell—into an underlying type, $\pi_i$. [6] Unobserved heterogeneity in job finding, $\sigma_{\varepsilon^\pi}$, is calibrated using evidence from Mueller

[5] We draw a random sample of 50,000 individuals from the baseline population of job seekers used in the empirical analysis.

[6] In the absence of unobserved heterogeneity, this mapping is given by $\pi_i = 1 - (1 - RS_i)^{1/6}$. With unobserved heterogeneity, no closed-form solution exists, and we instead recover $\pi_i$ using a numerical solver.

and Spinnewijn [2025], targeting the relative contribution of observed and unobserved heterogeneity to job-finding risk. We assume geometric depreciation, $\theta_d^\pi = (\theta^\pi)^d$, and choose $\theta^\pi$ to match the decline in job-finding rates over the first 12 months in the Flemish data.

Table 1: Calibrated Parameter Values in Baseline

| Object | Description | Value | Source/Targeted Moments |
|---|---|---|---|
| $G(\pi)$ | Distribution of observable job finding | — | Empirical distribution of risk scores |
| $\sigma_{\varepsilon^\pi}$ | Unobserved heterogeneity in job finding | 0.037 | Mueller and Spinnewijn [2025] |
| $\alpha_\tau$ | Constant in $\tau = \alpha_\tau + \beta_\tau \pi + e_\tau$ | -0.045 | RDD estimates by risk score |
| $\beta_\tau$ | Slope in $\tau = \alpha_\tau + \beta_\tau \pi + e_\tau$ | 0.104 | RDD estimates by risk score |
| $\sigma_{e^\tau}$ | Orthogonal heterogeneity in treatment effects | 0.022 | Variance of CATEs |
| $\theta^\pi$ | Duration dependence in job finding | 0.918 | 12-month decline in job finding |
| $\theta^\tau$ | Duration dependence in treatment effects | 1.000 | Estimates of risk-score RDDs |
| $\bar{B}$ | Budget constraint | 0.530 | Share of unemployed targeted by VDAB |

For treatment effects, we assume that observable heterogeneity relates linearly to risk scores according to

$$\tau = \alpha_\tau + \beta_\tau RS + e_\tau,$$

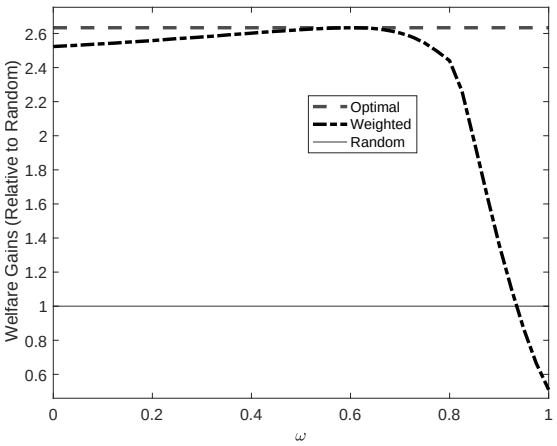
where $\alpha_\tau$ is a constant, $\beta_\tau$ captures the correlation between treatment effects and risk scores, and $e_\tau$ represents heterogeneity orthogonal to risk scores. We calibrate $\alpha_\tau$ and $\beta_\tau$ by estimating an RDD across deciles of the risk score distribution and fitting a linear approximation to these estimates (see Figure A2 in the Appendix). To discipline the orthogonal component, we target the variance of CATEs shown in Figure 5a.[7] In the baseline, we assume no additional unobserved heterogeneity in treatment effects, $\sigma_{\varepsilon^\tau} = 0$, but assess robustness to this assumption. We also impose no duration dependence in treatment effects, consistent with the RDD estimates at 5 and 19 weeks.

We lack a direct empirical target for the utility cost of unemployment and therefore set $\theta_d^u = (\theta^u)^d = 1$ in the baseline, while exploring sensitivity to this assumption. Because our analysis focuses on relative welfare comparisons, the treatment cost $c$ and the initial utility loss $u$ cancel out; we therefore normalize both to one and omit these parameters from the table. Finally, we impose a budget constraint consistent with current VDAB policy, allowing treatment of 53 percent of the sample. We further assume that distributions of $\varepsilon^\pi$ and $\varepsilon^\tau$ are uniformly distributed with mean zero. Further note that we assume that the environment becomes stationary after 12 months, i.e. there is no further duration dependence in job finding. Appendix Figure A1 shows the distribution of $\pi$s and $\tau$s as well as the duration dependence in job finding, which matches closely the empirical observed duration dependence.

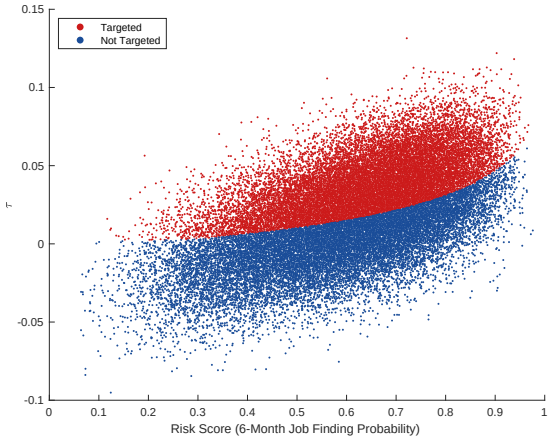
**Evaluation of different targeting rules.** We consider several different targeting rules for newly unemployed workers at duration $d = 0$:

1. **Optimal Targeting ($T_0^{OPT}$)**: The constrained optimal targeting policy, which targets individuals based on the return to targeting, $R^{T_0}(\pi, \tau)$.

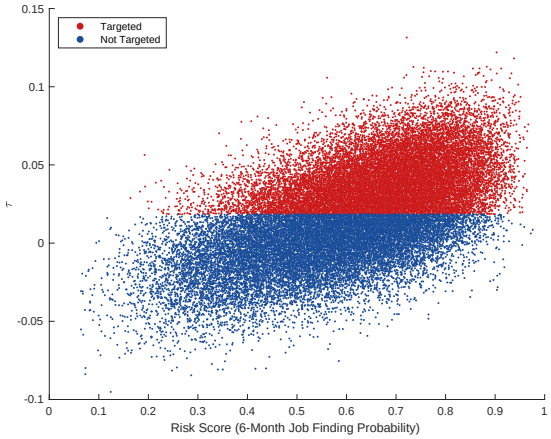
[7] More precisely, the CATEs reported in Figure 5a are based on days worked. In the calibration, however, we target the dispersion of CATEs for the probability of starting work within four weeks, which more closely aligns with the model's monthly job-finding probability.



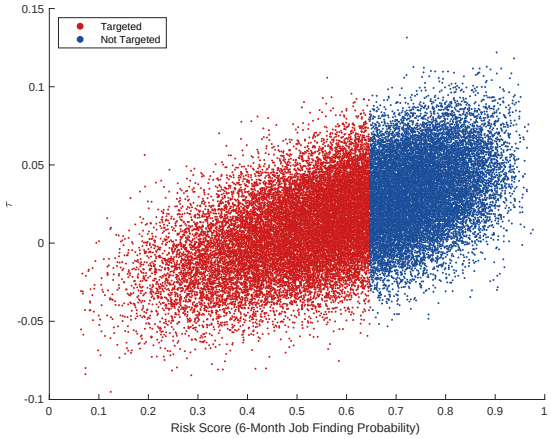

(a) Welfare Gains Relative to Random Targeting




(b) Targeted Types With Optimal Policy




(c) Targeted Types With $\omega = 0$




(d) Targeted Types With $\omega = 1$

Figure 8: Welfare Gains of Targeting and Targeted Types under Alternative and Optimal Policies

*Note:* Panel (a) shows welfare gains from targeting relative to the benchmark of random assignment. Panel (b) displays the worker types targeted under the optimal policy. Panels (c) and (d) illustrate the targeted types under sub-optimal policies that focus exclusively on treatment effects ($\omega = 0$) or exclusively on job-finding probabilities ($\omega = 1$), respectively. All policies are subject to the same budget constraint.

2. **Weighted Targeting ($T_0^{WGT}$)**: Targeting based on a geometric weighted average of $\pi$ and $\tau$

$$Q(\pi, \tau, \omega) \quad = \quad \left(\frac{1}{\pi}\right)^{\omega} (\tau)^{1-\omega}, \tag{7}$$

where $Q(\pi, \tau, 1) = \frac{1}{\pi}$ targets based on $\pi$ only and $Q(\pi, \tau, 0) = \tau$ targets based on $\tau$ only.

3. **Random Targeting ($T_0^{AVG}$)**: Targeting the average unemployed worker, i.e. a random sub-sample of unemployed workers.

Panel (a) of Figure 8 reports the returns to targeting under each policy, relative to a benchmark that randomly treats the average job seeker. Several results stand out. First, the welfare gain under the optimal policy is nearly three times larger than under the benchmark. This reflects substantial heterogeneity in both observed risk scores and treatment effects across individuals. Panel (b) illustrates the types targeted under the optimal policy and highlights the underlying trade-off.

Second, the optimally weighted targeting rule attains its maximum at $\omega = 0.60$, although performance is very similar at $\omega = 0.50$, indicating that a simple rule based on the index $\tau/\pi$ provides an excellent approximation.[8] Consistent with this, Table 2 shows that the optimal weight is exactly $\omega = 0.50$ in a stationary environment without duration dependence in job finding.

Importantly, a policy that targets exclusively on job-finding probabilities ($\omega = 1$) performs *worse* than the benchmark of random targeting. This is because individuals at highest risk of long-term unemployment exhibit negligible treatment effects. Conversely, a policy that targets exclusively on treatment effects ($\omega = 0$) performs nearly as well as the optimal policy, as it avoids selecting individuals with low returns to treatment. Panels (c) and (d) display the corresponding targeted types for $\omega = 0$ and $\omega = 1$, respectively. Taken together, these results show that targeting based solely on risk scores falls well short of the optimum, as it fails to balance long-term unemployment risk against treatment effectiveness.

Figure 9 reports the welfare gains for the optimal targeting rule, relative to the benchmark of random targeting at duration $d = 0$, separately by unemployment duration. The gains decline strongly over the first few months of the unemployment spell up to month 4, where fewer than half of the initial unemployed population remains without a job. At that point, the optimal targeting rule effectively assigns treatment to all individuals who remain unemployed. The initial decline is driven by the fact that the optimal policy can be very selective at the beginning of the spell where all unemployed workers are available for treatment. As the unemployment spell progresses, fewer unemployed are left to treat so the targeted policy becomes closer to random policy that treats the average surviving unemployed worker.

While the welfare gains of *optimal* targeting relative to the average benchmark declines with duration, the welfare gains of *random* targeting rise over the spell, as shown in Figure 9. This pattern is driven by the steep decline in job-finding rates observed in the Belgian data, which makes interventions later in the unemployment spell more valuable.[9] Of course, this result is influenced by our calibration assumption that treatment efficacy does not depreciate over the unemployment spell

[8] As noted above, $\pi$ denotes the 6-month risk score converted into a 1-month probability.
[9] See also Appendix Figure A3 for how the welfare gains increase for a given type.

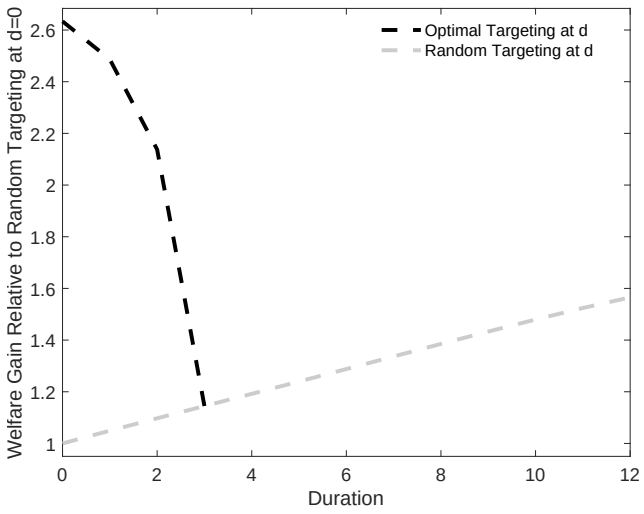



Figure 9: Welfare Gains of Targeting Policies at $d$, Relative to Random Targeting at $d = 0$

*Note:* The figure reports the welfare gains from applying the optimal targeting rule at unemployment duration $d$, relative to random targeting at $d = 0$. Welfare gains are greyed out after month 4, when the optimal policy can no longer exhaust the full budget because too few workers remain unemployed.

at the individual level. We further investigate the implications of this assumption in the sensitivity analysis below.

**Sensitivity analysis.** Table 2 presents a sensitivity analysis with respect to the main calibration parameters. It reports the returns to targeting, together with the optimal weight in the weighted targeting rule in equation 7. Across all specifications, the gains from optimal targeting remain large, indicating that the high return to targeting is robust to alternative calibrations. These gains become even more pronounced when the budget is constrained such that only the top 30 percent of individuals can be treated.

The optimal weight is also remarkably stable across calibrations, and is exactly equal to 0.5 in column (3), where job finding exhibits no duration dependence. This confirms our analytical result that, in the neighborhood of no depreciation in job-finding rates, the optimal targeting index is proportional to $\tau/\pi$. Moreover, policies that target solely based on low risk scores or solely on high treatment effects are clearly suboptimal, but nonetheless deliver substantial welfare improvements relative to the benchmark of targeting the average unemployed worker.

The result that the welfare gains are decreasing with duration under the optimal policy is also robust across different calibration choices. As Table 2 shows, with targeting only 30% of the sample, the welfare gains of the optimal policy are decline much less over the first few months of the spell. We verified, however, that the welfare gains of waiting longer than 3 months are more strongly decreasing in this scenario. The welfare gains of random targeting is increasing over the spell is sensitive to calibration choices and is stronger with positive duration dependence in the utility loss from unemployment (column 4), but is decreasing under no duration dependence in job finding

Table 2: Welfare Gains of Targeting, Across Alternative Calibration Scenarios

| | (1) Baseline | (2) Target 30% | (3) $\sigma_{\varepsilon^\pi}=0$, $\theta^\pi=1$ | (4) $\theta^u>1$ | (5) $\theta^\tau<1$ | (6) $corr(\pi,\tau)=0$ | (7) $\tau$ $(IV)$ | (8) $\sigma_{\varepsilon^\tau}>0$ | (9) At $d=25$: no continuation loss |
|---|---|---|---|---|---|---|---|---|---|
| Optimal Targeting vs. Random | 2.88 | 3.81 | 2.66 | 3.01 | 2.88 | 2.09 | 2.01 | 2.63 | 2.42 |
| Optimally Weighted Targeting vs. Random | 2.88 | 3.81 | 2.66 | 3.01 | 2.88 | 2.09 | 2.01 | 2.61 | 2.42 |
| Optimal Weight | 0.60 | 0.60 | 0.50 | 0.63 | 0.60 | 0.60 | 0.60 | 0.62 | 0.51 |
| Targeting High Treatment Effects vs. Random | 2.78 | 3.31 | 2.62 | 2.89 | 2.78 | 2.03 | 1.92 | 2.52 | 2.39 |
| Targeting Low Risk Scores vs. Random | 0.56 | -0.07 | 0.45 | 0.57 | 0.56 | 1.45 | 0.85 | 0.68 | 0.56 |
| Optimal at $d=3$ vs. Optimal at $d=0$ | 0.41 | 0.95 | 0.58 | 0.43 | 0.31 | 0.70 | 0.63 | 0.48 | 0.44 |
| Random at $d=3$ vs. Random at $d=0$ | 1.18 | 1.18 | 0.83 | 1.29 | 0.89 | 1.46 | 1.27 | 1.26 | 1.06 |

*Notes:* This table reports the welfare gains of targeting policies under alternative calibration scenarios, expressed relative to welfare gains of a benchmark policy. Numbers in parentheses refer to optimal weight in the weighted targeting policy. Column (1) shows the baseline calibration. Column (2) restricts the budget so that only the top 30 percent of individuals can be treated. Column (3) combines $\sigma_{\varepsilon^\pi}=0$ with $\theta^\pi=1$. Column (4) assumes that the utility cost of unemployment increases by 50 percent at month 12. Column (5) assumes treatment effects decline at the individual at the same rate as job finding. Column (6) imposes zero correlation between observed types, $corr(\pi,\tau)=0$. Column (7) uses instrumental-variables estimates of treatment effects, $\tau$ (IV), instead of reduced-form estimates. Column (8) allows for unobserved heterogeneity in treatment effects by setting $\sigma_{\varepsilon^\tau}=0.79\sigma_\tau$, mirroring the relative importance of unobserved and observed heterogeneity for job finding probabilities. Column (9) assumes no continuation loss after period 25 in the model.

(column 3) and under negative duration dependence in treatment effects (column 5), and less strongly increasing when the utility loss is assumed to drop to zero after two years (column 9). While the latter assumption is counterfactual, it can be interpreted as an environment in which the policymaker cares only about UI benefit payments, which expire after a fixed duration.

An important robustness check considers a scenario in which the correlation between risk scores and treatment effects is set to zero (i.e., $\beta_\tau=0$), while re-calibrating $\alpha_\tau$ to match the average treatment effect (see column (6)). In this case, the welfare gains under the optimal policy are somewhat reduced, but remain sizable. Importantly, targeting based solely on risk scores ($\omega=1$) performs substantially better in this environment and clearly outperforms the benchmark of random targeting. The remaining robustness checks in columns (7) and (8) have little effect on the results, including using IV-based rather than reduced-form treatment effects and allowing for unobserved heterogeneity in treatment effects.[10]

# 8 Conclusion

This paper has examined the optimal targeting of job search counseling in the context of the Flemish Public Employment Service. We have shown that the common practice of directing active labor market interventions toward the most at-risk job seekers (those with the lowest predicted probability of finding employment) is unlikely to be welfare-maximizing. The individuals most at risk are not the same as the individuals most responsive to intervention, and conflating the two leads to a substantial misallocation of scarce program resources.

Our three regression discontinuity designs, exploiting discontinuities in the VDAB's caseworker assignment process at the start of an unemployment spell, at five weeks, and at nineteen weeks, yield two key empirical findings. First, treatment effects are large and positive for low-risk job seekers but

[10]These alternative assumptions do not materially affect the qualitative ranking of targeting policies.

effectively zero for those classified as most at risk. Second, there is no evidence of individual-level depreciation in treatment effects over the unemployment spell, but strong dynamic selection causes the pool of surviving job seekers to be progressively less responsive even as it becomes more at-risk.

These empirical findings motivate our conceptual framework, which characterizes the optimal targeting problem of an unemployment agency that minimizes expected welfare losses subject to a fixed budget. The key insight is that the return to treating a given job seeker depends on both the magnitude of their treatment effect and the continuation loss from unemployment they would otherwise face, with the latter being larger for higher-risk individuals. The optimal targeting index thus involves the ratio of predicted treatment effects to baseline job-finding probabilities, placing positive weight on both responsiveness and risk.

The quantitative evaluation demonstrates that incorporating treatment effect heterogeneity into the targeting rule could nearly triple the welfare gains relative to random assignment. Even a simple rule based on treatment effects alone substantially outperforms the status quo. By contrast, targeting exclusively on risk scores can perform worse than random targeting, because risk scores and treatment effects are strongly negatively correlated: those with the lowest predicted job-finding probability are precisely the individuals with the smallest returns to intervention, so that risk-score targeting systematically directs resources away from those who would benefit most. The framework also sheds light on the optimal timing of intervention. Welfare gains from optimal targeting decline strongly over the first few months of unemployment, as dynamic selection progressively erodes the heterogeneity in the surviving pool and the targeted policy converges toward treating the average worker. Employment agencies should therefore invest in identifying responsive individuals early in the spell, when the scope for welfare-improving targeting is greatest.

More broadly, our results highlight a general principle that extends beyond unemployment policy: in settings where individual responsiveness to treatment is heterogeneous and negatively correlated with baseline need, targeting based on risk alone can be counterproductive. The degree to which this principle applies in other active labor market programs, particularly training programs which tend to exhibit the opposite heterogeneity pattern, remains an important question for future research.

# References


**Ahn, Hie Joo**, "The Role of Observed and Unobserved Heterogeneity in the Duration of Unemployment," *Journal of Applied Econometrics*, 2023, *38* (1), 3–23.

_ **and James D. Hamilton**, "Heterogeneity and Unemployment Dynamics," *Journal of Business & Economic Statistics*, 2020, *38* (3), 554–569.

**Altmann, Steffen, Anita Marie Glenny, Robert Mahlstedt, and Alexander Sebald**, "The Direct and Indirect Effects of Online Job Search Advice," Discussion Paper 15830, Institute of Labor Economics (IZA) 2022.

_ **, Armin Falk, Simon Jäger, and Florian Zimmermann**, "Learning about Job Search: A Field Experiment with Job Seekers in Germany," *Journal of Public Economics*, 2018, *164*, 33–49.

**Alvarez, Fernando, Katarína Borovičková, and Robert Shimer**, "Decomposing Duration Dependence in a Stopping Time Model," *Review of Economic Studies*, November 2024, *91* (6), 3151–3189.

**Athey, Susan, Julie Tibshirani, and Stefan Wager**, "Generalized Random Forests," *The Annals of Statistics*, 2019, *47* (2), pp. 1148–1178.

**_ , Niall Keleher, and Jann Spiess**, "Machine Learning Who to Nudge: Causal vs. Predictive Targeting in a Field Experiment on Student Financial Aid Renewal," *Journal of Econometrics*, 2025, *249*, 105945.

**Belot, Michèle, Philipp Kircher, and Paul Muller**, "Do the Long-Term Unemployed Benefit from Automated Occupational Advice during Online Job Search?," *The Economic Journal*, January 2026, *136* (673), 184–206.

**Card, David, Jochen Kluve, and Andrea Weber**, "What Works? A Meta-Analysis of Recent Active Labor Market Program Evaluations," *Journal of the European Economic Association*, 2018, *16* (3), 894–931.

**Desiere, Sam, Katrin Langenbucher, and Ludo Struyven**, "Statistical Profiling in Public Employment Services: An International Comparison," Technical Report 224, OECD Social, Employment and Migration Working Papers, Paris 2019.

**Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw**, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 2001, *69* (1), 201–209.

**Haushofer, Johannes, Paul Niehaus, Carlos Paramo, Edward Miguel, and Michael Walker**, "Targeting Impact versus Deprivation," *American Economic Review*, 2025, *115* (6), 1936–1974.

**Kitagawa, Toru and Aleksey Tetenov**, "Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice," *Econometrica*, 2018, *86* (2), 591–616.

**Koning, Pierre**, "The Effectiveness of Public Employment Service Workers in the Netherlands," *Empirical Economics*, October 2009, *37*, 393–409.

**Liu, Songqi, Jing L. Huang, and Mo Wang**, "Effectiveness of Job Search Interventions: A Meta-Analytic Review," *Psychological Bulletin*, 2014, *140* (4), 1009–1041.

**Mueller, Andreas I. and Johannes Spinnewijn**, "The Nature of Long-Term Unemployment: Predictability, Heterogeneity, and Selection," *Journal of Political Economy*, 2025, *133* (12), 3846–3902.

**_ , _ , and Giorgio Topa**, "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias," *American Economic Review*, January 2021, *111* (1), 324–363.

**VDAB**, "Jaarverslag 2024," 2024.

**Yadlowsky, Steve, Scott Fleming, Neil Shah, Emma Brunskill, and Stefan Wager**, "Evaluating Treatment Prioritization Rules via Rank-Weighted Average Treatment Effects," *Journal of the American Statistical Association*, 2025, *120* (549), 38–51.

# Appendix

## A Conceptual Framework

### A.1 Proofs of Propositions

**Proof of proposition 1.**

*Proof.* Since targeting is restricted to duration $d$, the treatment indicator simplifies to $T_d(\pi, \tau) \in \{0, 1\}$, with $T_\delta(\pi, \tau) = 0$ for all $\delta \neq d$. The agency's problem is therefore to choose a treatment rule $T_d(\pi, \tau)$ to minimize expected utility losses subject to the budget constraint:

$$\min_{T_d} \int \int L^{T_d}(\pi, \tau)\, g(\pi, \tau)\, d\pi\, d\tau \quad \text{s.t.} \quad \bar{B} \geq \int \int c\, T_d(\pi, \tau)\, S_d(\pi)\, g(\pi, \tau)\, d\pi\, d\tau,$$

where we use $S_d(\pi)$ since survival to duration $d$ is evaluated under no treatment for $\delta < d$.

*Lagrangian.* Attaching multiplier $\lambda \geq 0$ to the budget constraint, the Lagrangian is

$$\mathcal{L} = \int \int \Big[ L^{T_d}(\pi, \tau) + \lambda\, c\, T_d(\pi, \tau)\, S_d(\pi) \Big] g(\pi, \tau)\, d\pi\, d\tau.$$

Since the Lagrangian is an integral of separable terms, minimizing over the treatment rule $T_d(\pi, \tau)$ reduces to minimizing the integrand independently for each type $(\pi, \tau)$:

$$\min_{T_d(\pi,\tau) \in \{0,1\}} \; L^{T_d}(\pi, \tau) + \lambda\, c\, T_d(\pi, \tau)\, S_d(\pi) \quad \text{for each } (\pi, \tau).$$

*Pointwise optimality.* Define the net benefit of treating type $(\pi, \tau)$ as

$$\Delta(\pi, \tau) \;\equiv\; \big[ L(\pi) - L^{T_d}(\pi, \tau) \big] - \lambda\, c\, S_d(\pi).$$

It is optimal to set $T_d(\pi, \tau) = 1$ if and only if $\Delta(\pi, \tau) \geq 0$, i.e.

$$\frac{L(\pi) - L^{T_d}(\pi, \tau)}{c\, S_d(\pi)} \;\geq\; \lambda.$$

Recognising that the left-hand side equals $R^{T_d}(\pi, \tau)$ from equation (3), this condition reads

$$T_d(\pi, \tau) = 1 \iff R^{T_d}(\pi, \tau) \geq \lambda.$$

*Threshold and budget clearing.* Setting $\bar{R}^{T_d} = \lambda$, the budget constraint holds with equality at the optimum (since it is always beneficial to exhaust the budget), yielding

$$\bar{B} = \int \int c\, \mathbb{I}\big[ R^{T_d}(\pi, \tau) \geq \bar{R}^{T_d} \big]\, S_d(\pi)\, g(\pi, \tau)\, d\pi\, d\tau,$$

which determines $\bar{R}^{T_d}$ implicitly. Continuity of $g(\pi, \tau)$ ensures that the set of types exactly at the threshold has measure zero, so the $\geq$ and $>$ cases are equivalent and the threshold is well-defined. $\square$

**Proof of proposition 2.**

*Proof.* Since utility losses are constant and identical across job seekers, $u_d = u$, the continuation loss from unemployment from period $d+1$ onward is

$$L_{d+1}(\pi, \tau) = u \sum_{\delta=d+1}^{\infty} S_{\delta|d+1}(\pi),$$

where $S_{d+1|d+1}(\pi, \tau) = 1$ and for $\delta > d+1$, $S_{\delta|d+1}(\pi) = \prod_{j=1}^{\delta-d-1}(1-\pi_{d+j})$ is the probability of remaining unemployed from $d+1$ to $\delta$. This equals $u$ times the expected remaining duration from $d+1$ onward. At $\theta_\pi = 1$, all $\pi_{d+j} = \pi$ and $S_{\delta|d+1} = (1-\pi)^{\delta-d-1}$, so

$$L_{d+1} = u \sum_{k=0}^{\infty} (1-\pi)^k = \frac{u}{\pi}.$$

The one-period return to targeting is

$$R^{T_d}(\pi, \tau) \;=\; \frac{\tau_d}{c} \cdot L_{d+1}(\pi, \tau).$$

**Part (ii).** We first establish how $L_{d+1}$ varies with $\theta_\pi$ in a neighborhood of $\theta_\pi = 1$, as Part (i) follows as a corollary. At $\theta_\pi = 1$, $L_{d+1} = u/\pi$ for all $d$, so we expand:

$$L_{d+1} \;\approx\; \frac{u}{\pi} + \left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=1} (\theta_\pi - 1).$$

*Derivative of $L_{d+1}$ with respect to $\theta_\pi$.* Setting $k = \delta - d - 1$ and using $\pi_{d+j} = \theta_\pi^{d+j}\pi$, we write

$$L_{d+1} = u + u \sum_{k=1}^{\infty} \prod_{j=1}^{k} (1 - \theta_\pi^{d+j}\pi),$$

where the $k = 0$ term ($\delta = d+1$) equals $u \cdot 1$ and is independent of $\theta_\pi$, contributing zero to the derivative. For $k \geq 1$, by the product rule, differentiating the $j$-th factor gives $-(d+j)\theta_\pi^{d+j-1}\pi$, while the remaining $k-1$ factors each equal $(1-\pi)$ at $\theta_\pi = 1$. Therefore:

$$\left.\frac{\partial}{\partial \theta_\pi} \prod_{j=1}^{k} (1 - \theta_\pi^{d+j}\pi)\right|_{\theta_\pi=1} = -(1-\pi)^{k-1} \sum_{j=1}^{k} (d+j)\pi.$$

Summing over $k$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=1} \;=\; -u\pi \sum_{k=1}^{\infty} (1-\pi)^{k-1} \sum_{j=1}^{k} (d+j).$$

We evaluate this by swapping the order of summation. Since $j$ runs from 1 to $k$ and $k$ runs from 1 to $\infty$, swapping gives $j$ running from 1 to $\infty$ and $k$ running from $j$ to $\infty$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=1} = -u\pi \sum_{j=1}^{\infty}(d+j)\sum_{k=j}^{\infty}(1-\pi)^{k-1}.$$

Evaluating the inner sum $\sum_{k=j}^{\infty}(1-\pi)^{k-1} = (1-\pi)^{j-1}/\pi$, the factor of $\pi$ cancels:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=1} = -u \sum_{j=1}^{\infty}(d+j)(1-\pi)^{j-1}.$$

Re-indexing $m = j-1$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=1} = -u \sum_{m=0}^{\infty}(d+1+m)(1-\pi)^{m}.$$

Splitting the sum using $\sum_{m=0}^{\infty}(1-\pi)^m = 1/\pi$ and $\sum_{m=0}^{\infty} m(1-\pi)^m = (1-\pi)/\pi^2$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=1} = -u\left(\frac{d+1}{\pi} + \frac{1-\pi}{\pi^2}\right) = -\frac{u(d\pi+1)}{\pi^2}.$$

Therefore, the first-order Taylor expansion of $L_{d+1}$ around $\theta_\pi = 1$ gives:

$$L_{d+1} \approx \frac{u}{\pi} - \frac{u(d\pi+1)}{\pi^2}(\theta_\pi - 1).$$

*Ratio across consecutive durations.* Applying the same expansion at $d+1$:

$$L_{d+2} \approx \frac{u}{\pi} - \frac{u((d+1)\pi+1)}{\pi^2}(\theta_\pi - 1).$$

Forming the ratio and linearizing around $\theta_\pi = 1$ using the quotient rule $f/g \approx 1 + (f'(1) - g'(1))/g(1) \cdot (\theta_\pi - 1)$, where $f(1) = g(1) = u/\pi$:

$$\frac{L_{d+2}}{L_{d+1}} \approx 1 + \frac{-\frac{u((d+1)\pi+1)}{\pi^2} - \left(-\frac{u(d\pi+1)}{\pi^2}\right)}{u/\pi}(\theta_\pi - 1).$$

The numerator difference simplifies as

$$-\frac{u}{\pi^2}\big[((d+1)\pi+1) - (d\pi+1)\big] = -\frac{u\pi}{\pi^2} = -\frac{u}{\pi},$$

so that

$$\frac{L_{d+2}}{L_{d+1}} \approx 1 - \frac{u/\pi}{u/\pi}(\theta_\pi - 1) = 2 - \theta_\pi.$$

*Across-duration result for $R^{T_d}$.* Since $R^{T_d} = (\tau_d/c) \cdot L_{d+1}$ and $\tau_{d+1} = \theta_\tau \tau_d$ exactly:

$$\frac{R^{T_{d+1}}}{R^{T_d}} \;=\; \frac{\tau_{d+1}}{\tau_d} \cdot \frac{L_{d+2}}{L_{d+1}} \;\approx\; \theta_\tau(2-\theta_\pi),$$

and therefore:

$$\frac{R^{T_{d+1}} - R^{T_d}}{R^{T_d}} \;\approx\; \theta_\tau(2-\theta_\pi) - 1 \;\approx\; \theta_\tau - \theta_\pi,$$

where the last step drops the second-order term $(\theta_\tau - 1)(1-\theta_\pi)$ and holds in a neighborhood of $\theta_\pi = 1$. The result is exact in $\theta_\tau$.

**Part (i).** Since $R^{T_d} = (\tau_d/c) \cdot L_{d+1}$ and $L_{d+1} \approx u/\pi$ to zeroth order, we have

$$R^{T_d}(\pi, \tau) \;\propto\; \frac{\tau_d \cdot u}{c\,\pi},$$

so that $R^{T_d}$ is proportional to $\tau_d$ and inversely proportional to $\pi$. Log-linearizing around any baseline $(\pi, \tau)$:

$$\frac{R^{T_d}(\pi', \tau') - R^{T_d}(\pi, \tau)}{R^{T_d}(\pi, \tau)} \;\approx\; \frac{\tau' - \tau}{\tau} - \frac{\pi' - \pi}{\pi}. \qquad \square$$

## A.2 Additional Propositions

**Proposition 3.** ***Approximate returns to targeting under geometric heterogeneity.*** *Suppose baseline job-finding rates, utility losses, and treatment effects evolve geometrically: $\pi_{d+1} = \theta_\pi \pi_d$, $u_{d+1} = \theta_u u_d$, and $\tau_{d+1} = \theta_\tau \tau_d$, with initial values $(\pi_0, \tau_0, u_0) \equiv (\pi, \tau, u)$. One-period targeting returns $R_d^T$ (defined in equation 3) are approximately proportional to treatment effects and utility losses, and inversely proportional to baseline job-finding probabilities,*

*(i) across types. For two individuals with initial characteristics $(\pi, \tau, u)$ and $(\pi', \tau', u')$ at the same duration $d$:*

$$\frac{R_d^T(\pi', \tau', u') - R_d^T(\pi, \tau, u)}{R_d^T(\pi, \tau, u)} \;\approx\; \frac{\tau' - \tau}{\tau} + \frac{u' - u}{u} - \frac{\pi' - \pi}{\pi}.$$

*(ii) across durations. For the same individual across consecutive durations:*

$$\frac{R_{d+1}^T(\pi, \tau, u) - R_d^T(\pi, \tau, u)}{R_d^T(\pi, \tau, u)} \;\approx\; \theta_\tau + \theta_u - \theta_\pi - 1,$$

*where the approximation is exact in $\theta^\tau$ and holds in a neighborhood of $\theta^\pi = \theta^u = 1$.*

*Proof.* Since utility losses now vary with duration according to $u_{d+k} = \theta_u^{d+k} u$, the continuation loss from unemployment from period $d+1$ onward is

$$L_{d+1}(\pi, \tau, u) = \sum_{\delta=d+1}^{\infty} u_\delta \, S_{\delta|d+1}(\pi),$$

where $S_{d+1|d+1}(\pi,\tau) = 1$ and for $\delta > d+1$, $S_{\delta|d+1}(\pi) = \prod_{j=1}^{\delta-d-1}(1-\pi_{d+j})$ is the probability of remaining unemployed from $d+1$ to $\delta$. Setting $k = \delta - d - 1$ and using $\pi_{d+j} = \theta_\pi^{d+j}\pi$ and $u_{d+k} = \theta_u^{d+k}u$, we write

$$L_{d+1} = u\sum_{k=0}^{\infty}\theta_u^{d+k}\prod_{j=1}^{k}(1-\theta_\pi^{d+j}\pi).$$

At $\theta_\pi = \theta_u = 1$, all $\pi_{d+j} = \pi$ and $\theta_u^{d+k} = 1$, so

$$L_{d+1} = u\sum_{k=0}^{\infty}(1-\pi)^k = \frac{u}{\pi}.$$

The one-period return to targeting is

$$R^{T_d}(\pi,\tau,u) \;=\; \frac{\tau_d}{c}\cdot L_{d+1}(\pi,\tau,u).$$

**Part (ii).** We expand $L_{d+1}$ to first order around $\theta_\pi = \theta_u = 1$:

$$L_{d+1} \;\approx\; \frac{u}{\pi} + \left.\frac{\partial L_{d+1}}{\partial\theta_\pi}\right|_{\theta_\pi=\theta_u=1}(\theta_\pi - 1) \;+\; \left.\frac{\partial L_{d+1}}{\partial\theta_u}\right|_{\theta_\pi=\theta_u=1}(\theta_u - 1).$$

*Derivative of $L_{d+1}$ with respect to $\theta_\pi$.* We write

$$L_{d+1} = u\left[1 + \sum_{k=1}^{\infty}\theta_u^{d+k}\prod_{j=1}^{k}(1-\theta_\pi^{d+j}\pi)\right],$$

where the $k=0$ term is independent of $\theta_\pi$, contributing zero to the derivative. For $k \geq 1$, by the product rule, differentiating the $j$-th factor $(1-\theta_\pi^{d+j}\pi)$ with respect to $\theta_\pi$ gives $-(d+j)\theta_\pi^{d+j-1}\pi$, while the remaining $k-1$ factors each equal $(1-\pi)$ at $\theta_\pi = 1$. The prefactor $\theta_u^{d+k}$ equals 1 at $\theta_u = 1$. Therefore:

$$\left.\frac{\partial}{\partial\theta_\pi}\prod_{j=1}^{k}(1-\theta_\pi^{d+j}\pi)\right|_{\theta_\pi=\theta_u=1} = -(1-\pi)^{k-1}\sum_{j=1}^{k}(d+j)\pi.$$

Summing over $k$:

$$\left.\frac{\partial L_{d+1}}{\partial\theta_\pi}\right|_{\theta_\pi=\theta_u=1} \;=\; -u\pi\sum_{k=1}^{\infty}(1-\pi)^{k-1}\sum_{j=1}^{k}(d+j).$$

Swapping the order of summation, with $j$ running from 1 to $\infty$ and $k$ from $j$ to $\infty$:

$$\left.\frac{\partial L_{d+1}}{\partial\theta_\pi}\right|_{\theta_\pi=\theta_u=1} \;=\; -u\pi\sum_{j=1}^{\infty}(d+j)\sum_{k=j}^{\infty}(1-\pi)^{k-1}.$$

Evaluating the inner sum $\sum_{k=j}^{\infty}(1-\pi)^{k-1} = (1-\pi)^{j-1}/\pi$, the factor of $\pi$ cancels:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=\theta_u=1} \;=\; -u\sum_{j=1}^{\infty}(d+j)(1-\pi)^{j-1}.$$

Re-indexing $m = j-1$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=\theta_u=1} \;=\; -u\sum_{m=0}^{\infty}(d+1+m)(1-\pi)^{m}.$$

Splitting using $\sum_{m=0}^{\infty}(1-\pi)^m = 1/\pi$ and $\sum_{m=0}^{\infty} m(1-\pi)^m = (1-\pi)/\pi^2$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=\theta_u=1} \;=\; -u\left(\frac{d+1}{\pi}+\frac{1-\pi}{\pi^2}\right) \;=\; -\frac{u(d\pi+1)}{\pi^2}.$$

*Derivative of $L_{d+1}$ with respect to $\theta_u$.* Differentiating $\theta_u^{d+k}$ directly gives $(d+k)\theta_u^{d+k-1}$, which equals $(d+k)$ at $\theta_u = 1$. The product $\prod_{j=1}^{k}(1-\theta_\pi^{d+j}\pi)$ equals $(1-\pi)^k$ at $\theta_\pi = 1$. Therefore:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_u}\right|_{\theta_\pi=\theta_u=1} \;=\; u\sum_{k=0}^{\infty}(d+k)(1-\pi)^{k}.$$

Splitting using $\sum_{k=0}^{\infty}(1-\pi)^k = 1/\pi$ and $\sum_{k=0}^{\infty} k(1-\pi)^k = (1-\pi)/\pi^2$:

$$\left.\frac{\partial L_{d+1}}{\partial \theta_u}\right|_{\theta_\pi=\theta_u=1} \;=\; u\left(\frac{d}{\pi}+\frac{1-\pi}{\pi^2}\right) \;=\; \frac{u((d-1)\pi+1)}{\pi^2}.$$

*First-order expansion of $L_{d+1}$.* Combining both derivatives:

$$L_{d+1} \;\approx\; \frac{u}{\pi} - \frac{u(d\pi+1)}{\pi^2}(\theta_\pi - 1) + \frac{u((d-1)\pi+1)}{\pi^2}(\theta_u - 1).$$

*Ratio across consecutive durations.* Applying the same expansion at duration $d+1$ (replacing $d$ by $d+1$ throughout):

$$L_{d+2} \;\approx\; \frac{u}{\pi} - \frac{u((d+1)\pi+1)}{\pi^2}(\theta_\pi - 1) + \frac{u(d\pi+1)}{\pi^2}(\theta_u - 1).$$

Forming the ratio using $f/g \approx 1 + (f'(1) - g'(1))/g(1) \cdot (\theta_\pi - 1, \theta_u - 1)$ with $f(1) = g(1) = u/\pi$, the differences of the derivative coefficients are:

$$\Delta_\pi \;=\; -\frac{u((d+1)\pi+1)}{\pi^2} - \left(-\frac{u(d\pi+1)}{\pi^2}\right) \;=\; -\frac{u\pi}{\pi^2} \;=\; -\frac{u}{\pi},$$

$$\Delta_u \;=\; \frac{u(d\pi+1)}{\pi^2} - \frac{u((d-1)\pi+1)}{\pi^2} \;=\; \frac{u\pi}{\pi^2} \;=\; \frac{u}{\pi}.$$

Note that both $d$-dependent terms cancel exactly in each difference, so that

$$\frac{L_{d+2}}{L_{d+1}} \;\approx\; 1 - \frac{u/\pi}{u/\pi}(\theta_\pi - 1) + \frac{u/\pi}{u/\pi}(\theta_u - 1) \;=\; 1 + \theta_u - \theta_\pi.$$

*Across-duration result for $R^{T_d}$.* Since $R^{T_d} = (\tau_d/c) \cdot L_{d+1}$ and $\tau_{d+1} = \theta_\tau \tau_d$ exactly:

$$\frac{R^{T_{d+1}}}{R^{T_d}} \;=\; \frac{\tau_{d+1}}{\tau_d} \cdot \frac{L_{d+2}}{L_{d+1}} \;\approx\; \theta_\tau(1 + \theta_u - \theta_\pi),$$

and therefore:

$$\frac{R^{T_{d+1}} - R^{T_d}}{R^{T_d}} \;\approx\; \theta_\tau(1 + \theta_u - \theta_\pi) - 1 \;\approx\; \theta_\tau + \theta_u - \theta_\pi - 1,$$

where the last step drops the second-order terms $(\theta_\tau - 1)(\theta_u - 1)$, $(\theta_\tau - 1)(1 - \theta_\pi)$, and $(\theta_\tau - 1)(\theta_u - \theta_\pi)$, and holds in a neighborhood of $\theta_\pi = \theta_u = 1$. The result is exact in $\theta_\tau$.

**Part (i).** Since $L_{d+1} \approx u/\pi$ to zeroth order, we have

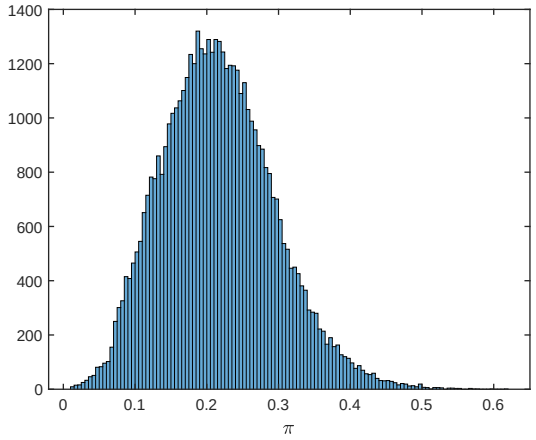
$$R^{T_d}(\pi, \tau, u) \;\propto\; \frac{\tau_d \cdot u}{c\,\pi},$$

so that $R^{T_d}$ is proportional to both $\tau_d$ and $u$, and inversely proportional to $\pi$. Log-linearizing around any baseline $(\pi, \tau, u)$ at fixed duration $d$:

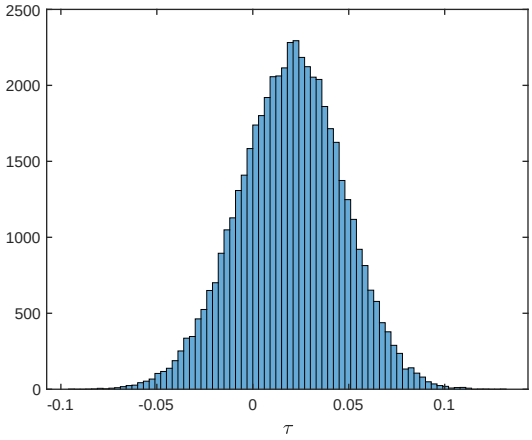
$$\frac{R^{T_d}(\pi', \tau', u') - R^{T_d}(\pi, \tau, u)}{R^{T_d}(\pi, \tau, u)} \;\approx\; \frac{\tau' - \tau}{\tau} + \frac{u' - u}{u} - \frac{\pi' - \pi}{\pi}. \qquad \square$$

## A.3 Additional Figures

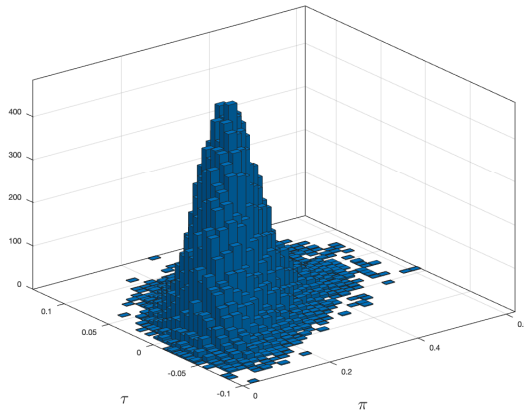
Figure A1: Distribution of $\pi$ and $\tau$ and Duration Dependence in Job Finding in Calibrated Model



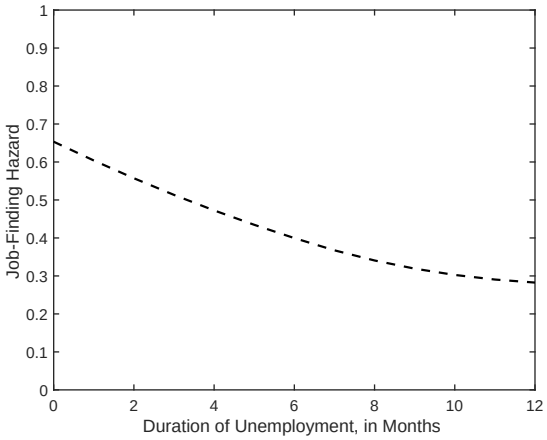

(a) Histogram of $\pi$




(b) Histogram of $\tau$



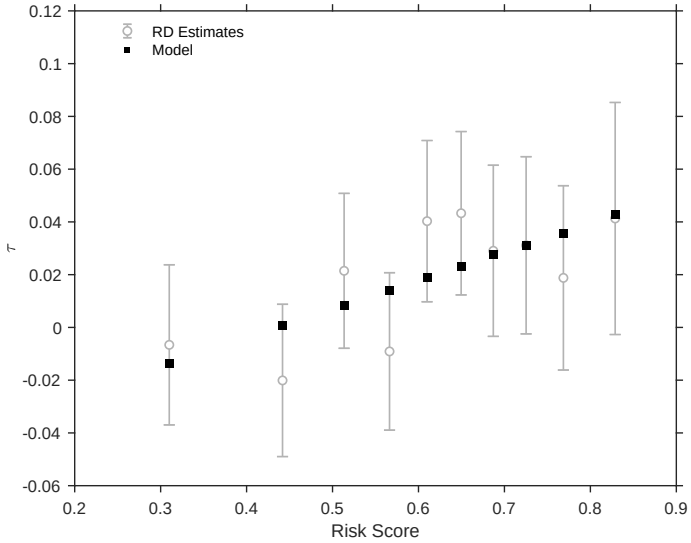

(c) Histogram of $\pi$ and $\tau$



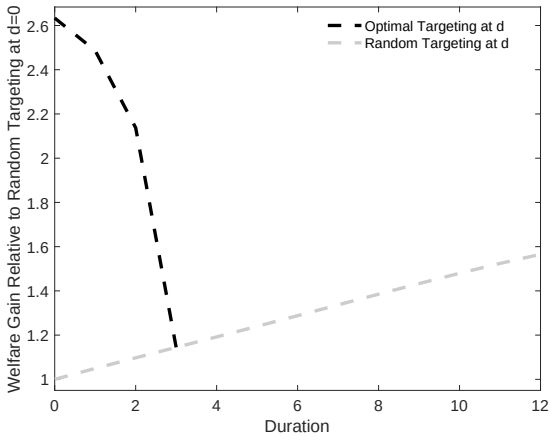

(d) Duration Dependence in Job Finding

*Notes:* The histograms shows the distributions of the 50,000 individuals sampled for our model calibration.

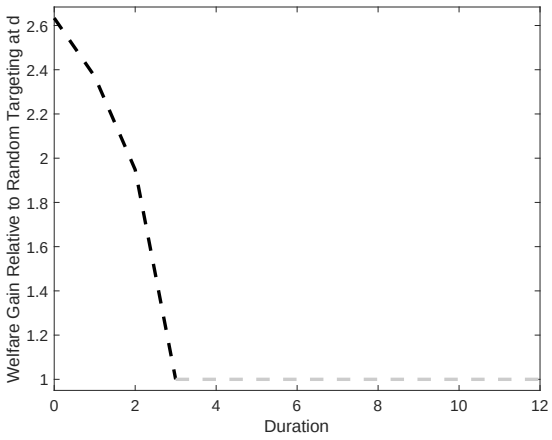
Figure A2: Treatment Effects by Decile of RS Distribution, RD Estimates vs. Model



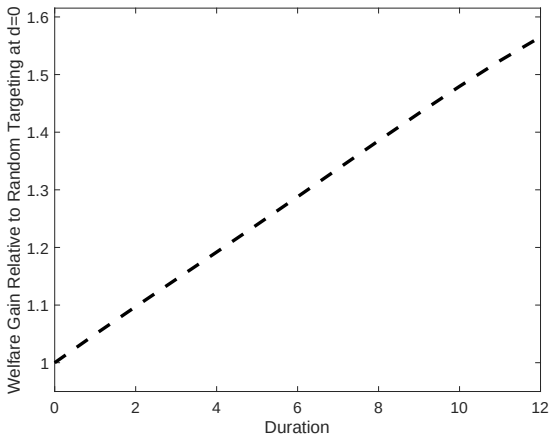

*Note:* The outcome variable for the RD estimates is the probability of working within the next four weeks.



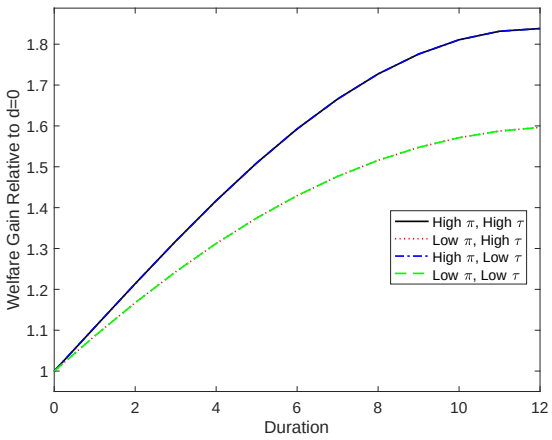

(a) Optimal Targeting at $d$ vs. Random at $d = 0$




(b) Optimal Targeting at $d$ vs. Random at $d$




(c) Random Targeting at $d$ vs. Random at $d = 0$

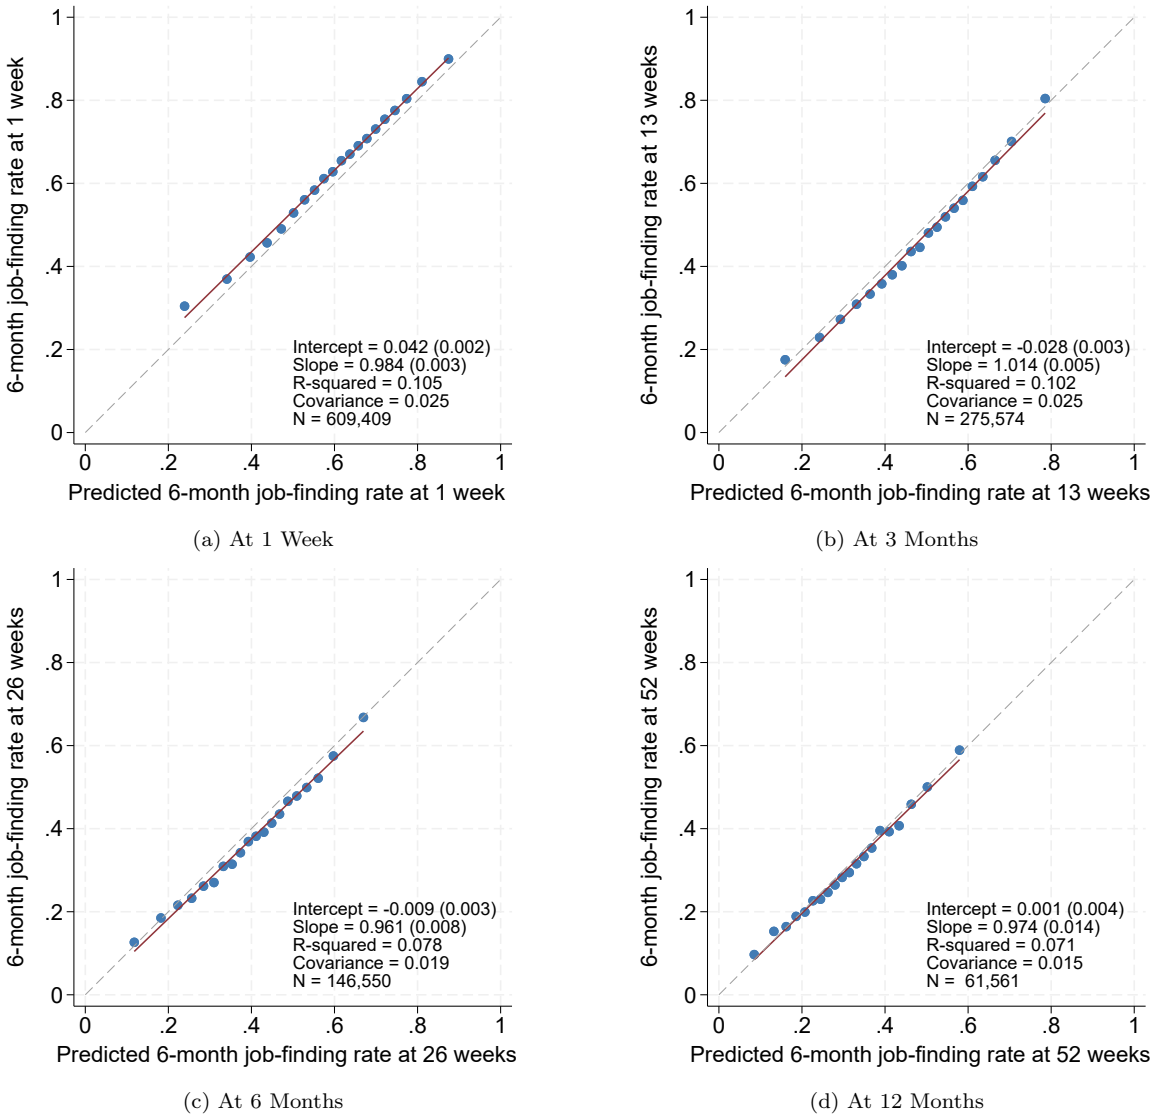



(d) Targeting at $d$ vs. $d = 0$, By Type

Figure A3: Welfare Gains of Targeting, By Duration

# B Data

## B.1 Risk Score Accuracy




(a) At 1 Week

(b) At 3 Months

(c) At 6 Months

(d) At 12 Months

Figure A4: Employment Predictions vs. Outcomes

*Note:* This figure shows binned scatter plots comparing predicted 6-month job-finding rates and empirical 6-month job-finding rates.

# C Effects of Job Search Counseling

## C.1 Risk Score RD: Bunching

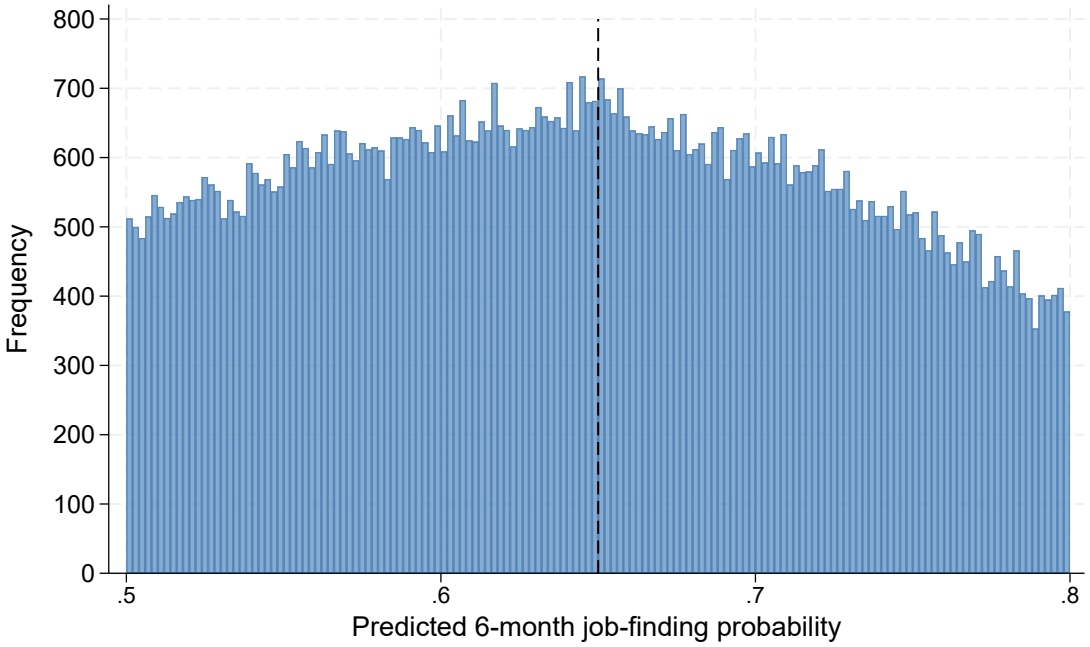



Figure A5: Distribution of Risk Scores

*Note:* This figure shows the frequency distribution of VDAB's predicted 6-month job-finding rates around the cut-off.

## C.2 Risk Score RD: Kernel and Bandwidth

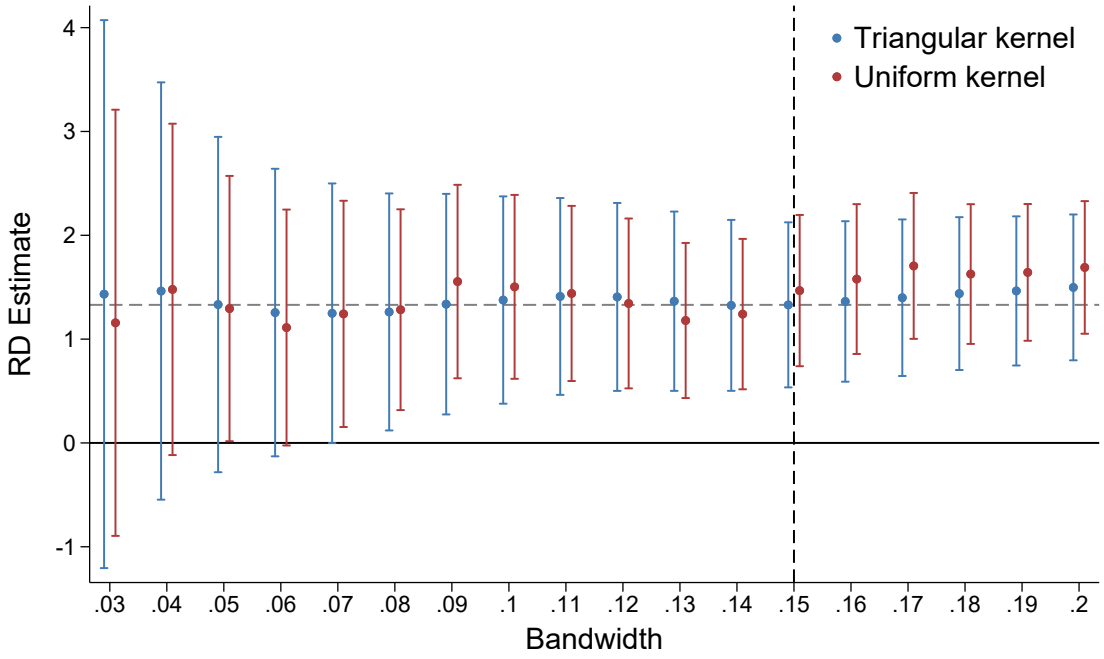



Figure A6: Estimates by Kernel and Bandwidth

*Note:* This figure plots IV estimates of the effect of caseworkers on days worked per week estimated using the risk score RD over weeks 6-17 since the start of unemployment. The vertical line represents the bandwidth used for our baseline specification. The horizontal gray line marks our baseline estimate. We display 95% confidence intervals for the estimated effects.

## C.3 Risk Score RD: Balance Tests

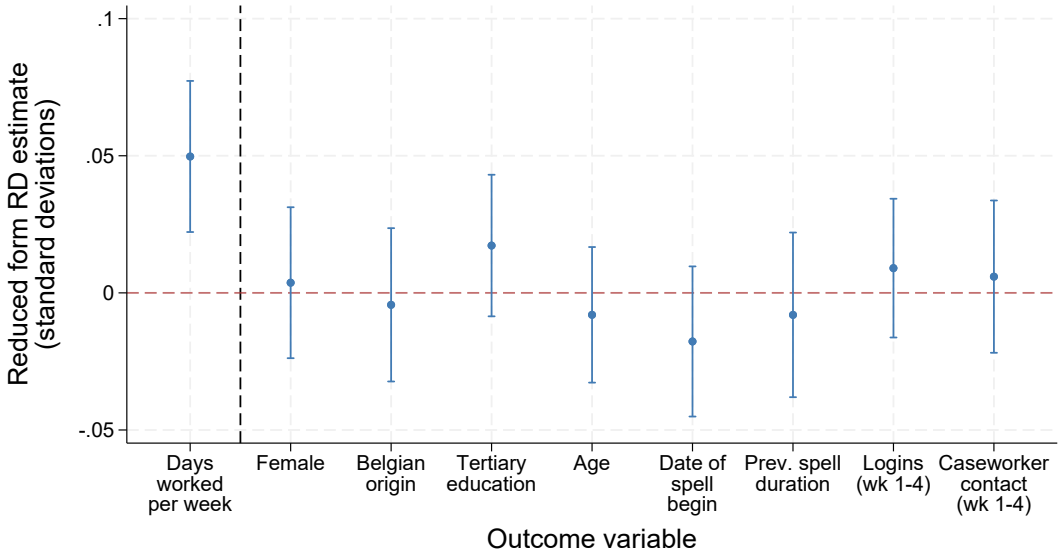



Figure A7: Estimates by Kernel and Bandwidth

*Note:* This figure plots reduced form estimates of the effect of being below the cut-off on days worked per week (our baseline outcome) as well as various predetermined covariates. We display 95% confidence intervals for the estimated effects.

## C.4 Risk Score RD: Other Outcomes

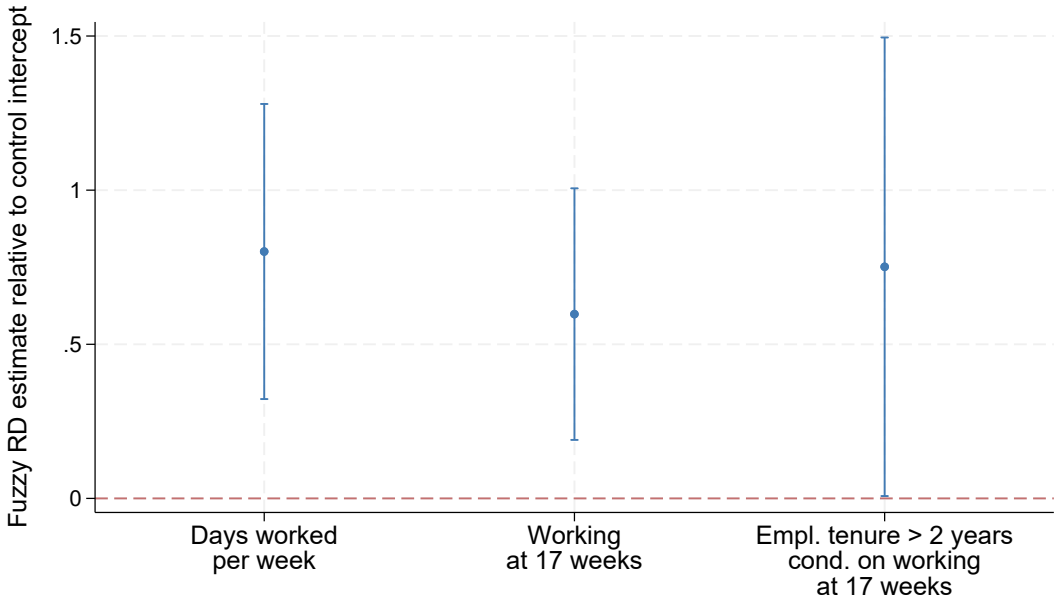
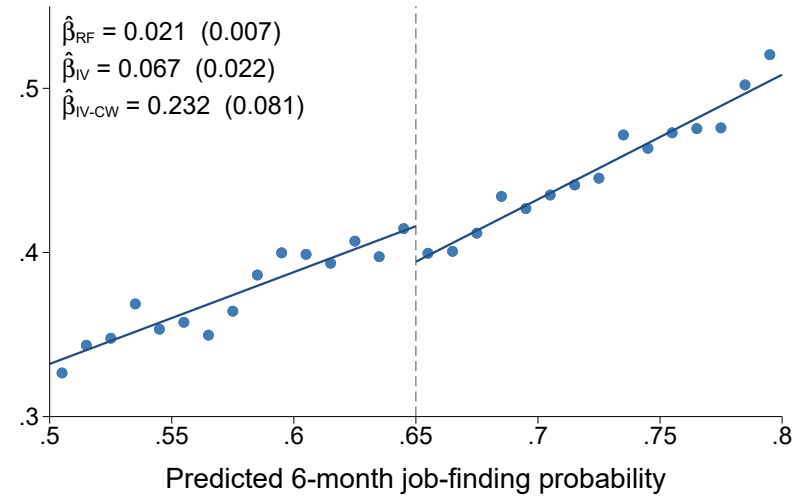
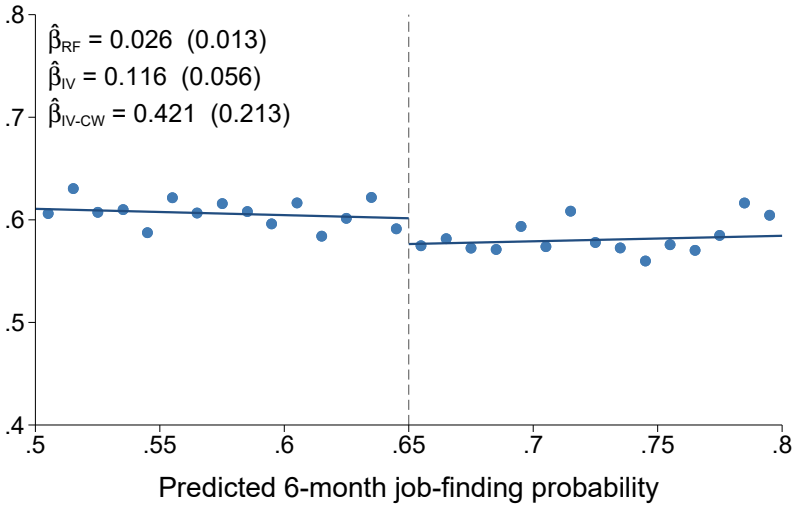



Figure A8: Other Employment Outcomes

*Note:* This figure plots IV estimates of the effect of caseworkers on three outcomes using the risk score RD. The first is days worked per week from weeks 6–17 (the baseline outcome). The second is employment at week 17, conditional on being unemployed after 5 weeks. The third is entry into an employment spell lasting at least two years conditional on entering employment in weeks 6–17. We display 95% confidence intervals for the estimated effects.




(a) Employment




(b) Employment Tenure

Figure A9: Risk Score Regression Discontinuity

*Note:* The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment, thereby excluding those who had already exited unemployment or been assigned to a caseworker. Panel (a) plots the proporiton of employed job seekers after 17 weeks conditional on being unemployed after 5 weeks. Panel (b) plots the proportion of job seekers who are employed in an employment spell lasting at least two years, conditional on entering employment in weeks 6–17. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean outcome. This variable is then plotted. The controls are: quarter-year, gender, education, age, arrondissement, average previous unemployment duration, indicators for whether the last job or current unemployment spell began on the first of the month, and the day of the week on which the last job and the current unemployment spell began. We report 95% confidence intervals. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

# D Heterogeneity in Effects

## D.1 Employment Spell RD: Bunching

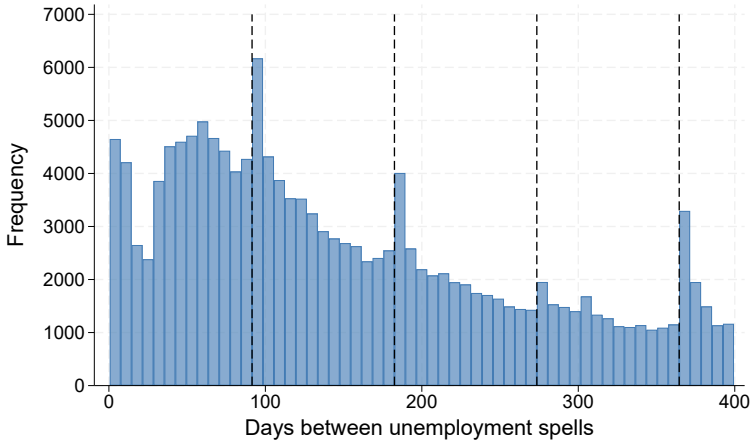



Figure A10: Frequency Distribution of Days Between Unemployment Spells (0–400 Days)

*Note:* This figure plots the distribution of the days between unemployment spells, applying the sample restrictions used in the employment spell RD and using a bin width of 7 days. The vertical lines mark 3, 6, 9 and 12 months.

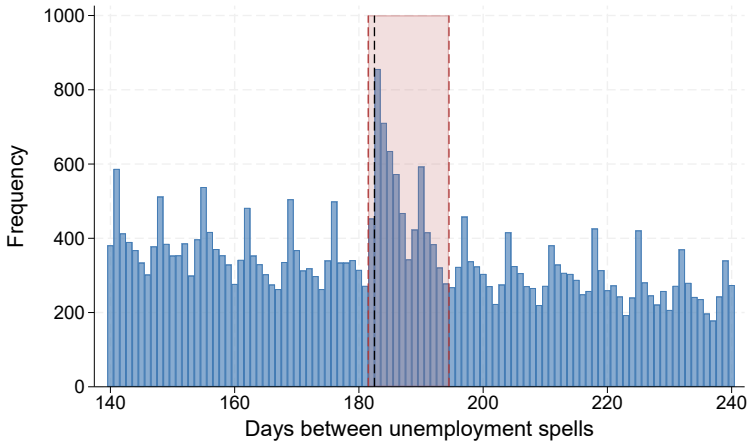
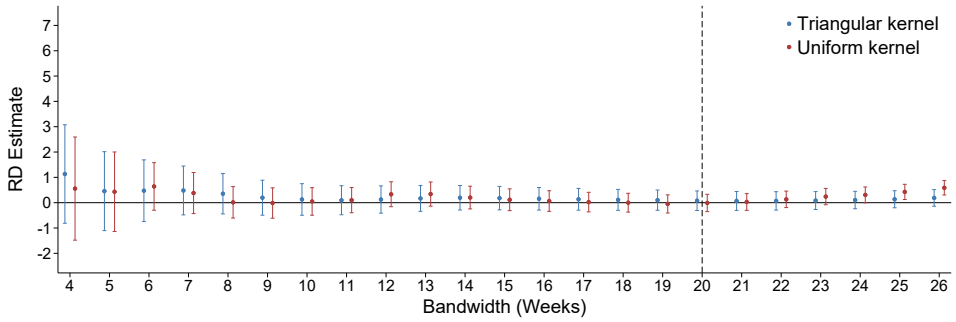


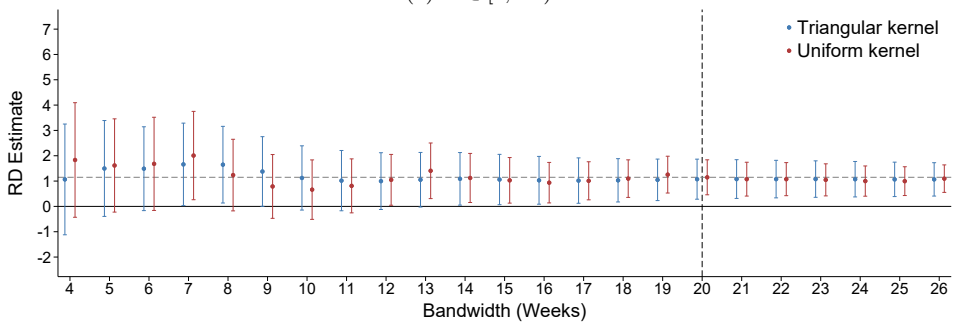

Figure A11: Frequency Distribution of Days Between Unemployment Spells (140–240 Days)

*Note:* This figure plots the distribution of the days between unemployment spells, applying the sample restrictions used in the employment spell RD and using a bin width of 1 day. The highlighted region (182–194 days) is the 'donut' hole, the excluded region in the baseline specification.

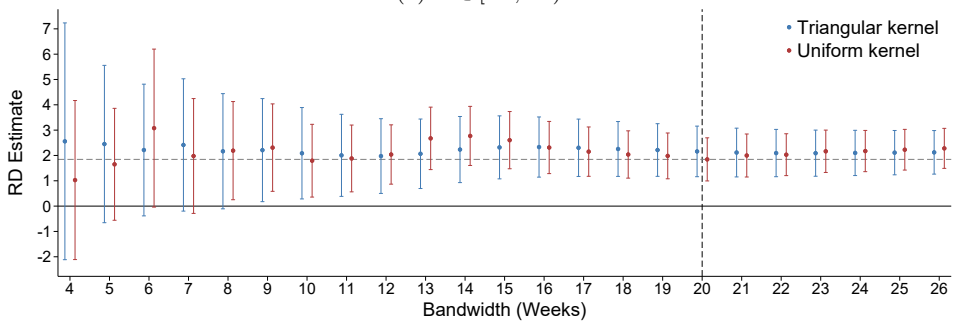
## D.2 Employment Spell RD: Kernel and Bandwidth




(a) $P \in [0, 0.6)$




(b) $P \in [0.6, 0.7)$

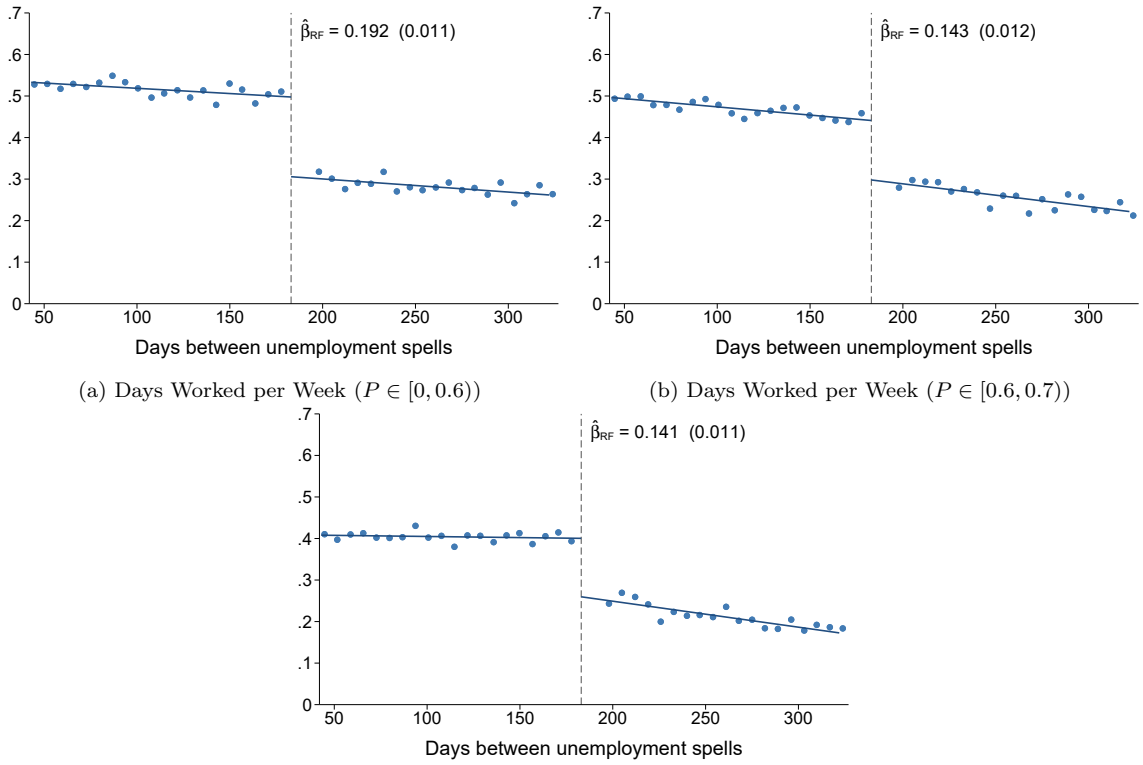



(c) $P \in [0.7, 1]$

Figure A12: Estimates by Kernel and Bandwidth

*Note:* These panels plot IV estimates of the effect of caseworkers on days worked per week estimated using the employment spell RD over weeks 1-4 since the start of unemployment. The vertical lines represent the bandwidth used for our baseline specification. The horizontal gray line marks our baseline estimate. We display 95% confidence intervals for the estimated effects.

## D.3 Employment Spell RD: First Stage




(a) Days Worked per Week ($P \in [0, 0.6)$)

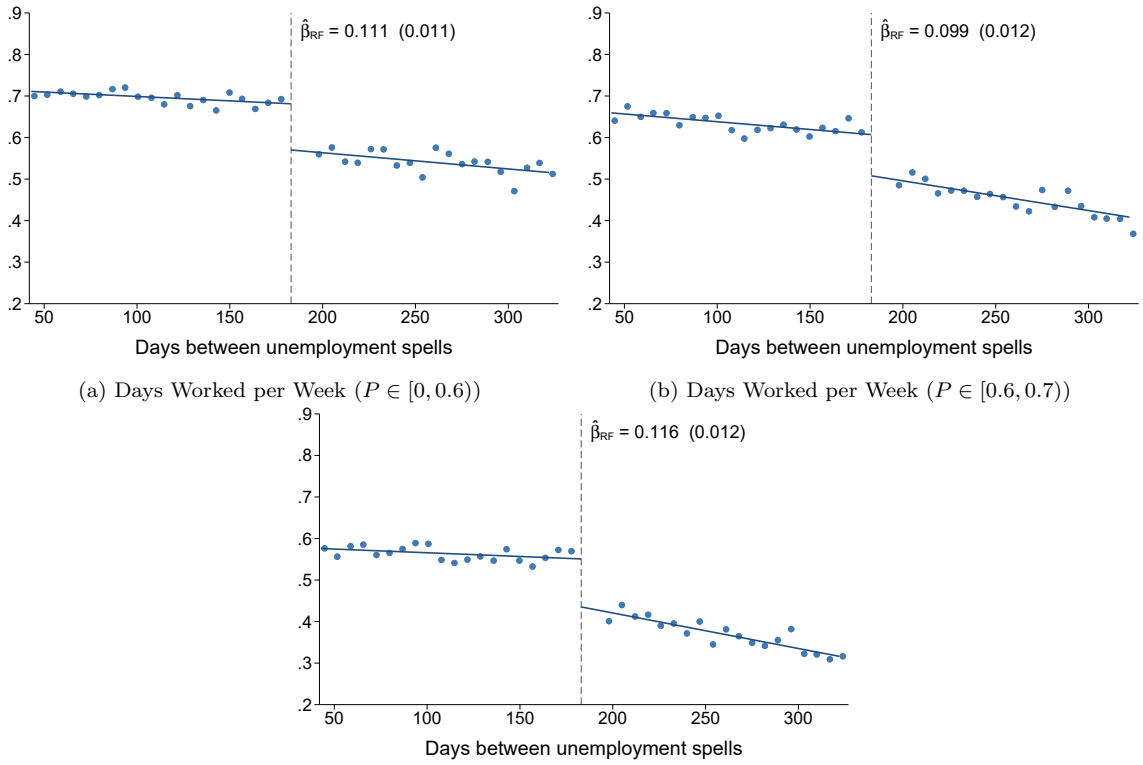
(b) Days Worked per Week ($P \in [0.6, 0.7)$)

(c) Days Worked per Week ($P \in [0.7, 1]$)

Figure A13: Caseworker Contact (Weeks 1–4)

*Note:* The sample is restricted to job seekers who had contact with a caseworker in their previous spell. Job seekers are then split into three groups by the probability of finding a job within 6 months $P$ that was predicted for them after 1 week of unemployment. Panels (a) to (c) the proportion of job seekers who have contact with a caseworker in weeks 1–4 of unemployment for each risk group. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

## D.4 Employment Spell RD: Dynamics



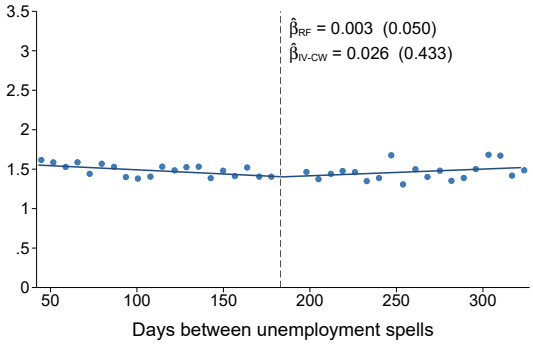

(a) Days Worked per Week ($P \in [0, 0.6)$)

(b) Days Worked per Week ($P \in [0.6, 0.7)$)

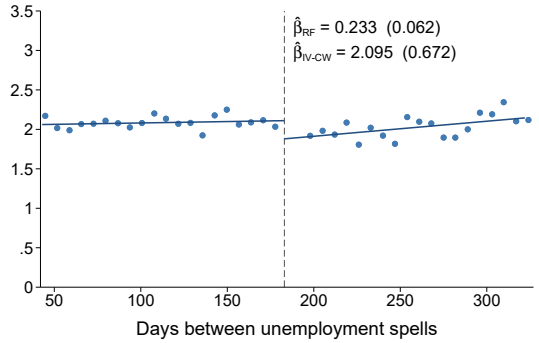
(c) Days Worked per Week ($P \in [0.7, 1]$)

Figure A14: Caseworker Contact (Weeks 1–16)

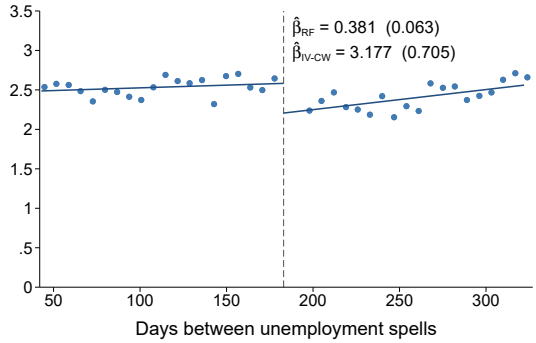
*Note:* The sample is restricted to job seekers who had contact with a caseworker in their previous spell. Job seekers are then split into three groups by the probability of finding a job within 6 months $P$ that was predicted for them after 1 week of unemployment. Panels (a) to (c) the proportion of job seekers who have contact with a caseworker in weeks 1–16 of unemployment for each risk group. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.




(a) Days Worked per Week ($P \in [0, 0.6)$)




(b) Days Worked per Week ($P \in [0.6, 0.7)$)




(c) Days Worked per Week ($P \in [0.7, 1]$)

Figure A15: Days Worked per Week (Weeks 1–16)

*Note:* The sample is restricted to job seekers who had contact with a caseworker in their previous spell. Job seekers are then split into three groups by the probability of finding a job within 6 months $P$ that was predicted for them after 1 week of unemployment. Panels (a) to (c) show the RD plots for each risk group, where the y-axis shows mean days worked per week plus the residual of days worked per week regressed on a set of controls. The window of observation is weeks 1–16 since the start of unemployment. The controls used are: quarter-year, gender, education, age, arrondissement, average previous unemployment duration, indicators for whether the last job or current unemployment spell began on the first of the month, and the day of the week on which the last job and the current unemployment spell began. We report 95% confidence intervals. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

## D.5 Employment Spell RD: Program Participation

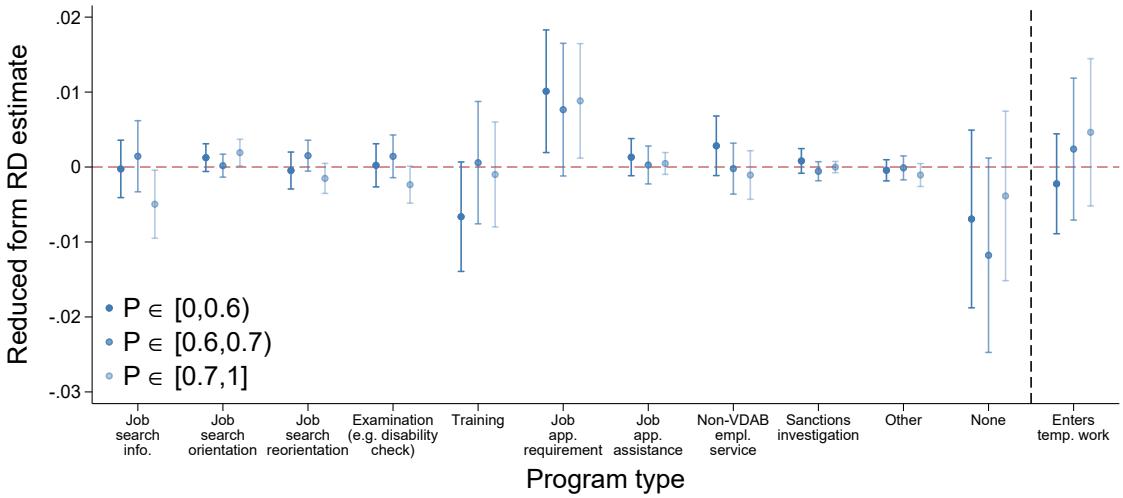



Figure A16: Program Participation

*Note:* This figure plots reduced form estimates of the effect of being below the cut-off on participation in various programs by risk group. We display 95% confidence intervals for the estimated effects.