At-Risk or Responsive?
Targeting in Unemployment Policy

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

July 28, 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: whom 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 little 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 assignment.

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).

1 Introduction

A central challenge for governments operating under limited resources is how to allocate interventions where they create the greatest social value. Public employment services must decide which unemployed workers receive intensive job search assistance [Berger et al., 2001; Black et al., 2003]. Health systems must decide which patients should receive treatment first [Obermeyer et al., 2019], while poverty alleviation programs determine which households receive transfers [Haushofer et al., 2025]. Increasingly, these decisions are guided by predictive algorithms that estimate which individuals are most at risk of poor outcomes.

Yet prediction and targeting are fundamentally different problems [Kleinberg et al., 2015; Athey, 2017]. Intuitively, a planner allocating scarce resources should care not only about who is most at risk, but also about who stands to benefit most from the intervention. Whether these two objects coincide is ultimately an empirical question. Nevertheless, many policy applications implicitly assume that they do. Public employment services across Europe, for example, increasingly prioritize job seekers predicted to remain unemployed longest, thereby treating measures of risk as measures of treatment value [Desiere et al., 2019].

This paper studies the distinction between prediction and targeting and shows that it is first-order for the design of unemployment policy. We consider the allocation of job search counseling by the Flemish Public Employment Service (VDAB), which uses a machine-learning model to prioritize workers predicted to be at greatest risk of long-term unemployment. This setting provides an unusually rich opportunity to study how scarce interventions should be allocated when both unemployment risk and treatment responsiveness are heterogeneous. We make three contributions. First, we develop a conceptual framework that characterizes the trade-offs of targeting job search counseling across both workers and unemployment durations. Second, we provide quasi-experimental evidence on how treatment effects vary across these two dimensions, exploiting three complementary regression discontinuity (RD) designs that separately identify average effects, heterogeneity by unemployment risk, and the evolution of treatment effects over the unemployment spell. Third, we combine the conceptual framework with the empirical estimates to evaluate alternative targeting policies, showing how employment services should jointly decide whom to treat and when to intervene.

We develop our framework in the context of unemployment, where job search counseling increases the probability that workers transition from unemployment into employment. However, the results generalize to other settings where policymakers seek to accelerate transitions out of adverse states (e.g., illness, poverty, or homelessness). 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. We show that the social return to treatment depends jointly on its effect on the transition probability and on the continuation value of remaining in the adverse state. In fact, these two objects combine in a remarkably simple sufficient statistic. To a first-order approximation, the return to treatment is proportional to the treatment effect τ divided by the underlying transition probability p,

$$R \approx \frac{\tau}{p}.$$

Intuitively, a given increase in the probability of finding employment is more valuable for workers who would otherwise remain unemployed for longer. This simple characterization provides the organizing principle for the remainder of the paper.

First and foremost, it provides a unified way of thinking about both *whom* to target and *when* to intervene. Across individuals, it highlights the fundamental trade-off between risk and responsiveness. Workers with lower job-finding probabilities have more to gain from leaving unemployment, but they need not be those who respond most to job search counseling. The optimal allocation therefore depends on the balance between these two forces rather than either object in isolation.

The same logic applies across unemployment durations. Interestingly, the return can change dynamically over the unemployment spell, because treatment effects and job-finding probabilities may change as unemployment lengthens, but also because the composition of the unemployed changes through dynamic selection. The optimal timing of intervention therefore depends jointly on duration dependence in treatment effects and in baseline job-finding probabilities. Together, these results provide estimable sufficient statistics for allocating scarce interventions across both individuals and time.

To bring the framework to the data, we study the allocation of job search counseling by the VDAB. Since 2018, the VDAB has relied on a machine-learning model using more than 400 granular predictors to estimate each unemployed worker's *risk score*, defined as the probability of finding employment within six months. Job seekers whose predicted job-finding probability falls below a threshold of 0.65 are prioritized for counseling: they are proactively contacted by the service line and considered for assignment to a caseworker. Those above the threshold are classified as self-reliant and are not actively contacted for the next three months. The VDAB's system exemplifies the international trend toward algorithmic profiling in active labor market policy [Desiere et al., 2019]. 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.

The combination of institutional detail, data richness, and quasi-experimental variation provides an unusually clean setting in which to study how risk-based targeting performs in practice. In particular, the institutional setting generates three complementary sources of quasi-experimental variation that map directly into the two allocation problems studied in the conceptual framework. First, a regression discontinuity around the algorithmic risk-score threshold identifies the average effects of job search counseling for workers near the margin of treatment. Second, a discontinuity generated by the automatic reassignment of recently unemployed workers to their previous caseworker allows us to estimate treatment effects across the distribution of predicted unemployment risk, providing evidence on whom to target. Third, the risk-score discontinuity reappears with the opposite sign later in the spell, when the VDAB reaches out to workers who have remained unemployed without previously being assigned to a caseworker. This design helps identify how the value of intervention evolves over the unemployment spell.

Our empirical analysis delivers three main sets of findings.

First, job search counseling has substantial average effects on employment. Workers prioritized for counseling find employment sooner, spending significantly more days in work. If anything, they also transition into more stable employment relationships, suggesting that the intervention improves job matches rather than simply accelerating acceptance of lower-quality jobs. Quantitatively, scaling the employment effects by the discontinuity in caseworker contact implies that counseling raises days worked by roughly 80% over the following three months for the marginal job seeker. This establishes that the allocation problem is quantitatively important.

Second, treatment effects vary sharply across workers. Contrary to the logic underlying current targeting practice, we find a strong negative relationship between unemployment risk and treatment responsiveness. We find a precise zero effect of job search counseling for workers with the highest risk of long-term unemployment — the very workers prioritized by the existing targeting algorithm. By contrast, for workers with better employment prospects, we find large and statistically significant effects. The planner therefore faces a first-order trade-off between targeting workers who are most at risk and those for whom treatment generates the largest returns.

Beyond this comparison across risk-score groups, we find substantial additional heterogeneity in treatment effects, estimated using an instrumental forest [Athey et al., 2019]. Importantly, the estimated treatment effects based on the reassignment discontinuity also predict heterogeneity in our baseline RD, where both the population and the treatment environment differ, although this cross-design comparison is statistically imprecise. We further show that the pattern is not explained by differential assignment to programs, which mitigates concerns that the heterogeneity is driven by lock-in effects or differences in the type of support received.

To quantify the importance of this heterogeneity, we can decompose the variation in returns to treatment using the sufficient statistic, $R \approx \tau/p$. Based on a Shapley decomposition of the variance in the targeting index τ/p, we find that 78% of the variation is attributable to heterogeneity in treatment effects. Moreover, the negative covariance between treatment effects and inverse job-finding probabilities offsets part of the variation arising from differences in baseline job-finding probabilities. As a result, targeting exclusively on unemployment risk misses considerable variation in expected returns to counseling.

Third, we find little evidence that treatment effects deteriorate over the unemployment spell for comparable workers. We exploit the reversal of the risk-score discontinuity when the VDAB contacts remaining unemployed workers later in the spell, allowing us to compare counseling effects for workers around the same initial risk-score threshold but treated at different unemployment durations. The estimated effects are similar across the first and second treatment rounds, although both estimates are imprecise. We subject this comparison to several validation checks, including tests for differential survival into the second round and for discontinuities in predicted treatment responsiveness among untreated survivors. These checks support the interpretation that the similar estimates reflect limited duration dependence in treatment effects for comparable workers.

This does not imply, however, that returns to treatment are constant over the unemployment spell. Workers with stronger employment prospects leave unemployment earlier, so the remaining pool becomes increasingly selected toward workers with lower job-finding probabilities and lower responsiveness to counseling. Through the lens of our sufficient statistic $R \approx \tau/p$, these forces have opposite implications for the timing of intervention. Declining treatment effects reduce the return to counseling, while declining baseline job-finding probabilities increase the continuation value of treatment for those who remain unemployed. Quantitatively, we find that average treatment effects decline by roughly 65% over the first year, while job-finding probabilities decline by roughly 57%. These two observable forces are of comparable magnitude, but the decline in job finding clearly dominates once we condition on observables — treatment effects show no decline for comparable workers, while their job-finding prospects continue to deteriorate — implying that the return to treating the average surviving job seeker rises over the spell.

Finally, we use the empirical estimates to calibrate our conceptual framework and evaluate alternative targeting policies. The welfare gains from optimal targeting are nearly three times those from 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 probability performs nearly as well as the fully optimal policy. By contrast, the current practice of targeting exclusively on risk scores performs worse than random assignment. The shortfall is quantitatively large: in our calibrated model, risk-score targeting yields only about half (53%) of the welfare gains of random assignment.

Extending the analysis to the timing dimension confirms that the average return to treating surviving job seekers rises over the unemployment spell. At the same time, the largest gains from selective targeting arise early, when many workers still exhibit high expected returns to treatment. As unemployment progresses, dynamic selection reduces heterogeneity in returns, making selective targeting progressively less valuable. Thus, while the return to treating the average surviving job seeker may rise over the spell, the gains from choosing whom to treat are greatest early on, when the pool of unemployed workers is largest and most heterogeneous.

Taken together, the results imply that employment services should not treat risk predictions as targeting rules. Risk scores are valuable inputs into allocation decisions, but optimal targeting also requires information on treatment responsiveness and on how returns evolve over the unemployment spell. In our setting, this points toward earlier intervention and substantially greater weight on predicted responsiveness, rather than relying primarily on predictions of long-term unemployment risk.

A final component of targeting we briefly investigate in our calibrated model concerns *how* targets should be estimated. The VDAB's algorithm, like most risk-profiling tools, is trained on realized employment outcomes, which reflect a mixture of baseline job-finding prospects and past treatment effects. This conflation creates adverse selection into treatment: individuals with high treatment effects will have higher predicted job-finding rates, pushing them above the threshold and out of the treatment pool. The appropriate risk measure for targeting should instead predict job finding in the absence of treatment, not in its presence. Our positive empirical finding — that treatment effects are increasing in predicted job-finding probability — may itself partly reflect this algorithmic bias. However, we find the bias to be moderate: correcting risk scores for the effects of past treatment weakens the relationship between treatment effects and predicted job finding by roughly a third, but does not overturn it. Importantly, the bias also does not invalidate our analysis of current practice: the VDAB assigns counseling based on the observed risk scores, so our quasi-experimental estimates and our evaluation of score-based targeting refer to the scores actually used for assignment.

Related Literature. Our paper contributes to several strands of literature. First, it connects two large literatures on unemployment policy that have mostly developed separately. One literature studies heterogeneity in job-finding prospects and the predictability of long-term unemployment [Ahn and Hamilton, 2020; Mueller et al., 2021; Ahn, 2023; Alvarez et al., 2024; Gregory et al., 2025; Mueller and Spinnewijn, 2025]. In the spirit of this work, we find substantial and predictable heterogeneity in unemployment risk. However, for targeting purposes, predicting who is likely to remain unemployed is not sufficient: the planner must also know who responds to intervention. A second literature estimates the effects of active labor market programs, documenting substantial heterogeneity across programs, populations, and unemployment durations [Card et al., 2018]. Consistent with our findings, existing evidence suggests that job search counseling often has larger effects for workers closer to employment, while training and information-based interventions may display different patterns [Koning, 2009; Liu et al., 2014; Altmann et al., 2018; Belot et al., 2026]. Our contribution is to bring these two objects together: we estimate heterogeneity in treatment effects across the risk distribution and show how risk and responsiveness should jointly determine the allocation of scarce employment services.

Second, our paper contributes to work on profiling, algorithmic assistance, and digital tools in public employment services. Statistical profiling has long been used to identify unemployment insurance claimants at risk of long spells and direct them toward reemployment services [Berger et al., 2001; Black et al., 2003; Desiere et al., 2019]. More recent work studies how predictive algorithms and recommender systems can improve job search and matching, including by recommending vacancies, occupations, or search strategies to job seekers [Altmann et al., 2022; Le Barbanchon et al., 2023; Behaghel et al., 2024; Belot et al., 2026]. We differ from this literature in focusing on the allocation of intensive counseling rather than the design of search recommendations themselves. Our results show that even when risk scores predict unemployment well, using them as targeting rules can perform poorly if they are negatively related to treatment responsiveness.

Third, our paper relates to a broader literature on targeting interventions when social need and treatment impact need not coincide. The closest paper is Haushofer et al. [2025], who study the trade-off between targeting deprivation and targeting impact in the context of cash transfers in Kenya. They show that targeting solely on deprivation is unattractive unless redistributive preferences are sufficiently strong, a conclusion that mirrors our finding that targeting solely on unemployment risk can be dominated by rules that also account for responsiveness. More broadly, an influential literature examines when policy problems are prediction problems and when they instead require causal knowledge, documenting both the promise of algorithmic predictions for high-stakes decisions and their pitfalls when the prediction target diverges from the object of policy interest [Kleinberg et al., 2015; Athey, 2017; Kleinberg et al., 2018; Obermeyer et al., 2019]. Related work in other domains similarly contrasts predictive targeting based on baseline outcomes with causal targeting based on heterogeneous treatment effects [Athey et al., 2025]. We contribute to this literature by developing a simple sufficient-statistic framework for targeting interventions that accelerate transitions out of adverse states, and by applying it in a dynamic setting where the planner must decide not only whom to treat but also when to intervene.

Finally, our analysis is related to methodological work on treatment targeting and policy learning under heterogeneous treatment effects [Manski, 2004; Kitagawa and Tetenov, 2018; Athey et al., 2025; Yadlowsky et al., 2025]. This literature provides tools for estimating and evaluating treatment assignment rules. Our contribution is complementary. We show that, in dynamic transition settings such as unemployment, the relevant targeting object combines predicted treatment effects with baseline transition probabilities. This structure clarifies why targeting on risk alone can fail, why treatment effects alone are also incomplete, and how sufficient statistics can guide practical allocation rules in public employment services.

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:

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

where π denotes an individual's baseline observable job-finding type, and ε^π captures unobserved heterogeneity in baseline job-finding prospects. Similarly, τ denotes the individual's predictable treatment effect, while ε^τ represents unobserved heterogeneity in treatment responsiveness. The indicator $T_d(\pi, \tau) \in \{0, 1\}$ specifies whether an individual of observed type (π, τ) 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 (π, τ). $\bar{D}$ 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 $\theta_{\pi,d}$ governs how baseline job-finding probabilities evolve with unemployment duration, while $\theta_{\tau,d}$ 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 ε^π and the unobserved treatment type ε^τ 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., $u_d = u\,\theta_{u,d}$, where the level u may differ across job seekers. We assume that u is observable to the agency, so that types are implicitly indexed by (π, τ, u) and targeting policies may condition on u; to keep notation light, we leave this dimension implicit in the type distribution $g(\pi, \tau)$ and make it explicit only in the objects that depend on it. 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. Note that we start from a world without any intervention and characterize the optimal policy for providing a single intervention. Providing treatment thus does not generate budgetary savings on other interventions that job seekers would otherwise receive over the spell. The agency's objective is to minimize the expected welfare losses associated with unemployment as follows:

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

We can prove the following proposition.

Proposition 1. ***Optimal targeting across individuals and durations.*** *Consider the agency's problem of allocating treatment across types and durations, where each type (π,τ) can be treated for one period and at most once, and assume that the distribution function $g(\pi,\tau)$ is continuous. Define the return to targeting type (π,τ) with utility-loss level u at duration d as the reduction in continuation losses per unit of treatment cost,*

$$R_d^T(\pi,\tau,u) \ \equiv \ -\frac{L_d^T(\pi,\tau,u) - L_d(\pi,u)}{c}. \tag{3}$$

Then there exists a threshold return $\bar{R} \geq 0$, determined by the budget constraint, such that the optimal policy:

(i) treats all types whose return at their optimal targeting duration exceeds the threshold,

$$T_{d^\star}(\pi,\tau,u) = 1 \iff R^T_{d^\star}(\pi,\tau,u) \geq \bar{R};$$

(ii) targets each treated type at the duration that maximizes its survival-weighted excess return,

$$d^\star(\pi,\tau,u) = \arg\max_d \ S_d(\pi,\tau)\big[R_d^T(\pi,\tau,u) - \bar{R}\big].$$

In the absence of unobserved heterogeneity and duration dependence ($\sigma_{\varepsilon^\pi} = \sigma_{\varepsilon^\tau} = 0$ and $\theta_{\pi,d} =$

$\theta_{\tau,d} = \theta_{u,d} = 1$ for all d), the return simplifies to

$$R_d^T(\pi,\tau,u) \;=\; \frac{u}{c}\cdot\frac{\tau}{\pi} \qquad \text{for all } d,$$

so that the optimal policy treats all types with $u\,\tau/\pi$ above a threshold — reducing to a ranking in τ/π when utility losses are identical across job seekers — and, because returns are constant over the spell while survival declines, targets every treated type at entry, $d^\star = 0$.

Proof. See Appendix A.1.

Proposition 1 separates the targeting problem into two economic margins. Part (i) is a threshold rule across individuals: the agency ranks types by their return per unit of treatment cost and treats those whose return, evaluated at their optimal targeting duration, exceeds the common threshold $\bar{R}$, which measures the shadow value of the budget. Within a given duration, this is the familiar logic of ranking policies by their returns. Part (ii) governs the timing of treatment and adds an option-value consideration: treating a type at duration d only reaches the fraction S_d of that type still unemployed, so postponing treatment sacrifices the opportunity to treat those who exit in the meantime. The agency therefore compares excess returns $R_d^T - \bar{R}$ across durations, weighted by survival. This implies that delaying treatment is optimal only when returns rise sufficiently over the spell to compensate for the shrinking pool of treatable job seekers.

The special case in Proposition 1 illustrates these forces. Without unobserved heterogeneity, a one-period treatment at duration d raises the job-finding probability of the treated by $\tau_d \equiv \tau\,\theta_{\tau,d}$, so the return to targeting can be written as

$$R_d^T(\pi,\tau,u) = \frac{\tau_d}{c} L_{d+1}(\pi,u),$$

where $L_{d+1}(\pi,u) = \sum_{\delta=d+1}^{\infty} u_\delta\, S_{\delta|d+1}(\pi)$ denotes the continuation loss from unemployment from period $d+1$ onward and $S_{\delta|d+1}(\pi)$ is the survival probability to period $\delta \geq d+1$, conditional on remaining unemployed until $d+1$: returns are high where treatment effects and continuation losses are large. Without duration dependence, the continuation loss further reduces to u/π, giving the simplification in the proposition. The return is then the treatment effect relative to its cost, times the avoided unemployment loss, which is larger for *at-risk* job seekers with low job-finding probabilities — so the optimal policy ranks job seekers by the index $u\,\tau/\pi$, which reduces to τ/π when utility losses are identical, trading off risk against responsiveness. The timing implication is equally stark: because the return is the same at every duration while survival declines, delaying treatment shrinks the group that can be reached without raising its return, and all targeting optimally occurs at entry. Away from this benchmark, both forces in part (ii) of the proposition are active: returns may rise or fall over the spell, and the agency weighs this against the declining survival rate. The next proposition characterizes how the returns vary across job seekers and across the spell.

Proposition 2. ***Approximate returns to targeting.*** *Suppose baseline job-finding probabilities and treatment effects evolve geometrically at the individual level, $P_d = (\pi+\varepsilon^\pi)\,\theta_\pi^d$ and $\tau_d = (\tau+\varepsilon^\tau)\,\theta_\tau^d$, and that utility losses, which may differ across job seekers, evolve geometrically, $u_d = u\,\theta_u^d$. For type (π,τ), let $\mathbb{E}_d[\,\cdot\,]$ denote the expectation over $(\varepsilon^\pi,\varepsilon^\tau)$ among untreated survivors at duration d, and* *define the average job-finding probability and the average treatment effect among these survivors,* $\bar{P}_d \equiv \mathbb{E}_d[P_d]$*, and* $\bar{\tau}_d \equiv \mathbb{E}_d[\tau_d]$*, with observed gross changes* $\bar{\theta}_{\pi,d} \equiv \bar{P}_{d+1}/\bar{P}_d$ *and* $\bar{\theta}_{\tau,d} \equiv \bar{\tau}_{d+1}/\bar{\tau}_d$*.*[1] *One-period targeting returns* R_d^T *(defined in equation 3) are approximately proportional to average treatment effects and utility losses, and inversely proportional to average job-finding probabilities, among the surviving unemployed:*

(i) across types. For two types with initial characteristics (π, τ, u) *and* (π', τ', u') *at the same duration* d*, with survivor averages* $(\bar{P}_d, \bar{\tau}_d)$ *and* $(\bar{P}'_d, \bar{\tau}'_d)$*:*

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

(ii) across durations. For the same type across consecutive durations:

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

Part (i) holds in a neighborhood of $\theta_\pi = \theta_u = 1$*. The ratio form underlying part (ii),* $R_{d+1}^T/R_d^T \approx \bar{\theta}_{\tau,d}\big(1 + \theta_u - \bar{\theta}_{\pi,d}\big)$*, is exact in* θ_τ*; the displayed difference form additionally drops the cross term* $(\bar{\theta}_{\tau,d} - 1)(\theta_u - \bar{\theta}_{\pi,d})$ *and thus also requires* θ_τ *close to one. Without unobserved heterogeneity,* $\bar{P}_d = \pi\,\theta_\pi^d$*,* $\bar{\tau}_d = \tau\,\theta_\tau^d$*,* $\bar{\theta}_{\pi,d} = \theta_\pi$*, and* $\bar{\theta}_{\tau,d} = \theta_\tau$*, so that parts (i) and (ii) hold with individual characteristics in place of survivor averages.*

Proof. See Appendix A.1.

Proposition 2 expresses the returns to targeting in terms of objects that are directly estimable at each duration: the average treatment effect $\bar{\tau}_d$ and the average job-finding probability $\bar{P}_d$ among job seekers of a given observable type who are still unemployed at duration d. Part (i) shows that the cross-sectional logic of Proposition 1 carries over to these survivor averages: returns are approximately proportional to average treatment effects and to the utility losses from unemployment, and inversely proportional to average job-finding probabilities. Unobserved heterogeneity within observable types matters for the returns to targeting only through the survivor averages themselves, up to terms that are second order in its dispersion.

Part (ii) characterizes how the returns evolve over the unemployment spell: returns rise or fall according to whether the observed treatment effects among survivors decline more or less quickly than their observed job-finding probabilities, net of the depreciation of utility losses. Moreover, as we show in the proof, the observed decline in job finding compounds two distinct forces through an exact identity, $\bar{\theta}_{\pi,d} = \theta_\pi\big[1 - \mathrm{Var}_d(P_d)/(\bar{P}_d(1 - \bar{P}_d))\big]$, where Var_d denotes the variance among survivors at duration d: true duration dependence (θ_π) and dynamic selection (the variance term), which arises because job seekers with high job-finding probabilities exit unemployment first, leaving behind an increasingly negatively selected pool. Importantly, the two forces matter for the returns to

[1] Continuation losses are interpreted as truncated at a finite maximum spell length: for $\theta_\pi < 1$, individual job-finding probabilities vanish geometrically over the spell and a positive fraction of job seekers would never exit unemployment, so that without discounting the untruncated continuation loss would be infinite. We further assume that $\pi + \varepsilon^\pi$ is bounded away from zero and one.

targeting only through their combined — and directly estimable — effect on the survivor averages, so the agency need not disentangle them. If observed treatment effects fall over the spell while observed job finding remains stable, returns are highest early in the spell. Conversely, if treatment effects are stable among survivors but observed job finding deteriorates, returns increase over 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 job seekers would exit unemployment even without treatment, which reduces the value of intervening at short durations.

Importantly, Proposition 2 characterizes the individual returns to targeting, not the optimal policy. The optimal policy is characterized in Proposition 1 and, in the absence of duration dependence, prescribes targeting at the start of the spell, since delaying treatment forgoes the opportunity to treat job seekers who exit in the meantime. Targeting later in the spell can therefore only be optimal if the returns to targeting increase over the spell. Furthermore, the characterization of the returns in Proposition 2 is based on approximations. In Section 7, we therefore complement the analytical results with a full quantitative evaluation of the optimal targeting policy in a calibrated version of the model.

3 Context and Data

This section describes the data and institutional details relevant for our empirical analysis. 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].

3.1 Institutional Setting

Caseworkers and Service Line Operators. The primary point of contact between the VDAB and a job seeker 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 benefit sanctions. For job seekers who have not been assigned to a caseworker, communication with the VDAB occurs through the 'service line', the VDAB's call center. Operators 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.

Assignment to caseworkers. Our empirical analysis leverages three sources of variation in the assignment to caseworkers. We describe the assignment procedure here in detail and provide a schematic overview in Figure 1.[2] Upon registration with the VDAB, job seekers receive an introductory 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. This is the first source of variation we exploit in what we will refer to as the *'Heterogeneity RD'*.

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.[3] After four weeks of unemployment, all 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 an operator 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.[4] We exploit this variation in our *'Baseline RDD'*. 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. We exploit this variation in a third RDD, which we refer to as the *'Dynamic RD'*. 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

[2]The described process was in place from October 2018 until May 2023. 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.

[3]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.

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

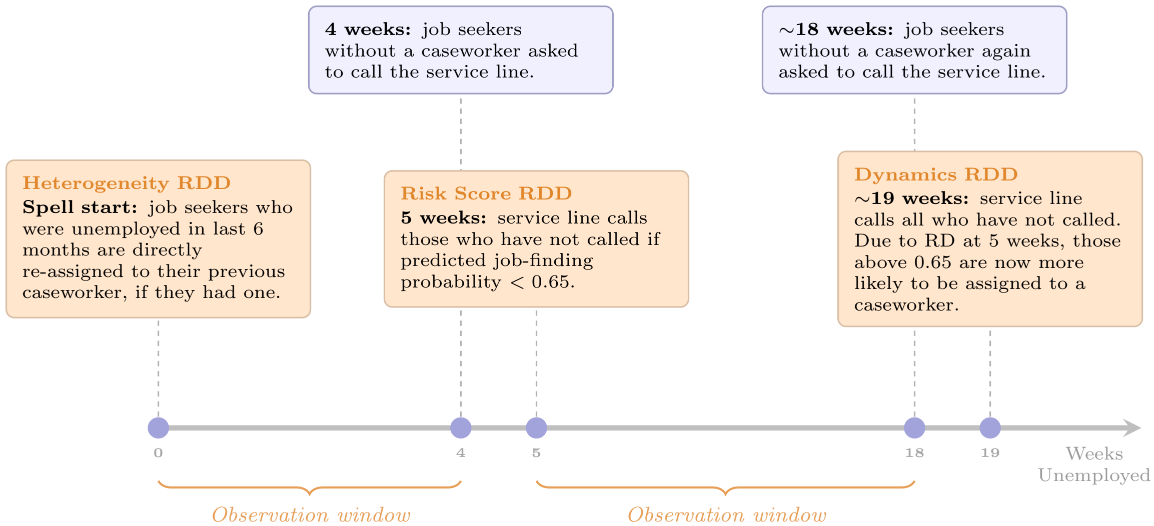

Figure 1: Assignment to Caseworkers

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 are updated automatically on a daily basis. See Ernst et al. [2024] for a more detailed description of the prediction model and the process of assigning job seekers to counseling.

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. To avoid the peak of the COVID crisis, we check robustness for unemployment spells in the first half year and the last half year of our sample. 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), occupational experience and occupational preferences.

Risk scores. We have data on the 6-month job-finding probabilities predicted by the VDAB for job seekers in our sample at different durations. Panel (a) of Figure 2a illustrates the dispersion in these risk scores, plotting the distribution of predicted job-finding probabilities for all job seekers in the second week of the spell. The 90th percentile faces an 80 percent chance of finding a job in the next 6 months, while for the 10th percentile this is only 40 percent. Panel (b) of Figure 2b illustrates the accuracy of the prediction model, showing the average 6-month job finding rate nicely tracks the

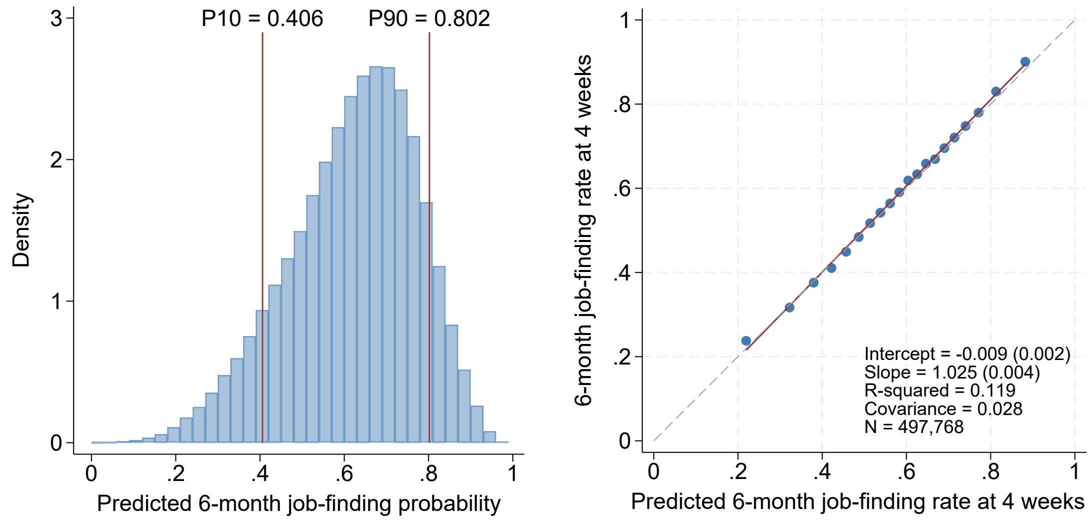

(a) Distribution of Predictions

(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.

predicted 6-month job find rate, also at the top and the bottom of the distribution.[5] Our empirical relies on the risk scores that were provided to the service line and used to profile job seekers.[6]

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 or in subsidized work and training programs lasting 1-6 months.[7] 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 check the robustness of our results to always counting temporary and subsidized work as employment as well as never counting them as employment.

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.

[5]We consider job-finding rates at 4 weeks of unemployment. In Appendix B.2, 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.

[6]Predictions from 3 October 2020 to 20 October 2022 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.

[7]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.

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.[8]

4 Baseline Design

We first turn to the discontinuity at a risk score of 0.65 in week 5 of the unemployment spell, which plays a central role in the VDAB's targeting strategy for assigning job seekers to caseworkers. This discontinuity provides the cleanest and most powerful source of quasi-experimental variation for estimating the effects of job search counseling in the Flemish context. It also serves as a benchmark for our alternative RD designs, which we use below to assess heterogeneity in treatment effects by risk score and dynamics in treatment effects over the unemployment spell.

4.1 Empirical Strategy

After five weeks of unemployment, job seekers with a predicted 6-month job-finding rate below 0.65 are prioritized for contact by the service line. During the resulting interview, the service line operator decides whether the job seeker should be assigned to a caseworker. The threshold therefore generates a discontinuous increase both in the probability of interacting with the service line and in the probability of subsequent assignment to a caseworker.

We employ an 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. The sample is restricted to job seekers who were invited to call the service line after four weeks of unemployment but did not call within a week and were therefore added to the service line's call list.

The service line begins calling job seekers below the cut-off after five weeks of unemployment. Because the next round of service-line interviews begins approximately three months later, we focus on weeks 6–17 of the unemployment spell. This observation window captures the effects of the initial

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

targeting decision while avoiding contamination from later interventions. Our baseline outcome is the number of days worked per week during this 12-week period. This measure captures both the probability of finding employment and the speed with which employment is obtained. We also consider two treatment measures in the fuzzy RD analysis: whether the job seeker had contact with the service line and whether the job seeker had contact with a caseworker during weeks 6–17 of unemployment.

For our baseline specification, we include observations within a bandwidth of 0.15 on either side of the cut-off and apply a triangular kernel, thereby assigning greater weight to observations closer to the threshold. Estimates are robust to using a uniform kernel and a wide range of bandwidths, as shown in Appendix Figure C.2.

As risk scores are generated automatically by an algorithm, strategic manipulation of the running variable is unlikely to be a concern. Nevertheless, discontinuities could arise mechanically if the prediction algorithm assigns scores in a non-smooth manner around the threshold. Appendix Figure C.1 shows no evidence of bunching in the distribution of risk scores at the cut-off. We further verify in Appendix Figure C.3 that predetermined covariates are balanced around the threshold.

4.2 Results

The discontinuity identifies the effect of being prioritized for outreach by the service line. Since service line contact also affects subsequent assignment to caseworkers, the RD estimates the effect of a bundled intervention consisting of service-line counseling and, for a subset of job seekers, caseworker support. We therefore report reduced-form estimates as well as two fuzzy RD estimates that scale the reduced form by the corresponding first stages.

Panels (a) and (b) of Figure 3 plot the proportions of job seekers around the cut-off who had contact with the service line and with a caseworker, respectively. Both panels reveal sharp discontinuities at the 0.65 threshold. The probability of contact with the service line increases by 31 percentage points (s.e. 0.5), from a baseline of approximately 5 percent just above the threshold. This translates into a 9 percentage point increase (s.e. 0.6) in the probability of contact with a caseworker, relative to a baseline of just under 20 percent. The relationship between the risk score and both contact measures is close to linear on either side of the threshold, contributing to the robustness of the estimated first stages. While we aggregate contacts over the full observation window of weeks 6–17, we explore the timing of these interactions in detail in Section 6.

Panel (c) of Figure 3 plots the number of days worked per week during weeks 6–17 of unemployment. As expected, employment outcomes improve steadily with the predicted job-finding probability. More importantly, we observe a clear discontinuity at the 0.65 threshold, with individuals just below the cut-off working more days on average than those just above it. The reduced-form estimate implies that being prioritized for outreach by the service line increases days worked per week by 0.12 (s.e. 0.034).

Combining the reduced-form estimate with the first-stage estimates yields two fuzzy RD estimates. The first, reported as $\hat{\beta}_{IV}$, uses the discontinuity as an instrument for contact with the service line. Under the standard fuzzy RD assumptions, this estimate identifies the local average

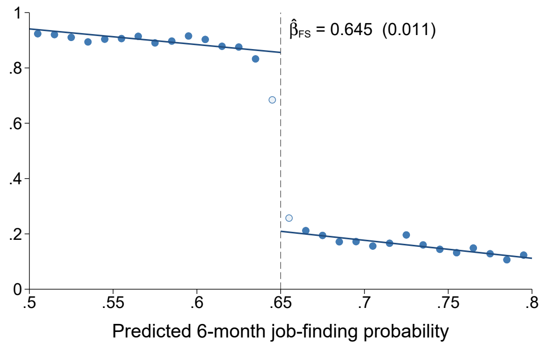

(a) First Stage 1: Contact with the Service Line

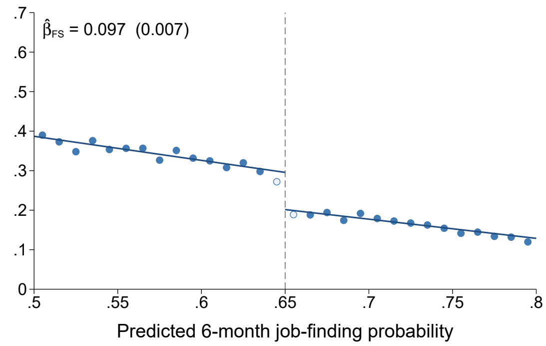

(b) First Stage 2: Contact with a Caseworker

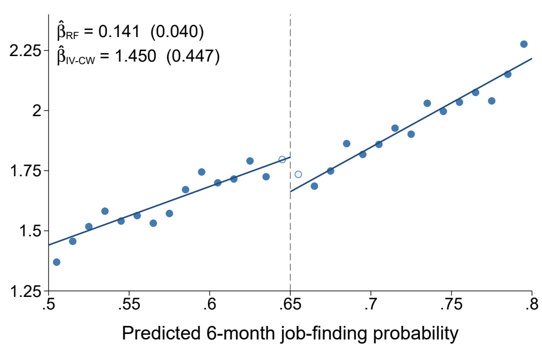

(c) Days Worked per Week

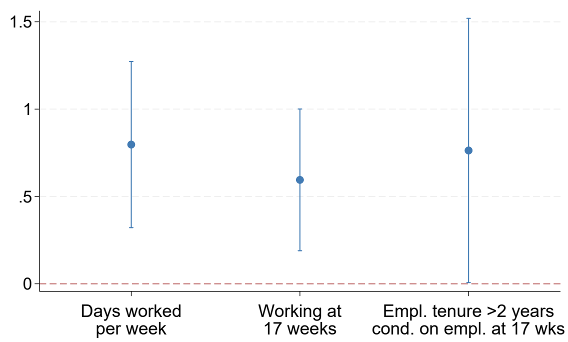

(d) Other Employment Outcomes

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. Panel (d) plots the effects of caseworkers on three different employment outcomes estimated using the fuzzy risk score RD. Estimates are divided by the threshold intercept of the local linear fit on the right side of the cut-off. The first outcome is the baseline outcome used in panels (a)–(c). The second is a binary indicator of 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses. We report 95% confidence intervals.

effect of service-line contact for individuals whose contact status is affected by the threshold [Hahn et al., 2001]. We estimate that service-line contact increases days worked per week by 0.37 (s.e. 0.11).

The second estimate, reported as $\hat{\beta}_{IV-CW}$, uses the discontinuity as an instrument for contact with a caseworker. The corresponding estimate implies an increase of 1.33 days worked per week (s.e. 0.40), or an 80 percent increase relative to baseline. This estimate should not be interpreted as the causal effect of caseworker contact alone. Service-line interactions may themselves affect job finding, anticipation of future meetings may alter job-search behaviour before actual caseworker contact occurs, and assignment decisions by service-line operators may affect the composition of job seekers meeting caseworkers around the threshold. Nevertheless, $\hat{\beta}_{IV-CW}$ provides a useful benchmark for comparison with the alternative RD designs introduced below.

Additional Outcomes The positive employment effects are not unique to our baseline measure of days worked per week. Panel (d) of Figure 3 shows that job seekers just below the threshold are also more likely to be employed after 17 weeks, the end of our observation window.

An important concern is whether the increase in employment reflects job seekers accepting lower-quality jobs in order to leave unemployment more quickly. To examine this possibility, we consider the duration of subsequent employment spells. If anything, job seekers just below the threshold appear to enter longer-lasting employment relationships than those just above, as reported in panel (d) of Figure 3. The corresponding RD plot is provided in Appendix Figure A11. These findings are consistent with job search counseling improving not only the speed with which job seekers find employment, but also the quality or stability of the resulting job matches. However, this pattern may also reflect treatment effect heterogeneity. If job seekers who are ex ante more likely to find stable employment are also more responsive to treatment, the treated group will mechanically exhibit more stable employment conditional on being employed at 17 weeks.

Finally, we examine a number of intermediate outcomes that may help shed light on the mechanisms underlying the employment effects (see Appendix Figure C.4). First, job seekers below the threshold are substantially more likely to log into the VDAB job-search platform as well as report job applications. These patterns are consistent with the intervention increasing search activity, either through counseling provided by the service line or through the monitoring and guidance associated with subsequent caseworker contact. Second, we find evidence that treated job seekers adjust their occupational search preferences, suggesting that the intervention affects not only search intensity but also search direction.

The results in this section establish that job search counseling can have substantial effects on employment outcomes. The key question for targeting, however, is not whether the intervention works on average, but for whom and at what point in the unemployment spell it is most effective. We turn to these questions next.

5 Heterogeneity in Effects

While the estimated effects in our baseline RD are substantial, they may mask important heterogeneity in responsiveness to job search counseling. Understanding such heterogeneity is central to evaluating the VDAB's targeting strategy, which assigns treatment based on predicted job-finding prospects. However, the baseline RD does not allow us to estimate treatment effects separately by risk score, as risk score is the running variable. We therefore introduce a second research design that exploits a different source of quasi-random variation and allows us to study how responsiveness varies across the risk distribution.

5.1 Empirical Strategy

We exploit a separate discontinuity that arises when job seekers re-enter unemployment. Individuals who become unemployed again within six months of a previous spell are automatically reassigned to their former caseworker if they had one during the earlier spell. Individuals returning after more than six months are not. This reassignment rule creates a discontinuity in caseworker contact around the six-month threshold that we use to estimate employment effects. We refer to this empirical strategy as the "heterogeneity RD."

Unlike our baseline RD, this design allows us to estimate treatment effects separately by predicted risk score. We therefore split job seekers into three groups based on their predicted probability of finding a job within six months, generated on day 8 of unemployment: low-risk ($P_i < 0.6$), medium-risk ($P_i \in [0.6, 0.7)$, and high-risk ($P_i \geq 0.7$).[9] The middle group is chosen to correspond most closely to the population around the 0.65 threshold in the baseline RD.

For each risk group, we estimate the following local linear RD specification:

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

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. The sample is restricted to job seekers who had contact with a caseworker in their previous unemployment spell.

We focus on employment outcomes during the first four weeks of unemployment, as reassignment occurs immediately whereas the standard VDAB intervention begins only after four weeks. We estimate local linear specifications within a bandwidth of twenty weeks around the threshold and exclude observations in a narrow donut around six months to account for substantial bunching at common employment durations. Appendix D.1 documents clear bunching at three, six, nine and twelve months, suggesting that these durations serve as reference points for employers and employees. Appendix D.3 shows that our results are robust to alternative bandwidths, kernels, and donut specifications.

[9] We use the prediction generated on day 8 of unemployment, as earlier predictions are considered less reliable. Predictions on day 8 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 and therefore differ little from those who find work shortly after day 8.

A remaining concern is that individuals returning to unemployment within six months differ systematically from those returning later. Reassuringly, outside the bunching region the density evolves smoothly around the threshold and observable characteristics are balanced on either side of the cutoff. Moreover, because reassignment only applies to job seekers who previously interacted with a caseworker, we can construct placebo estimates for otherwise similar individuals who are not eligible for reassignment. We discuss these results below.

The heterogeneity RD differs from the baseline RD both in sample and treatment. The baseline RD focuses on individuals around the risk-score threshold who do not contact the VDAB when instructed to do so, whereas the heterogeneity RD focuses on job seekers who recently experienced a previous unemployment spell. In addition, treatment in the baseline RD occurs later in the spell and combines contact with service-line mediators and caseworkers, whereas the heterogeneity RD captures reassignment to a previous caseworker. We therefore view the heterogeneity RDD as a complementary source of evidence on treatment effect heterogeneity rather than providing directly comparable estimates.

5.2 Results

We define the first stage as contact with a caseworker during the first four weeks of unemployment. Appendix D.4 shows strong first stages for all three risk groups. The discontinuity increases caseworker contact by roughly 14–19 percentage points, with the largest increase among job seekers with the lowest predicted job-finding probabilities.

Panels (a)-(c) of Figure 4 plot employment outcomes around the six-month reassignment threshold separately by risk group. A striking pattern emerges. For job seekers with the lowest predicted job-finding probabilities – those most at risk of long-term unemployment and most likely to be targeted under the VDAB's standard approach – employment outcomes evolve smoothly through the threshold. The estimated effect on days worked is essentially zero (.002; s.e. .034).

In contrast, we find large and statistically significant employment effects for the two groups with higher predicted job-finding probabilities. The estimated effect equals .173 (s.e. .049) for job seekers with $P_i \in [0.6, 0.7)$ and rises further to .264 (s.e. .055) for those with $P_i \in [0.7, 1]$. We reject equality of treatment effects between the lowest and highest risk groups with a p-value below 0.001.

Panel d reports the corresponding fuzzy RD estimates, $\hat{\beta}_{IV-CW}$, using the discontinuity as an instrument for caseworker contact. The IV estimates reinforce the monotonic relationship between predicted job-finding prospects and responsiveness to counseling. For job seekers around the .65 threshold, the estimated effect closely matches the baseline RD estimate. More importantly, treatment effects decline sharply as predicted risk of long-term unemployment increases.

As mentioned, the results are robust to different RD specifications, as shown in Appendix D.3. Moreover, while the heterogeneity RD is restricted to a four-week outcome window, Appendix D.7 shows that the same monotonic pattern persists when outcomes are measured over sixteen weeks.

Taken together, these results reveal a striking negative relationship between unemployment risk and responsiveness to counseling. The job seekers whom the VDAB most actively targets exhibit little or no employment response, whereas those at substantially lower risk respond strongly. The

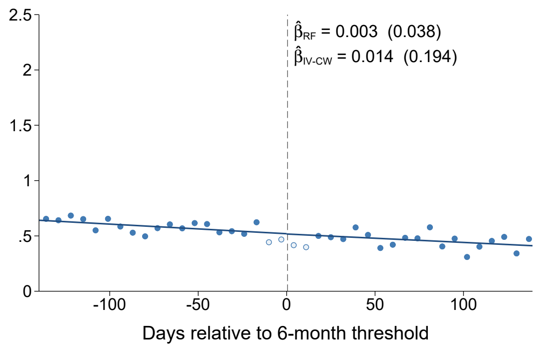

(a) Days Worked per Week ($P \in [0, 0.6)$)

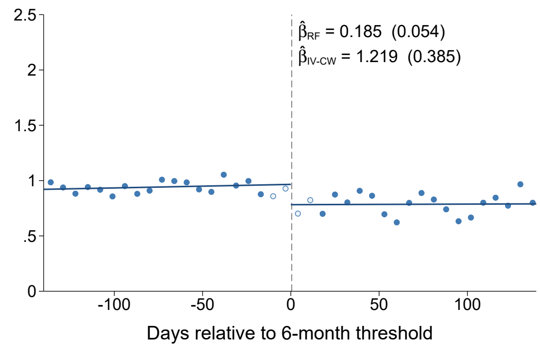

(b) Days Worked per Week ($P \in [0.6, 0.7)$)

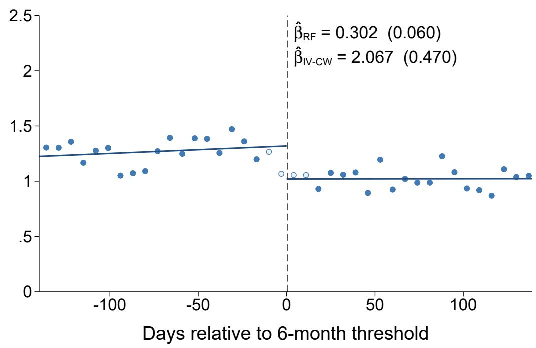

(c) Days Worked per Week ($P \in [0.7, 1]$)

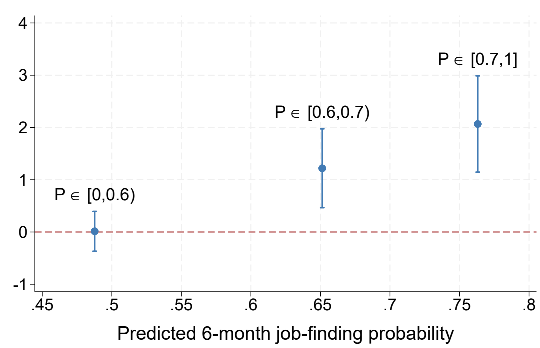

(d) Effect on Days Worked per Week (IV Estimates)

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. 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 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 hypotheses that the IV estimate from panel (a) is equal to the estimates from panels (b) and (c) with p-values of 0.005 and 0.000, respectively. Panel (d) plots, for each group, the IV RD estimate against the mean of P. We report 95% confidence intervals. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

trade-off between targeting individuals who are most at risk and targeting those who are most responsive therefore appears to be first-order in this setting. Before interpreting these results as evidence of genuine heterogeneity in responsiveness, we consider two alternative explanations.

Heterogeneity in Treatments A natural alternative explanation for the observed heterogeneity is that different job seekers receive different interventions. In particular, high-risk job seekers may be more likely to be referred to intensive training programs. Such programs often involve lock-in effects, reducing job search and employment in the short run while potentially improving longer-run outcomes. If so, the weaker short-run employment effects for high-risk job seekers would reflect differences in treatment content rather than differences in responsiveness to counseling.

However, we find little evidence that caseworkers systematically assign different interventions across risk groups. Appendix D.9 examines a broad range of tasks and programs and reveals remarkably similar treatment patterns. As discussed before, the only intervention that consistently increases at the discontinuity is the requirement to apply for jobs, and this effect is present for all three risk groups. We therefore find little support for the view that the observed heterogeneity is driven by differences in treatment content. Instead, it appears that the individuals most likely to be targeted by the VDAB are also the least responsive to the intervention they receive.

Algorithmic Bias An additional concern is that predicted risk may mechanically correlate with treatment effects because the risk score is trained on outcomes that themselves reflect prior treatment. Put differently, the VDAB algorithm predicts job-finding probabilities using historical data generated under an existing targeting regime. If treatment affects job finding, then the algorithm does not simply learn who would find work in the absence of intervention; it also learns who benefited from intervention in the past. The relevance of this concern becomes clear when considering the relationship between predicted job-finding probabilities P_i and treatment effects τ_i.

To see this, suppose $P_i = \pi_i + \tau_i T_i$ and thus:

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

Hence, in the presence of treatment T_i, the algorithm learns a combination of baseline job-finding prospects and treatment responsiveness rather than the former alone. Under random assignment, we obtain $Cov(\tau_i T_i, \tau_i) = Pr(T_i = 1)Var(\tau_i)$. In this case, even if untreated job-finding prospects π_i are unrelated to treatment effects, observed job-finding probabilities P_i will be positively correlated with treatment effects.

This observation has important implications for targeting. Ideally, targeting would distinguish between untreated job-finding prospects π_i and treatment effects τ_i. However, when treatment is assigned on the basis of predicted outcomes P_i instead, individuals with low treatment effects τ_i may be adversely selected into treatment.[10] The key question, however, is whether this mechanism is quantitatively important. We return to this issue using the calibrated model in Section 7 and

[10] In practice, π_i is not directly observed. Potential solutions include estimating risk scores using only untreated individuals or combining predictions of job-finding probabilities and treatment effects to recover untreated prospects more accurately.

find that, while the mechanism is important in theory, its quantitative contribution is modest in our setting.

5.3 Other Observable Heterogeneity

The heterogeneity by risk score reveals an important trade-off between targeting job seekers who are most at risk and those who are most responsive to intervention. However, risk scores represent only one dimension of observable heterogeneity. More generally, efficient targeting requires identifying which observable characteristics predict treatment effects and whether these predictions improve upon targeting based on risk scores alone. We therefore adopt a data-driven approach and estimate conditional average treatment effects (CATEs) using an instrumental forest algorithm following Athey et al. [2019]. This allows us both to predict treatment effects from observables and to identify the characteristics most important for explaining treatment effect heterogeneity.

We estimate the instrumental forest using the heterogeneity RD. Since current implementations do not accommodate the local-linear structure of the RD, we restrict attention to observations within ten weeks of the six-month cutoff and use caseworker contact as the treatment variable and assignment below the cutoff as the instrument.

Panel (a) of Figure 5 plots the distribution of predicted CATEs. The distribution reveals substantial predictable heterogeneity in treatment effects. Indeed, the relative dispersion in predicted treatment effects is considerably larger than the relative dispersion in risk scores. This distinction is important because Proposition 2 implies that targeting depends on proportional differences: a job seeker with a sufficiently larger treatment effect should be prioritized even if their job-finding probability is correspondingly higher.

Panel (b) of Figure 5 assesses both the internal and external validity of the predicted CATEs. For internal validity, we estimate caseworker effects using the heterogeneity RD separately for each CATE tercile. We find that the estimated treatment effect is significantly larger in the highest tercile, confirming that the CATE estimates capture meaningful treatment effect heterogeneity. Note that, if anything, the CATE estimates underestimate the heterogeneity in RD estimates by risk score. As shown in Appendix, these differences are primarily explained by the inability of the instrumental forest procedure to accommodate linear trends in the running variable

The stronger test is external validity. The CATEs are estimated using the heterogeneity RD, which captures the short-run effects of reassigning recently unemployed job seekers to a previous caseworker. The baseline RD instead captures a different intervention, occurring later in unemployment, for a different population. It is therefore far from obvious that the same characteristics should predict treatment effects across the two settings. Nevertheless, we find that CATE terciles estimated from the heterogeneity RD also predict treatment effects in the baseline RD. We view this result as evidence that the heterogeneity uncovered by the heterogeneity RD reflects stable differences in responsiveness rather than idiosyncrasies of a particular treatment environment.

In addition to quantifying the heterogeneity, the instrumental forest estimation allows us to analyze the observables that drive the heterogeneity. In the Appendix, we report the "variable importance" for different observables, which is the number of times the forest split on a variable

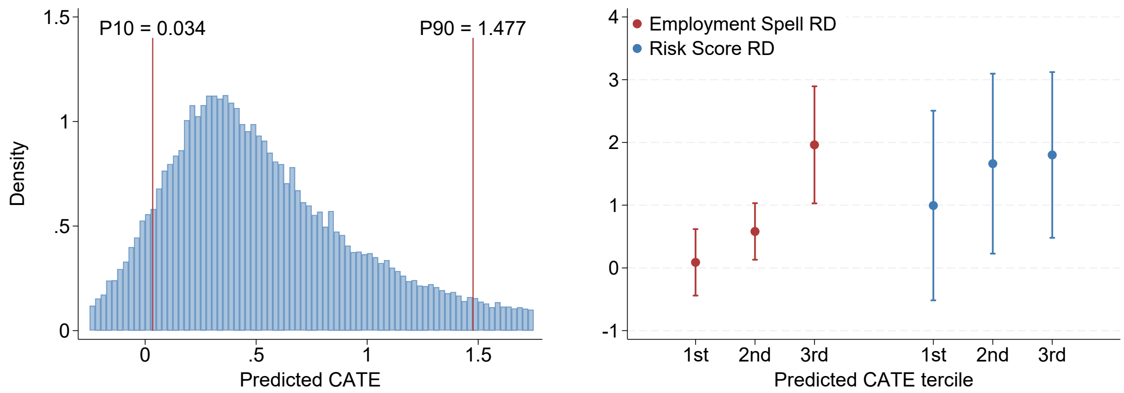

(a) Distribution of CATEs

(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 heterogeneity RD. We reject the null hypothesis that effects on the 1st and 3rd terciles are equal with a p-value of 0.001. 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 rejecting the null hypothesis that the effects on the 1st and 3rd terciles are equal is 0.422. The tercile boundaries used are the same for both sets of estimates. We display 95% confidence intervals for the estimated effects.

weighted by the depth at which the split occurred. Interestingly, the risk score is the single most important predictor of treatment effects, confirming the earlier evidence that responsiveness is strongly related to predicted unemployment risk. However, the forest also identifies demographic characteristics, unemployment history, and occupation preferences as important predictors. The fact that these variables continue to matter even after conditioning on the risk score suggests that substantial heterogeneity remains within risk-score groups. In other words, the risk score captures an important but incomplete dimension of treatment effect heterogeneity.

To illustrate this further, we compare individuals with high and low risk scores and high and low predicted treatment effects. Some characteristics affect unemployment risk and treatment responsiveness in the same direction, whereas others generate a trade-off between the two. For example, job seekers of non-Belgian origin appear both more likely to experience long-term unemployment and more responsive to intervention. Such characteristics strengthen the case for targeting, as they simultaneously increase need and expected returns from treatment. Other characteristics primarily affect responsiveness while having little impact on unemployment risk. For instance, female job seekers exhibit higher predicted treatment effects despite having similar risk scores. A targeting strategy based solely on unemployment risk would therefore overlook dimensions of heterogeneity that are highly relevant for efficient treatment assignment. We return to these implications when comparing alternative targeting strategies in the calibrated model below.

6 Duration Dependence in Effects

The conceptual framework highlights that the timing of treatment depends on potentially opposing forces underlying its return. On the one hand, treatment should be provided early if treatment effects decline with unemployment duration. On the other hand, treatment becomes more valuable later in the spell if job-finding prospects deteriorate and the consequences of continued unemployment become more severe. Whether employment services should intervene early or late therefore depends on the relative dynamics of treatment effects and job-finding probabilities.

In this section, we estimate both objects. We first introduce an alternative RD design that allows us to compare treatment effects for similar job seekers treated at different points in the unemployment spell. We then quantify how dynamic selection changes both treatment effects and job-finding probabilities among the surviving unemployed population.

6.1 Empirical Design

To study the dynamics of treatment effects, we exploit a second round of interventions that begins around week 18 of the unemployment spell. At that point, all job seekers who remain unemployed and have not previously been assigned to a caseworker are invited again to contact the service line. In week 19, the service line starts contacting those who did not respond. This second round generates a useful reversal in treatment assignment relative to the baseline RD.

In the baseline RD, we used the risk score threshold to estimate caseworker effects for individuals with a predicted six-month job-finding probability just below 0.65. Because job seekers below the threshold are more likely to be treated during this first round, those above the threshold are more likely to remain untreated and therefore more likely to receive treatment during the second round, when the risk-score threshold is no longer used.

We exploit this reversal by estimating the same specification as in the baseline RD, but with treatment assignment reversed around the threshold. The running variable remains the predicted six-month job-finding probability initially used by the service line in week 5. We refer to this empirical strategy as the "dynamic RD." Combined with the baseline RD, it allows us to compare treatment effects for job seekers with similar predicted job-finding prospects, but at different points in the unemployment spell.

6.2 Results

Panels (a) and (b) of Figure A37 first verify the reversal in treatment assignment that underlies the dynamic RD. Job seekers below the risk-score threshold are substantially more likely to interact with the service line or be assigned to a caseworker during the first treatment round, but substantially less likely to do so during the second round. The differences in contact probabilities are largest in the first few weeks after each treatment round begins and then rapidly disappear, although caseworker contact lags service line contact by up to a few weeks and the difference in contact rates decreases more gradually. The reversal in assignment generates quasi-experimental variation in caseworker contact among surviving job seekers around the same risk-score threshold, but at different unemployment

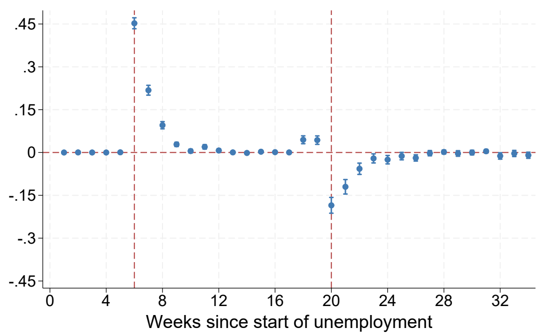

(a) Effect on Contact with Service Line (Weekly)

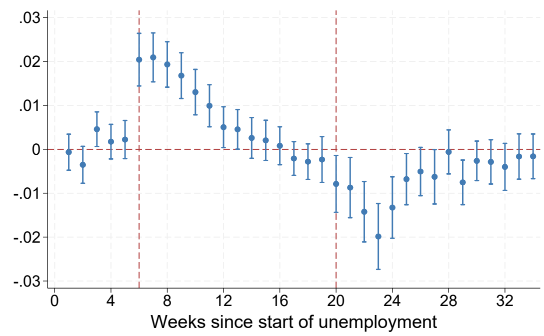

(b) Effect on First Contact with Caseworker (Weekly)

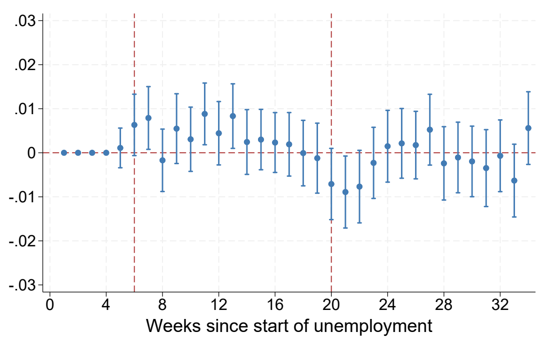

(c) Effect on Entry into Employment (Weekly)

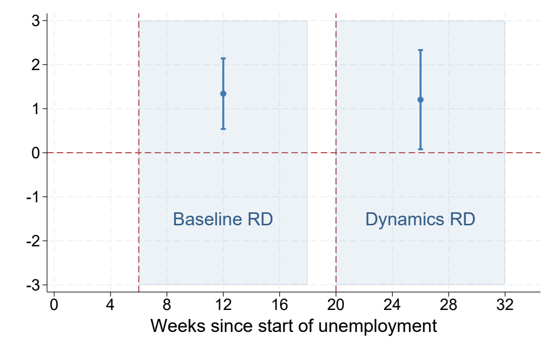

(d) IV Effect on Days Worked per Week (Over 12 Weeks)

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.

durations.

Panel (c) of Figure A37 examines the corresponding employment effects. Employment responds rapidly following both treatment rounds, with effects concentrated in the weeks immediately after the start of the contact round. While the weekly estimates are somewhat noisy, the timing and shape of the employment response does not differ substantially across the two rounds. If anything, the effects appear slightly more concentrated in the second round.

Panel (d) of Figure A37 summarizes these dynamics using IV estimates of the effect of caseworker contact on days worked per week over a twelve-week horizon. We estimate a treatment effect of 1.33 days worked per week in the first treatment round and 1.22 in the second. Appendix E.1 shows the corresponding RD plots for the dynamic RD, which closely mirror the baseline RD estimates. While the confidence intervals remain substantial, the estimates provide little evidence that treatment effects decline as unemployment lengthens.[11]

From the perspective of the conceptual framework, the comparison between the baseline and dynamic RD captures the relevant notion of duration dependence. Both designs identify treatment effects for job seekers with predicted job-finding probabilities around the same threshold, but at different points in the unemployment spell. To the extent that the risk score summarizes underlying job-finding prospects, the comparison therefore isolates how treatment effectiveness evolves with unemployment duration for comparable job seekers. The similarity of the estimated effects suggests little depreciation in treatment effectiveness over the first months of unemployment.

Interpreting the comparison between the baseline and dynamic RD as evidence on *true* duration dependence in treatment effects nevertheless requires ruling out dynamic selection on unobservables. While both designs focus on job seekers with similar predicted job-finding prospects, surviving job seekers in the second treatment round may still differ from those in the first round along dimensions not captured by the risk score. If treatment effects were correlated with such unobserved characteristics, differences between the two estimates could reflect changes in composition rather than genuine duration dependence. In any case, the estimates provide little indication of either substantial duration dependence or strong selection on unobservables. At minimum, neither force appears quantitatively dominant in our setting.

Separately, the institutional structure of the dynamic RD raises two concerns for interpreting the second-round estimates themselves. Both arise because treatment during the first round affects the composition of job seekers who remain eligible for treatment during the second round.

A first concern is differential survival into the second treatment round. Since treatment increases job finding in the first round, job seekers below the threshold are less likely to remain unemployed until week 20. If treatment effects are positively correlated with unobserved job-finding prospects, the surviving population above the threshold could therefore exhibit higher job-finding rates even in the absence of the second-round intervention. We view this concern as limited. The employment effects in Figure 6c emerge precisely when the second treatment round begins, while estimates in weeks 18 and 19 are close to zero. This timing is difficult to reconcile with a purely compositional explanation.

A second concern is that surviving untreated job seekers below and above the threshold may

[11] Appendix E.2 further gauges robustness of this conclusion to the bandwidth and weighting choices.

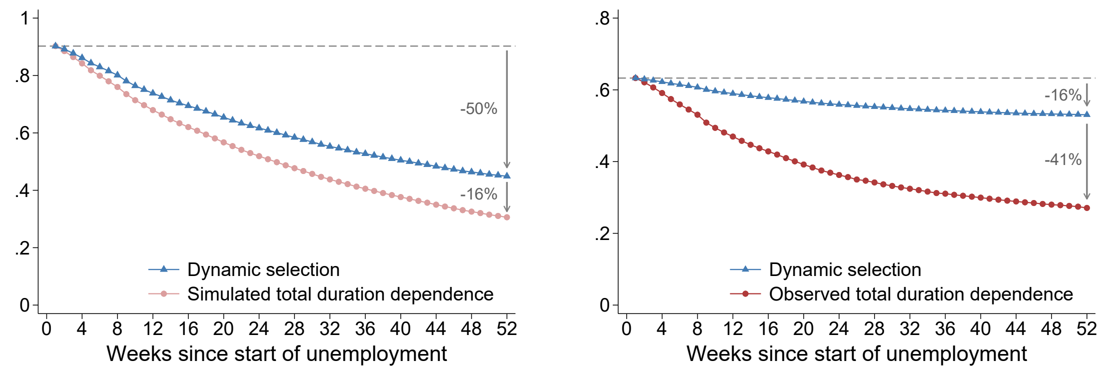

(a) Treatment Effects

(b) Job Finding

Figure 7: Duration Dependence in Treatment Effects and Job Finding

Note: Panel (a) simulates the evolution of the average treatment effect of caseworkers on days worked at different unemployment durations due solely to dynamic selection using heterogeneity in treatment effects and observed survival rates by predicted risk quintile. Treatment effects for each quintile are estimated using the heterogeneity RD. Total duration dependence is then estimated by calculating true duration dependence using the baseline and dynamic RD estimates and assuming it to be homogeneous within the sample and affecting treatment effects at a constant geometric rate with duration. Panel (b) plots the empirical 6-month job-finding rate for the surviving population at each duration (observed total duration dependence) against 6-month job-finding rate predicted by VDAB on day 8 of unemployment for the surviving population at each duration (a lower bound on dynamic selection). All predicted rates are multiplied by a constant so that the predicted and empirical rates are equal at a duration of 1 week.

differ systematically because mediators exercise discretion when assigning caseworkers during the first treatment round. In particular, some job seekers below the threshold may be classified as self-reliant and thus remain untreated despite being contacted. If these individuals differ in their responsiveness to treatment, the composition of untreated survivors could differ across sides of the threshold. To assess this possibility, we estimate predicted treatment effects among job seekers who survive untreated until week 18. We find no meaningful discontinuities at the threshold, either in overall predicted treatment effects or in their main observable determinants. This suggests that differential selection into the pool of untreated survivors is unlikely to explain our findings.

6.3 Dynamic Selection

The previous subsection provided little evidence that treatment effects decline substantially with unemployment duration for comparable job seekers. This does not imply, however, that the returns to treatment remain constant over the unemployment spell. As unemployment lengthens, the composition of the unemployed population changes. Job seekers with lower job-finding prospects become increasingly overrepresented among those who remain unemployed and, given the heterogeneity documented in Section 5, these individuals also tend to exhibit lower treatment responsiveness. Dynamic selection therefore affects both job-finding probabilities and treatment effects, making it a central determinant of the optimal timing of interventions.

Panel (a) of Figure 7a simulates the dynamic selection in treatment effects by splitting the sample into predicted risk quintiles and simulating the average treatment effect of the surviving population at different unemployment durations. This assumes treatment effects and survival rates are the same within quintile, but heterogeneous across quintiles.[12] 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. Thus, even in the absence of true duration dependence, average treatment effects decline substantially over the unemployment spell because the remaining unemployed population becomes increasingly negatively selected.

Panel (b) of Figure 7b then investigates 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. In comparison to Sweden as documented by Mueller and Spinnewijn [2025], the decline in job-finding caused by dynamic selection is somewhat smaller, while the overall decline in job-finding is much larger. This suggests that, unlike treatment effects, job-finding rates exhibit substantial duration dependence beyond what can be explained by dynamic selection alone.

Taken together, the results imply a decline in average treatment effects of roughly 65% over the first year of unemployment, compared to a decline in job-finding probabilities of roughly 57%. Through the lens of Proposition 2, these two forces have opposing implications for timing. The decline in treatment responsiveness favors earlier intervention, whereas the decline in job-finding probabilities increases the value of helping those who remain unemployed. On balance, the two observable forces are of comparable magnitude, suggesting only a modest advantage of targeting earlier rather than later in the unemployment spell. In contrast, the strong decline in job-finding probabilities that is not accounted for by observables (about 41%) points toward increasing returns to targeting later in the spell, conditional on risk scores.

The comparison above focuses on average returns to treatment at different unemployment durations. Optimal targeting, however, depends not only on average returns but also on their dispersion and predictability. The distribution of treatment returns evolves over the unemployment spell as dynamic selection changes the composition of the unemployed population. In the next section, we quantify these forces within a calibrated version of the framework.

7 Quantitative Evaluation

The preceding sections estimated the key ingredients for targeting job-search counseling: baseline job-finding prospects, treatment responsiveness, and how both evolve over the unemployment spell. This section combines these ingredients in a calibrated version of the framework to quantify the welfare gains from alternative targeting policies. The exercise serves two main purposes. First, we compare targeting policies to evaluate how much is lost by using unemployment risk alone as a targeting rule, and how closely the simple sufficient statistic $R \approx \tau/p$ approximates the full model-based return to treatment. Second, we quantify the timing trade-off, distinguishing between changes

[12] Quintile treatment effects are estimated using the heterogeneity RD. Quintile survival rates are estimated using the average predicted risk score within each quintile.

in the return to treating job seekers who remain unemployed and changes in the share of job seekers who survive to each duration.

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
σ_{ε^π}	Unobserved heterogeneity in job finding	0.037	Mueller and Spinnewijn [2025]
θ^π	Depreciation of job finding	0.918	12-month decline in job finding
α_τ	Intercept in $\tau = \alpha_\tau + \beta_\tau \pi + e_\tau$	-0.045	RDD estimates by risk score
β_τ	Slope in $\tau = \alpha_\tau + \beta_\tau \pi + e_\tau$	0.104	RDD estimates by risk score
σ_{e^τ}	Orthogonal heterogeneity in treatment effects	0.022	Variance of CATEs
θ^τ	Depreciation of treatment effects	1.000	Estimates of risk-score RDDs
$\bar{B}$	Budget constraint	0.530	Share of unemployed targeted by VDAB

7.1 Calibration

We calibrate the model to key moments in the data, as summarized in Table 1. For the job finding we start from the empirical distribution of predicted 6-month job-finding probabilities P_i from the VDAB, drawing a random sample of 50,000 individuals at the start of the spell. We calibrate the model at the monthly frequency and thus map each risk score into a monthly job-finding probability. We refer to this as the baseline job-finding probability, denoted by π_i. [13] We allow for additional unobserved heterogeneity in job finding, σ_{ε^π}, which is calibrated using estimates on the relative contribution of observed and unobserved heterogeneity to job-finding risk in Sweden from Mueller and Spinnewijn [2025]. We assume geometric depreciation, $\theta^\pi_d = (\theta^\pi)^d$, where θ^π is chosen so that the implied decline in job-finding from both the depreciation and the dynamic selection based on observed and unobserved heterogeneity matches the observed decline in job-finding rates over the first 12 months in the Flemish data. We assume no further depreciation in job finding after 12 months.

Turning to the treatment effects, we assume that observable heterogeneity is linear in the VDAB prediction:

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

where α_τ is a constant, β_τ captures the relation between treatment effects and risk scores, and e_τ represents heterogeneity orthogonal to risk scores. We assume that treatment increases the job finding probability for one month only and use the estimated reduced-form effects on the probability of starting work within four weeks from the heterogeneity RDD. We calibrate α_τ and β_τ by estimating the 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.[14] In the baseline specification, we assume no additional

[13] We recover π_i from the VDAB risk score P_i using a numerical solver. In the absence of depreciation and unobserved heterogeneity, this mapping would be simply given by $\pi_i = 1 - (1 - P_i)^{1/6}$. The VDAB predictions are based on realized job-finding outcomes under the existing policy environment, while the model object is baseline job finding in the absence of treatment. In the baseline calibration, we abstract from this distinction, but then explicitly relax it in sensitivity analysis.

[14] More precisely, the CATEs reported in Figure 5a are based on days worked. In the calibration, however, we target again the dispersion of CATEs for the probability of starting work within four weeks.

unobserved heterogeneity in treatment effects, $\sigma_{\varepsilon^\tau} = 0$ and no depreciation in treatment effects, consistent with the comparison of RDD estimates at 5 and 19 weeks. We assess robustness to both assumptions.

For evaluating welfare effects we lack direct empirical targets for the utility cost of unemployment u and the treatment cost c. Since our focus is on relative welfare comparisons, we normalize both to one and omit these parameters from the table. We also set $\theta_d^u = (\theta^u)^d = 1$ in the baseline, while exploring sensitivity to this assumption. Finally, we impose a budget constraint consistent with current VDAB policy, allowing treatment of 53 percent of the sample.[15]

7.2 Optimal Targeting and Timing

We first evaluate targeting rules for newly unemployed workers, fixing treatment at the start of the spell, $d = 0$. This exercise isolates the cross-sectional targeting problem: whom should the agency treat among the entering unemployed, abstracting from the timing dimension considered below. We compare three classes of policies, each subject to the same budget constraint:

1. **Optimal Targeting (T_0^{OPT}):** Constrained optimal targeting of individuals with the highest return, $R^{T_0}(\pi, \tau)$.

2. **Index Targeting (T_0^{IDX}):** Targeting based on a geometric weighted average of π and τ,[16]

$$Q(\pi, \tau, \omega) \quad = \quad \left(\frac{1}{\pi}\right)^{\omega} (\tau)^{1-\omega}. \tag{6}$$

3. **Random Targeting (T_0^{RDM}):** Targeting a random sub-sample of unemployed workers.

The index targeting allows us to evaluate simple rules, ranging from targeting on π ($\omega = 1$) to τ ($\omega = 0$) only, and also includes our sufficient statistic approximation when targeting on τ/π ($\omega = .5$). Panel (a) of Figure 8 reports the returns of different targeting policies relative to the random benchmark, which obtains the average returns to intervening. 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. A Shapley decomposition of the variance of the targeting index τ/π, reported in Appendix Table A1, attributes 78 percent to heterogeneity in treatment effects and the remaining 22 percent to heterogeneity in baseline job-finding probabilities. The corresponding panel (b) illustrates the types targeted under the optimal policy, clearly highlighting the underlying trade-off between targeting individuals that are responsive vs. at risk.

Second, the index rule can perform very well and realize basically all gains from the optimally targeted policy. The index τ/π is highly correlated with the exact return to targeting (correlation 0.85), supporting the sufficient-statistic approximation. The simplest rule based on this index provides an excellent approximation, but the index attains its maximum at $\omega = 0.60$. Remember that

[15] We assume that both the distributions of ε^π and ε^τ are uniformly distributed with mean zero. Appendix Figure A1 shows the distribution of π's and τ's as well as the duration dependence in job finding, which matches closely the empirical observed duration dependence.

[16] To be precise, for the purposes of computing the index, we assume that the treatment effects are non-negative.

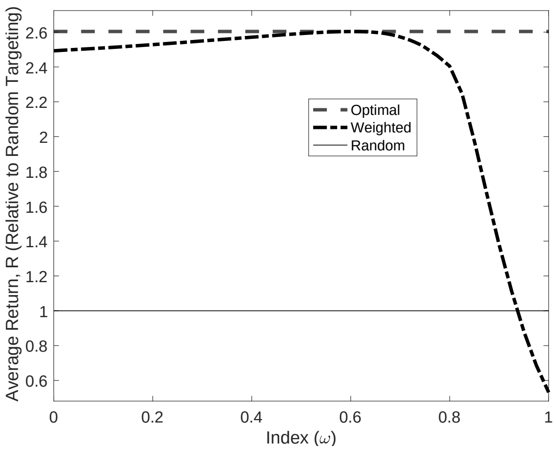

(a) Average Return Relative to Random Targeting

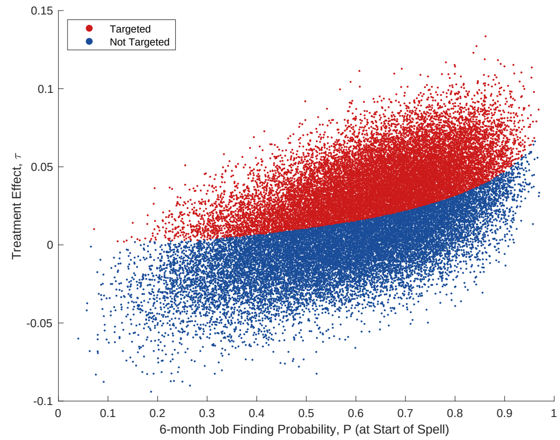

(b) Targeted Types With Optimal Policy

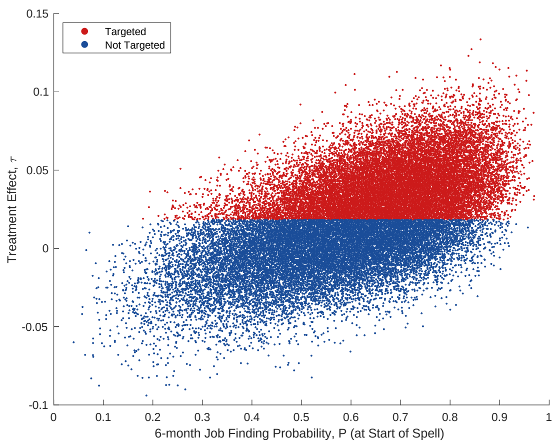

(c) Targeted Types With $\omega = 0$

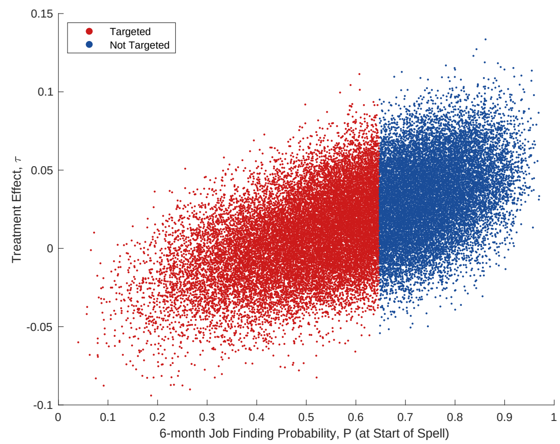

(d) Targeted Types With $\omega = 1$

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.

the optimal weight is exactly $\omega = 0.50$ in a stationary environment without duration dependence in job finding, as confirmed in Table 2.

Third, a policy that targets exclusively on job-finding probabilities ($\omega = 1$) performs poorly and even *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, and shows substantially less overlap between the latter and the optimal policy. 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.

We next turn to the timing dimension. The preceding comparison fixes treatment at the start of the spell. In practice, however, the agency may also decide whether to intervene immediately or wait until job seekers have remained unemployed for some time. This introduces a distinct trade-off. Conditional on still being unemployed, the return to treatment may change with duration because baseline job-finding probabilities and treatment effects evolve over the spell. At the same time, delaying treatment reduces the mass of job seekers who remain available for intervention, since some workers exit unemployment before treatment is delivered.

We distinguish between three timing exercises. First, we evaluate the return to treating a fixed share of the surviving unemployed at each duration. This comparison is conditional on survival: at each duration, the agency ranks only those job seekers who remain unemployed. Figure 9 reports the corresponding welfare gains relative to random targeting at the start of the spell. The return to treating surviving job seekers rises over the spell, as also illustrated in the Figure by the returns to random targeting at different durations d. This reflects the steep decline in baseline job-finding probabilities in the Flemish data, which raises the value of accelerating exit from unemployment, despite adverse selection in treatment responsiveness. This result naturally depends on our baseline assumption that treatment effects do not depreciate at the individual level.[17]

Second, we consider a different exercise in which the agency must spend the same overall budget, but is constrained to treat workers at a common duration d. This comparison incorporates the shrinking size of the survivor pool. In contrast to the conditional-survivor exercise, welfare gains decline sharply over the first months of the spell. Early in the spell, the unemployed pool is large and heterogeneous, allowing the agency to select many high-return workers. As duration increases, many of these workers have already exited unemployment, and the remaining pool is smaller and less favorable for selective targeting. After month 4, fewer than half of the initial unemployed population remains unemployed, so the agency can no longer exhaust the full budget even by treating all survivors and realizes the average return to treating surviving job seekers.

Finally, we allow the planner to choose jointly whom to treat and when to treat them, subject to the same overall budget and the constraint that each individual can be treated at most once. The optimal policy combines the two forces above. Later treatment can be attractive because conditional returns among survivors are high. But reserving treatment for later also risks losing

[17] Appendix Figure A3 shows how the return evolves for a given type.

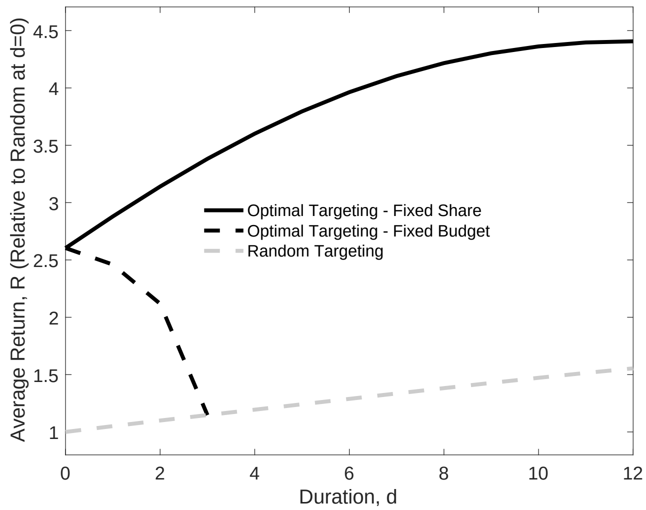

Figure 9: Welfare Gains of Targeting Policies at d, Relative to Random Targeting at $d = 0$

Note: The figure reports the welfare gains from applying three targeting rules: random targeting among survivors, optimal targeting of a fixed share of survivors, and optimal targeting under a fixed budget at unemployment duration d, relative to random targeting at $d = 0$. Welfare gains for the fixed-budget rule are greyed out after month 4, when the optimal policy can no longer exhaust the full budget because too few workers remain unemployed.

high-return individuals who exit unemployment before treatment is delivered. The relevant object is therefore not only the conditional return $R_d(\pi, \tau)$, but the survival-weighted value of assigning treatment to a type at duration d, as characterized in Proposition 1. Quantitatively, the optimal policy assigns a much larger share of treatments at the start of the spell and reserves later treatment only for few selected workers who remain unemployed, as shown in Panel (a) of Figure 10. Panel (b) shows that the workers treated early have higher average returns than those treated later.[18]

Taken together, the calibration clarifies the policy message. Risk scores are informative because they predict baseline job-finding probabilities, but they are not sufficient targeting rules. High-return policies combine information on risk and responsiveness. The timing of treatment depends on two distinct objects: the return to treating job seekers who remain unemployed, and the share of each type that survives to later durations. This explains why returns among survivors can rise with unemployment duration, while the globally optimal policy still treats most high-return job seekers early.

7.3 Algorithmic Bias

We finally use the calibrated model to assess whether the estimated relationship between predicted job-finding prospects and treatment effects could itself reflect the way the risk score is constructed. Because the VDAB score is trained on realized employment outcomes under the existing policy

[18] This is not inconsistent with the rising return to treating the average survivor. The global optimum first uses early treatment slots for the highest-return types. Later treatment is then allocated to the best remaining survivor-types for whom waiting is valuable, but these need not have higher returns than the individuals already treated at the start of the spell.

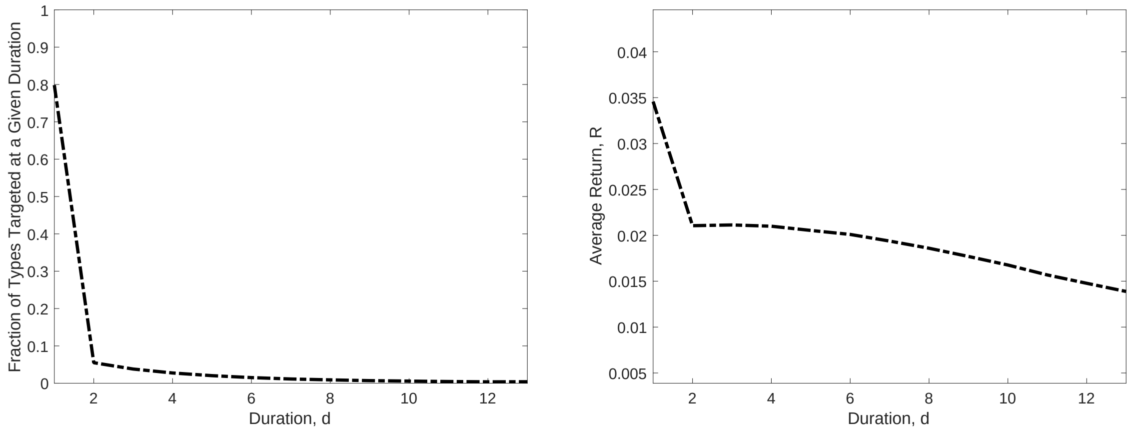

(a) Fraction Targeted at Duration d

(b) Average Return to Targeting of Targeted Types

Figure 10: Optimal Targeting Across Durations

Notes: Based on model simulations with 200,000 individuals.

environment, it may partly incorporate the effects of past treatment. As discussed in Section 5, this can mechanically raise predicted job-finding probabilities for groups with high treatment effects and vice versa, generating a positive relationship between the score and responsiveness.

To quantify this bias, we construct a counterfactual risk score that predicts job finding in the absence of treatment and compare its relationship with treatment effects to that of the observed score. This exercise assumes a specific assignment to treatment underlying the observed scores, as detailed in Appendix A.3. We find that the bias is present but moderate: the regression coefficient of treatment effects on the risk score — the moment targeted in the calibration — falls from 0.73 for the observed score to 0.50 for the treatment-free score. Thus, the observed relationship is amplified by the algorithmic feedback, but not generated by it. This result is robust to different underlying treatment assignments.

The positive relationship between treatment effects and predicted job-finding prospects thus remains strong once the bias is accounted for, so our conclusions are not an artifact of this feedback effect. Note also that the bias itself does not compromise our evaluation of current practice: the VDAB assigns counseling based on the observed risk scores, so both our quasi-experimental estimates and our assessment of score-based targeting apply to the scores actually used for assignment, whether or not these embed the effects of past treatment. Two further remarks are in order. First, the sensitivity analysis in Table 2 already covers the polar case in which the correlation between treatment effects and baseline job-finding prospects is set to zero (column 6) — the case that would apply if the observed relationship were entirely an artifact of algorithmic bias. Removing the correlation naturally improves the performance of risk-based targeting, which then no longer selects against treatment responsiveness, but it leaves the main conclusions intact: risk-only targeting still falls well short of policies that incorporate treatment effects. Second, while an ideal risk score would predict job finding in the absence of treatment, the VDAB does not adjust its risk scores for the effects of past treatment. Such an adjustment would be far from straightforward in practice, as it would require identifying which job seekers' scores are inflated by treatment received during the period on which the algorithm was trained.

7.4 Sensitivity analysis

Table 2 presents a sensitivity analysis with respect to the main calibration parameters. Panel A reports the welfare gains of targeting policies at the start of the unemployment spell. The gains from optimal targeting remain large across all specifications, indicating that the high return to targeting is robust to alternative calibrations. This includes changes in the heterogeneity and/or the dynamics in job-finding probabilities, treatment effects and utility loss (columns 2-6). The relative welfare gains under the optimal policy are somewhat reduced, but remain sizable, when using IV-based rather than reduced-form treatment effects (column 7) or when when the correlation between risk scores and treatment effects is set to zero ($\beta_\tau = 0$ in column 8). However, the gains become even more pronounced when the budget is constrained such that only the top 30 percent of individuals can be treated, as shown in column (9).

Similarly, the simple index rule targeting τ/π realizes most of the targeting gains across the different specification and, in fact, all of them in the specification with no depreciation in job finding (column 2). Moreover, targeting solely based on high treatment effects ($\omega = 0$) comes close to the optimal rule across all calibrations, whereas targeting solely based on low risk scores ($\omega = 1$) delivers only about half of the welfare gains of random targeting in the baseline calibration. Using the IV-based rather than reduced-form treatment effects reduces the level of the welfare gains somewhat but leaves the comparison across targeting policies unaffected (column 7). Only when the correlation between risk scores and treatment effects is set to zero, targeting only on risk scores performs better and, in fact, outperforms the benchmark of random targeting (column 6). Instead, when the budget is restricted to the top 30 percent of individuals (column 9), this policy generates negative welfare effects.

Panels B and C turn to the timing of interventions. Panel B compares the welfare gains of targeting at three months to those of targeting at entry. When the agency targets a fixed share of the surviving unemployed, the welfare gains relative to the cost of treatment are 30 percent higher at three months than at entry in the baseline calibration (first row), and random targeting exhibits a similar, albeit weaker, increase (third row). The cross-calibration patterns confirm the mechanisms behind these rising returns: the advantage of targeting later vanishes in the stationary environment (column 2) and when treatment effects decline over the spell at the same rate as job finding (column 4) — indeed, in both scenarios random targeting performs *worse* at three months than at entry — while it is amplified when the utility loss from unemployment rises over the spell (column 7) and dampened when continuation losses are truncated after two years (column 8). 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.

The picture reverses when the agency must exhaust a fixed budget at a single duration (second row): the welfare gains at three months drop to less than half of those at entry in the baseline calibration, because too few workers remain unemployed for the policy to stay selective. The exception

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

	(1) Baseline	(2) $\sigma_\varepsilon^\psi = 0$ $\theta^\psi = 1$	(3) $\sigma_\varepsilon^\tau > 0$	(4) $\theta^\tau < 1$	(5) $\tau(IV)$	(6) $corr(\psi, \tau) = 0$	(7) $\theta^u > 1$	(8) At $d = 25$: no continuation loss	(9) Target 30%
Panel A. Targeting at $d = 0$									
Optimal Targeting vs. Random	2.60	2.41	2.38	2.60	1.84	1.93	2.72	2.20	3.35
Sufficient Statistics Targeting vs. Random	2.59	2.41	2.35	2.59	1.83	1.93	2.70	2.20	3.30
Targeting High Treatment Effects vs. Random	2.49	2.37	2.26	2.49	1.74	1.85	2.58	2.17	2.87
Targeting Low Risk Scores vs. Random	0.53	0.44	0.66	0.53	0.76	1.45	0.54	0.56	-0.16
Panel B. Targeting at $d > 0$									
Optimal (Fixed Share) at $d = 3$ vs. $d = 0$	1.30	0.99	1.31	1.03	1.29	1.44	1.40	1.16	1.32
Optimal (Fixed Budget) at $d = 3$ vs. $d = 0$	0.44	0.61	0.52	0.33	0.66	0.75	0.46	0.48	0.97
Random at $d = 3$ vs. $d = 0$	1.15	0.82	1.24	0.87	1.22	1.45	1.26	1.06	1.15
Panel C. Targeting Across Durations									
Global Optimal Targeting vs. Random at $d = 0$	2.64	2.41	2.48	2.60	1.89	1.95	2.79	2.22	3.54
Fraction Targeted at $d = 0$	0.80	1.00	0.77	0.97	0.72	0.81	0.74	0.85	0.47

Notes: This table reports the welfare gains of targeting policies under alternative calibration scenarios. Panel A reports the welfare gains of policies applied at $d = 0$, relative to random targeting at $d = 0$, together with the optimal weight in the weighted targeting rule. Panel B reports the ratio of the welfare gains from applying a policy at $d = 3$ to those from applying the same policy at $d = 0$, holding fixed either the share of survivors targeted (fixed share) or the total budget (fixed budget). Panel C reports the welfare gains of the optimal policy across durations, relative to random targeting at $d = 0$, and the fraction of targeted individuals who are treated at $d = 0$. Column (1) shows the baseline calibration. Column (2) combines $\sigma_{\varepsilon^\pi} = 0$ with $\theta^\pi = 1$. Column (3) 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 (4) assumes treatment effects decline at the individual level at the same rate as job finding. Column (5) uses instrumental-variables estimates of treatment effects, τ (IV), instead of reduced-form estimates. Column (6) imposes zero correlation between observed types, $corr(\pi, \tau) = 0$. Column (7) assumes that the utility cost of unemployment increases by 50 percent at month 12. Column (8) assumes no continuation loss after period 25 in the model. Column (9) restricts the budget so that only the top 30 percent of individuals can be treated.

is the constrained budget of column (9), where only 30 percent of the initial cohort can be treated: waiting until three months is then almost costless, although we verified that the gains decline more strongly when waiting beyond three months.

Panel C considers the optimal policy when treatment can be targeted across durations. In the baseline calibration, the gains of the global optimum are only marginally larger than those of optimal targeting at entry (2.64 versus 2.60 in Panel A), and 80 percent of targeted individuals are treated immediately: the option value of spreading treatment over the spell is small. Consistent with the timing results in Panel B, this option value disappears in the stationary environment of column (2), where all targeted individuals are treated at entry, and is negligible when adding unobserved heterogeneity or depreciation in the treatment effects (columns 3 and 4).

Timing flexibility becomes considerably more valuable, however, when the budget is tight: with only the top 30 percent treated (column 9), the gains rise from 3.35 to 3.54 and fewer than half of the targeted individuals are treated at entry. Intuitively, a small budget allows the agency to remain selective throughout the spell, and dynamic selection makes waiting attractive: as the spell progresses, survival itself reveals the workers with the poorest job-finding prospects and thus the highest returns to treatment. For similar reasons, targeting later also becomes more attractive when the utility loss from unemployment rises over the spell (column 7).

Another margin of flexibility concerns the duration of treatment itself, which we have so far restricted to a single period per job seeker. Appendix A.4 relaxes this restriction and shows that allowing the agency to target the same individuals in two consecutive periods further increases the welfare gains from optimal targeting under a fixed budget, as the highest-return job seekers can now be reached more than once. This exercise is similar in nature to restricting the budget to the top 30 percent of individuals (column 9): holding total resources fixed, treating individuals repeatedly means treating fewer of them, so that treatment is concentrated on a more select group of job seekers.

Taken together, the sensitivity analysis confirms the robustness of our main findings. Across all calibration scenarios, targeting rules that combine predicted treatment effects with predicted job-finding probabilities generate welfare gains around two to three times as large as those of random assignment, with the simple index τ/π remaining close to optimal throughout. The timing results are similarly robust: although the returns to treating a given worker rise over the unemployment spell through dynamic selection and duration dependence in job finding, a budget-constrained agency optimally treats the large majority of targeted job seekers at the start of the spell in nearly all scenarios.

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 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," Technical Report Discussion Paper No. 15830, IZA – Institute of Labor Economics 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, "Beyond Prediction: Using Big Data for Policy Problems," *Science*, 2017, *355* (6324), 483–485.

_ **, 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.

Barbanchon, Thomas Le, Lena Hensvik, and Roland Rathelot, "How Can AI Improve Search and Matching? Evidence from 59 Million Personalized Job Recommendations," 2023. SSRN Working Paper No. 4604814.

Behaghel, Luc, Sofia Dromundo, Marc Gurgand, Yagan Hazard, and Thomas Zuber, "The Potential of Recommender Systems for Directing Job Search: A Large-Scale Experiment," IZA Discussion Paper 16781, IZA Institute of Labor Economics, Bonn 2024.

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.

Berger, Mark C., Dan A. Black, and Jeffrey A. Smith, "Evaluating Profiling as a Means of Allocating Government Services," in Michael Lechner and Friedhelm Pfeiffer, eds., *Econometric Evaluation of Labour Market Policies*, Vol. 13 of *ZEW Economic Studies*, Heidelberg: Physica-Verlag, 2001, pp. 59–84.

Black, Dan A., Jeffrey A. Smith, Mark C. Berger, and Brett J. Noel, "Is the Threat of Reemployment Services More Effective Than the Services Themselves? Evidence from Random Assignment in the UI System," *American Economic Review*, 2003, *93* (4), 1313–1327.

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*, June 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.

Ernst, Sebastian, Andreas I. Mueller, and Johannes Spinnewijn, "Risk Scores for Long-Term Unemployment and the Assignment to Job Search Counseling," *AEA Papers and Proceedings*, May 2024, *114*, 572–576.

Gregory, Victoria, Guido Menzio, and David Wiczer, "The Alpha Beta Gamma of the Labor Market," *Journal of Monetary Economics*, 2025, *150*, 103695. Labor Market Dynamics: Drivers and Consequences.

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.

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, "Human Decisions and Machine Predictions," *Quarterly Journal of Economics*, 2018, *133* (1), 237–293.

_ , Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer, "Prediction Policy Problems," *American Economic Review*, 2015, *105* (5), 491–495.

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.

Manski, Charles F., "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 2004, *72* (4), 1221–1246.

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.

Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations," *Science*, 2019, *366* (6464), 447–453.

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. Step 1: The agency's problem. The agency operates under a fixed budget $\bar{B}$. Given a per-period targeting cost c, it can treat $\bar{B}/c$ individuals for one period and for one period only. Abstracting from distortionary taxation and the optimal overall size of the program, the agency allocates treatment across unemployed individuals to minimize the expected welfare losses associated with unemployment:

$$\min_{T} \quad \int\int L^T(\pi,\tau)g(\pi,\tau)d\pi d\tau \qquad \text{(A1)}$$
$$s.t.\ \bar{B} \geq \int\int B^T(\pi,\tau)g(\pi,\tau)d\pi d\tau,$$

where $L^T(\pi,\tau)$ and $B^T(\pi,\tau)$ are the type-level expected utility loss and budgetary cost defined in Section 2. Utility-loss levels u may differ across types and are observable, so that policies may condition on u; we make this dependence explicit in continuation losses, returns, the optimal targeting duration, and the treatment rule below, while leaving the u-dimension implicit in the type distribution $g(\pi,\tau)$ and in the full-spell objects (the budgetary cost B^T does not depend on u).

Step 2: Reformulation in terms of treatment indicators. A policy that treats each type at most once, for one period, is described by indicators $T_d(\pi,\tau) \in \{0,1\}$ with $\sum_{d=0}^{\bar{D}} T_d(\pi,\tau) \leq 1$, where $\bar{D}$ is the maximum duration of unemployment after which no further treatment is allowed. Treatment at duration d leaves survival up to d unaffected (no treatment occurs before d) and changes outcomes only for the $S_d(\pi,\tau)$ survivors: losses incurred before d are policy-invariant, and the loss from d onward changes from $L_d(\pi,u)$ to $L_d^T(\pi,\tau,u)$, where $L_d^T(\pi,\tau,u)$ and $L_d(\pi,u)$ denote the continuation losses from period d onward with and without treatment. Hence

$$L^T(\pi,\tau) \;=\; L(\pi) + \sum_{d=0}^{\bar{D}} T_d(\pi,\tau)\, S_d(\pi,\tau)\big(L_d^T(\pi,\tau,u) - L_d(\pi,u)\big), \qquad B^T(\pi,\tau) \;=\; \sum_{d=0}^{\bar{D}} c\, T_d(\pi,\tau)\, S_d(\pi,\tau),$$

since the treatment cost c is incurred at duration d only for those still unemployed at d. Substituting into (A1) and dropping the policy-invariant baseline loss $\int\int L(\pi)g(\pi,\tau)\,d\pi\,d\tau$ from the objective, the problem becomes

$$\min_{\{T_d\}} \quad \sum_{d=0}^{\bar{D}} \int\int T_d(\pi,\tau)\, S_d(\pi,\tau)(L_d^T(\pi,\tau,u) - L_d(\pi,u))g(\pi,\tau)d\pi d\tau \tag{A2}$$
$$s.t.\ \bar{B} \geq \sum_{d=0}^{\bar{D}} \int\int c\, T_d(\pi,\tau)\, S_d(\pi,\tau) g(\pi,\tau) d\pi d\tau,$$
$$T_d(\pi,\tau) \in \{0,1\}, \quad \sum_{d=0}^{\bar{D}} T_d(\pi,\tau) \leq 1.$$

In this notation, the return to targeting defined in equation (3) satisfies $-c\, R_d^T(\pi,\tau,u) = L_d^T(\pi,\tau,u) - L_d(\pi,u)$.

Step 3: Lagrangian and separation by type. Attach a multiplier $\lambda \geq 0$ to the budget constraint. The contribution of treating type (π,τ) at duration d to the Lagrangian is

$$S_d(\pi,\tau)\, g(\pi,\tau)\Big[\big(L_d^T(\pi,\tau,u) - L_d(\pi,u)\big) + \lambda c\Big] = -\, c\, S_d(\pi,\tau)\, g(\pi,\tau)\big[R_d^T(\pi,\tau,u) - \lambda\big]. \tag{A3}$$

The agency minimizes the Lagrangian, so each treated cell should make (A3) as negative as possible. Because the Lagrangian is an integral of terms that are separable across types, and the exclusivity constraint allows at most one nonzero indicator per type, the problem separates: for each type (π,τ) the agency solves

$$\max_d \quad \underbrace{S_d(\pi,\tau)\big[R_d^T(\pi,\tau,u) - \lambda\big]}_{\text{net welfare gain per cost, weighted by survival}}, \tag{A4}$$

treating at the maximizing duration if the maximized value is positive, and leaving the type untreated otherwise. The per-survivor return R_d^T governs whether a cell is worth treating relative to the shadow cost λ, but the survival rate $S_d(\pi,\tau)$ governs how many individuals that treatment reaches and hence how much aggregate welfare the cell delivers: it is not enough for the return to be high if almost no one survives to that duration to receive treatment. The optimal targeting duration therefore solves

$$d^\star(\pi,\tau,u;\lambda) \in \arg\max_d \; S_d(\pi,\tau)\big[R_d^T(\pi,\tau,u) - \lambda\big], \tag{A5}$$

trading off the survival rate, declining in d, against the net return $R_d^T - \lambda$, which may rise in d as the surviving unemployed become more disadvantaged. Since $S_{d^\star}(\pi,\tau) > 0$, the maximized value in (A4) is positive if and only if the return at the optimal duration exceeds the multiplier:

$$S_{d^\star}(\pi,\tau)\big[R_{d^\star}^T(\pi,\tau,u) - \lambda\big] > 0 \quad \Longleftrightarrow \quad R_{d^\star}^T(\pi,\tau,u) > \lambda. \tag{A6}$$

Step 4: Existence of the threshold and budget clearing. Let $E(\lambda)$ denote the expenditure induced by the policy just described,

$$E(\lambda) \;\equiv\; \iint c\, S_{d^\star(\pi,\tau,u;\lambda)}(\pi,\tau)\, \mathbf{1}\big\{R_{d^\star}^T(\pi,\tau,u) > \lambda\big\}\, g(\pi,\tau)\, d\pi\, d\tau. \tag{A7}$$

$E(\lambda)$ is nonincreasing in λ for two reasons. First, the treated set shrinks: the maximized value in (A4) is decreasing in λ for every type, so any type left untreated at λ remains untreated at $\lambda' > \lambda$. Second, each treated type weakly delays treatment: for $\lambda' > \lambda$, optimality of $d^\star(\lambda)$ at λ and of $d^\star(\lambda')$ at λ' implies

$$S_{d^\star(\lambda)}\big[R^T_{d^\star(\lambda)} - \lambda\big] \geq S_{d^\star(\lambda')}\big[R^T_{d^\star(\lambda')} - \lambda\big] \quad \text{and} \quad S_{d^\star(\lambda')}\big[R^T_{d^\star(\lambda')} - \lambda'\big] \geq S_{d^\star(\lambda)}\big[R^T_{d^\star(\lambda)} - \lambda'\big];$$

adding the two inequalities yields $(\lambda' - \lambda)\big[S_{d^\star(\lambda)} - S_{d^\star(\lambda')}\big] \geq 0$, so $S_{d^\star(\pi,\tau,u;\lambda)}$ is nonincreasing in λ and each treated type's cost $c\, S_{d^\star}$ weakly falls. Moreover, $E(0)$ is the expenditure of the $\lambda = 0$ policy, which treats every type whose return is positive at some duration. Note that no sign restriction on returns is required: types whose return is nonpositive at every duration fail the criterion in (A6) for any $\lambda \geq 0$ and simply remain untreated. Finally, $E(\lambda) = 0$ for λ above the highest return. If $\bar{B} \geq E(0)$, the budget constraint is slack and the policy with $\lambda^\star = 0$ is optimal. Otherwise, continuity of $g(\pi, \tau)$ ensures that the set of types indifferent between treatment durations or exactly at the threshold has measure zero, so $E(\lambda)$ is continuous, and there exists $\lambda^\star > 0$ such that the budget binds, $E(\lambda^\star) = \bar{B}$. The same measure-zero property implies that the $\geq$ and $>$ cases in the treatment rule coincide, so the threshold $\bar{R} = \lambda^\star$ is well-defined.

Step 5: Global optimality. Let $T^\star$ denote the policy constructed above, which minimizes the Lagrangian pointwise at $\lambda^\star$ and exhausts the budget (complementary slackness: $\lambda^\star\big[E(\lambda^\star) - \bar{B}\big] = 0$). For any feasible policy T, writing $V(T)$ for the objective and $C(T)$ for the expenditure in (A2),

$$V(T^\star) \;=\; V(T^\star) + \lambda^\star\big[C(T^\star) - \bar{B}\big] \;\leq\; V(T) + \lambda^\star\big[C(T) - \bar{B}\big] \;\leq\; V(T),$$

where the first inequality holds because $T^\star$ minimizes the Lagrangian and the second because T is feasible ($C(T) \leq \bar{B}$) and $\lambda^\star \geq 0$. Hence $T^\star$ is optimal, establishing parts (i) and (ii) with $\bar{R} = \lambda^\star$: for each type, the agency selects the single targeting duration that maximizes the survival-weighted net return (A5) — the survival rate enters because the budget is spent on, and welfare accrues to, only those who remain unemployed at the treated duration — and it treats those types whose return at their optimal duration exceeds $\bar{R}$, the return forgone on the marginal treated type.

Step 6: Simplification. Suppose $\sigma_{\varepsilon^\pi} = \sigma_{\varepsilon^\tau} = 0$, and write $\pi_d \equiv \pi\,\theta_{\pi,d}$ and $\tau_d \equiv \tau\,\theta_{\tau,d}$ for the baseline job-finding probability and the treatment effect at duration d. Continuation losses then satisfy the recursions

$$L_d(\pi, u) = u_d + (1 - \pi_d)\, L_{d+1}(\pi, u), \qquad L^T_d(\pi, \tau, u) = u_d + (1 - \pi_d - \tau_d)\, L_{d+1}(\pi, u),$$

since a one-period treatment at d raises the exit hazard at d by τ_d and leaves hazards from $d + 1$ onward unchanged. Subtracting, $L_d(\pi, u) - L^T_d(\pi, \tau, u) = \tau_d\, L_{d+1}(\pi, u)$ exactly, so $R^T_d = (\tau_d/c)\, L_{d+1}(\pi, u)$. With $\theta_{\pi,d} = \theta_{u,d} = 1$ for all d, the continuation loss is stationary and solves $L = u + (1 - \pi)L$, i.e., $L_{d+1}(\pi, u) = u/\pi$ for all d; with $\theta_{\tau,d} = 1$ the treatment effect is constant, $\tau_d = \tau$. Hence

$$R^T_d(\pi, \tau, u) \;=\; \frac{u}{c} \cdot \frac{\tau}{\pi} \qquad \text{for all } d,$$

and the threshold rule in part (i) reduces to a threshold rule in $u\,\tau/\pi$, and in τ/π when utility losses are identical across types. Since $R_d^T - \bar{R}$ is then constant in d while $S_d(\pi) = (1-\pi)^d$ is strictly decreasing, the survival-weighted excess return of every treated type is uniquely maximized at $d^\star = 0$: with returns constant over the spell, delaying treatment only shrinks the group that can be reached. $\square$

Proof of proposition 2.

Proof. Write $p \equiv \pi + \varepsilon^\pi$ and $t \equiv \tau + \varepsilon^\tau$, so that $P_d = p\,\theta_\pi^d$ and $\tau_d = t\,\theta_\tau^d$, and let

$$L_{d+1}(p) \;=\; \sum_{\delta=d+1}^{\infty} u\,\theta_u^\delta \prod_{j=d+1}^{\delta-1} \big(1 - p\,\theta_\pi^j\big)$$

denote the individual continuation loss from period $d+1$ onward. Survivor expectations are $\mathbb{E}_d[X] \equiv \mathbb{E}[X\,S_d(\pi,\varepsilon^\pi)]/\mathbb{E}[S_d(\pi,\varepsilon^\pi)]$, with Var_d and Cov_d the corresponding variance and covariance. Throughout, "dispersion" refers to the standard deviations of $(\varepsilon^\pi, \varepsilon^\tau)$ among survivors, and orders of approximation are joint in $(\theta_\pi - 1)$, $(\theta_u - 1)$, and the dispersion.

Step 1: The type-level return. Since targeting conditions only on observables, treating a type at duration d treats all its surviving ε-realizations. A one-period treatment at d raises the individual exit hazard at d by $\tau_d = t\,\theta_\tau^d$ and leaves hazards from $d+1$ onward unchanged, so it lowers the individual continuation loss by $t\,\theta_\tau^d\,L_{d+1}(p)$. The type-level return defined in equation (3) — with the continuation losses L_d^T and L_d averaged over $(\varepsilon^\pi, \varepsilon^\tau)$ among survivors at duration d — therefore equals

$$R_d^T \;=\; \frac{\theta_\tau^d\,\mathbb{E}\big[t\,S_d\,L_{d+1}(p)\big]}{c\,\mathbb{E}[S_d]} \;=\; \frac{\theta_\tau^d}{c}\,\mathbb{E}_d\big[t\,L_{d+1}(p)\big]. \tag{A8}$$

Step 2: Dynamic selection. Untreated survival satisfies $S_{d+1} = S_d\,(1 - P_d)$, so survivor expectations tilt as $\mathbb{E}_{d+1}[X] = \mathbb{E}_d[(1-P_d)\,X]/\mathbb{E}_d[1 - P_d]$. Using $P_{d+1} = \theta_\pi\,P_d$,

$$\bar{P}_{d+1} = \theta_\pi\,\frac{\mathbb{E}_d[P_d(1-P_d)]}{1-\bar{P}_d} = \theta_\pi\,\frac{\bar{P}_d - \bar{P}_d^2 - \mathrm{Var}_d(P_d)}{1-\bar{P}_d} = \theta_\pi\,\bar{P}_d\left[1 - \frac{\mathrm{Var}_d(P_d)}{\bar{P}_d(1-\bar{P}_d)}\right],$$

an exact dynamic-selection identity: the observed change $\bar{\theta}_{\pi,d}$ compounds true duration dependence (θ_π) and dynamic selection (the variance term). Analogously, since $\bar{\tau}_d = \theta_\tau^d\,\mathbb{E}_d[t]$, the observed change in treatment effects satisfies $\bar{\theta}_{\tau,d} = \theta_\tau\,\mathbb{E}_{d+1}[t]/\mathbb{E}_d[t]$.

Step 3: Expansion of the individual continuation loss. Setting $k = \delta - d - 1$,

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

which equals $u\sum_k (1-p)^k = u/p$ at $\theta_\pi = \theta_u = 1$. For the θ_π-derivative, the $k=0$ term is independent of θ_π; for $k \geq 1$, differentiating the j-th factor gives $-(d+j)\theta_\pi^{d+j-1}p$, while the remaining $k-1$

factors each equal $(1-p)$ at $\theta_\pi = 1$ and the prefactor θ_u^{d+k+1} equals one at $\theta_u = 1$. Hence

$$\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=\theta_u=1} = -u\,p\sum_{k=1}^{\infty}(1-p)^{k-1}\sum_{j=1}^{k}(d+j) = -u\sum_{j=1}^{\infty}(d+j)(1-p)^{j-1} = -\frac{u\,(dp+1)}{p^2},$$

where the second equality swaps the order of summation and evaluates the inner geometric sum, $\sum_{k=j}^{\infty}(1-p)^{k-1} = (1-p)^{j-1}/p$, and the third re-indexes $m = j-1$ and uses $\sum_{m=0}^{\infty}(1-p)^m = 1/p$ and $\sum_{m=0}^{\infty} m(1-p)^m = (1-p)/p^2$. For the θ_u-derivative, differentiating the prefactor θ_u^{d+k+1} gives $(d+k+1)$ at $\theta_u = 1$, so

$$\left.\frac{\partial L_{d+1}}{\partial \theta_u}\right|_{\theta_\pi=\theta_u=1} = u\sum_{k=0}^{\infty}(d+k+1)(1-p)^k = \frac{u\,(dp+1)}{p^2} = -\left.\frac{\partial L_{d+1}}{\partial \theta_\pi}\right|_{\theta_\pi=\theta_u=1}.$$

Hence, to first order in $(\theta_\pi - 1)$ and $(\theta_u - 1)$,

$$L_{d+1}(p) \;\approx\; \frac{u}{p} + c_d(p)\,\big[(\theta_u - 1) - (\theta_\pi - 1)\big], \qquad c_d(p) \equiv \frac{u\,(dp+1)}{p^2}, \tag{A9}$$

and the key property, holding individual by individual, is that the coefficient grows by

$$c_{d+1}(p) - c_d(p) \;=\; \frac{u}{p}$$

across consecutive durations.

Step 4: Part (i), returns across types. Fix a duration d and evaluate to zeroth order in $(\theta_\pi - 1)$ and $(\theta_u - 1)$, so that $L_{d+1}(p) = u/p$. Writing $P_d = \bar{P}_d(1+x)$ with $\mathbb{E}_d[x] = 0$ and expanding, $\mathbb{E}_d[1/P_d] = \big(1+\mathrm{CV}_d^2\big)/\bar{P}_d + O\big(\mathbb{E}_d[x^3]\big)$ with $\mathrm{CV}_d^2 \equiv \mathrm{Var}_d(P_d)/\bar{P}_d^2$, so that by (A8), using $1/p = \theta_\pi^d/P_d$,

$$R_d^T = \frac{u\,\theta_\tau^d}{c}\,\mathbb{E}_d[t]\,\mathbb{E}_d[1/p]\,\big(1+\rho_d\big) = \frac{u\,\bar{\tau}_d\,\theta_\pi^d}{c\,\bar{P}_d}\,\big(1+\mathrm{CV}_d^2\big)\big(1+\rho_d\big) + O\big(\mathbb{E}_d[x^3]\big), \qquad \rho_d \equiv \frac{\mathrm{Cov}_d(t, 1/p)}{\mathbb{E}_d[t]\,\mathbb{E}_d[1/p]}.$$

Taking the ratio across two types at the same duration — the factor θ_π^d/c is common to both —

$$\frac{R_d^T(\pi',\tau',u')}{R_d^T(\pi,\tau,u)} = \frac{\bar{\tau}_d'\,u'/\bar{P}_d'}{\bar{\tau}_d\,u/\bar{P}_d}\cdot\frac{\big(1+\mathrm{CV}_d'^2\big)\big(1+\rho_d'\big)}{\big(1+\mathrm{CV}_d^2\big)\big(1+\rho_d\big)}.$$

The survivor dispersion corrections CV_d^2 and ρ_d are second order in the dispersion and, unlike in part (ii) below, do not difference out across types, so the second factor is $1 + O(\text{dispersion}^2)$. Log-linearizing the first factor yields part (i); reinstating the first-order terms in (A9) adds type-specific corrections of order $(\theta_\pi - 1)$ and $(\theta_u - 1)$, consistent with the neighborhood qualifier in the proposition.

Step 5: Part (ii), composition channel. We first track how the survivor average of $1/p$ drifts across durations. By the tilting identity and $P_d/p = \theta_\pi^d$,

$$\mathbb{E}_{d+1}[1/p] = \frac{\mathbb{E}_d[(1-P_d)/p]}{1-\bar{P}_d} = \frac{\mathbb{E}_d[1/p] - \theta_\pi^d}{1-\bar{P}_d},$$

so that, dividing by $\mathbb{E}_d[1/p]$ and using $\mathbb{E}_d[1/p] = \theta_\pi^d\,\mathbb{E}_d[1/P_d]$,

$$\frac{\mathbb{E}_{d+1}[1/p]}{\mathbb{E}_d[1/p]} = \frac{1-H_d}{1-\bar{P}_d}, \qquad H_d \equiv \frac{1}{\mathbb{E}_d[1/P_d]},$$

exactly, where H_d is the harmonic mean of the job-finding probability among survivors. Expanding as in Step 4, $H_d = \bar{P}_d\big(1-\mathrm{CV}_d^2\big) + O\big(\mathbb{E}_d[x^3]\big)$, so

$$\frac{\mathbb{E}_{d+1}[1/p]}{\mathbb{E}_d[1/p]} = 1 + \frac{\mathrm{Var}_d(P_d)}{\bar{P}_d(1-\bar{P}_d)} + O\big(\mathbb{E}_d[x^3]\big) = 1 + \Big(1 - \frac{\bar{\theta}_{\pi,d}}{\theta_\pi}\Big) + O\big(\mathbb{E}_d[x^3]\big),$$

where the last equality uses the dynamic-selection identity of Step 2: the second-order composition effect is absorbed exactly by the observed change $\bar{\theta}_{\pi,d}$, and the error is third order (the skewness of the survivor distribution of p).

Step 6: Part (ii), duration-dependence channel. Taking survivor expectations of (A9) at durations d and $d+1$,

$$\mathbb{E}_d[L_{d+1}] \approx u\,\mathbb{E}_d[1/p] + (\theta_u - \theta_\pi)\,\mathbb{E}_d[c_d(p)], \qquad \mathbb{E}_{d+1}[L_{d+2}] \approx u\,\mathbb{E}_{d+1}[1/p] + (\theta_u - \theta_\pi)\,\mathbb{E}_{d+1}[c_{d+1}(p)],$$

so that, dividing,

$$\frac{\mathbb{E}_{d+1}[L_{d+2}]}{\mathbb{E}_d[L_{d+1}]} \;\approx\; \frac{\mathbb{E}_{d+1}[1/p]}{\mathbb{E}_d[1/p]}\left[1 + (\theta_u - \theta_\pi)\left(\frac{\mathbb{E}_{d+1}[c_{d+1}(p)]}{u\,\mathbb{E}_{d+1}[1/p]} - \frac{\mathbb{E}_d[c_d(p)]}{u\,\mathbb{E}_d[1/p]}\right)\right].$$

Since $c_d(p)/u = d/p + 1/p^2$, the coefficient ratio equals $d + \mathbb{E}_d[1/p^2]/\mathbb{E}_d[1/p]$, so its change across consecutive durations is

$$1 + \left(\frac{\mathbb{E}_{d+1}[1/p^2]}{\mathbb{E}_{d+1}[1/p]} - \frac{\mathbb{E}_d[1/p^2]}{\mathbb{E}_d[1/p]}\right) \;=\; 1 + O(\text{dispersion}^2)\,:$$

the d-dependent terms cancel exactly — the survivor-average counterpart of $c_{d+1}(p) - c_d(p) = u/p$ — and the survivor drift of $\mathbb{E}[1/p^2]/\mathbb{E}[1/p]$ is second order in the dispersion, since tilting shifts survivor moments by covariance terms. Combining with Step 5 and dropping products of a second-order dispersion term with $(\theta_u - \theta_\pi)$ — third order jointly —

$$\frac{\mathbb{E}_{d+1}[L_{d+2}]}{\mathbb{E}_d[L_{d+1}]} \;\approx\; 1 + \Big(1 - \frac{\bar{\theta}_{\pi,d}}{\theta_\pi}\Big) + (\theta_u - \theta_\pi) \;\approx\; 1 + \theta_u - \bar{\theta}_{\pi,d},$$

where the last step replaces $1 - \bar{\theta}_{\pi,d}/\theta_\pi$ by $\theta_\pi - \bar{\theta}_{\pi,d}$, an error of $(\theta_\pi - \bar{\theta}_{\pi,d})(1-\theta_\pi)/\theta_\pi$ — second order in the dispersion times first order in $(\theta_\pi - 1)$, hence third order jointly. Finally, the expansion (A9) itself carries a Taylor remainder that is second order in $(\theta_\pi - 1, \theta_u - 1)$, and unlike the first-order coefficients, whose duration-differences cancel to u/p, its coefficients do not cancel across consecutive durations. The resulting deterministic error is second order in $(\theta_\pi - 1, \theta_u - 1)$ and proportional to $(\theta_\pi - 1)$: it vanishes identically at $\theta_\pi = 1$, where $L_{d+1}(p) = u\,\theta_u^{d+1}/(1-\theta_u(1-p))$ and the ratio equals θ_u exactly (see the Remark below); a direct second-order expansion shows it is of order $(\theta_\pi - 1)(\theta_\pi - \theta_u)$. Hence all error terms involving the unobserved heterogeneity are third order jointly, while the deterministic error is second order in $(\theta_\pi - 1, \theta_u - 1)$.

Step 7: Part (ii), treatment-effect heterogeneity. Write $\mathbb{E}_d[t\, L_{d+1}] = \mathbb{E}_d[t]\,\mathbb{E}_d[L_{d+1}] + \mathrm{Cov}_d(t, L_{d+1})$. Relative to the product of the means, the covariance term is second order in the dispersion; its change across consecutive durations is a third joint moment of the survivor distribution (tilting a covariance changes it by third central moments), hence third order, so it cancels from the ratio to the stated order. Using (A8), $\bar{\theta}_{\tau,d} = \theta_\tau\, \mathbb{E}_{d+1}[t]/\mathbb{E}_d[t]$ from Step 2, and Step 6,

$$\frac{R^T_{d+1}}{R^T_d} = \theta_\tau \cdot \frac{\mathbb{E}_{d+1}[t\, L_{d+2}]}{\mathbb{E}_d[t\, L_{d+1}]} \;\approx\; \theta_\tau \cdot \frac{\mathbb{E}_{d+1}[t]}{\mathbb{E}_d[t]} \cdot \frac{\mathbb{E}_{d+1}[L_{d+2}]}{\mathbb{E}_d[L_{d+1}]} \;\approx\; \bar{\theta}_{\tau,d}\left(1 + \theta_u - \bar{\theta}_{\pi,d}\right),$$

which is exact in θ_τ. Subtracting one and dropping the cross term $(\bar{\theta}_{\tau,d} - 1)(\theta_u - \bar{\theta}_{\pi,d})$ — second order only for θ_τ close to one, so that the difference form additionally requires θ_τ near one — yields

$$\frac{R^T_{d+1} - R^T_d}{R^T_d} \;\approx\; \bar{\theta}_{\tau,d} + \theta_u - \bar{\theta}_{\pi,d} - 1.$$

Step 8: Nesting. Without unobserved heterogeneity the survivor distribution is degenerate, so $\bar{P}_d = \pi_d$, $\bar{\tau}_d = \tau_d$, $\bar{\theta}_{\pi,d} = \theta_\pi$, $\bar{\theta}_{\tau,d} = \theta_\tau$, all dispersion terms vanish, and parts (i) and (ii) reduce to their deterministic counterparts. If ε^τ is mean-independent of ε^π, then $\mathbb{E}_d[t] = \mathbb{E}[t]$ does not vary with d, since the survival weights depend on ε^π only, so $\bar{\theta}_{\tau,d} = \theta_\tau$. □

Remark. When $\theta_\pi = 1$, the continuation loss has the closed form $L_{d+1}(p) = u\,\theta_u^{d+1}\big/\big(1 - \theta_u(1-p)\big)$, so that the type-level return is, exactly,

$$R^T_d \;=\; \frac{u\,\theta_u^{d+1}\,\theta_\tau^d}{c}\,\mathbb{E}_d\left[\frac{t}{1 - \theta_u(1-p)}\right] :$$

at the individual level, the return to treatment is proportional to $t\big/\big(1 - \theta_u(1-p)\big)$, the treatment effect scaled by the inverse of the effective exit rate from the loss-weighted spell. If in addition ε^τ is mean-independent of ε^π, an exact version of part (ii) follows: defining the generalized harmonic mean $\tilde{H}_d \equiv 1\big/\mathbb{E}_d\big[\big(1-\theta_u(1-p)\big)^{-1}\big]$, we have $R^T_{d+1}/R^T_d = \theta_\tau\,\theta_u\,\tilde{H}_d/\tilde{H}_{d+1}$ exactly, for any distribution of ε^π; at $\theta_u = 1$, $\tilde{H}_d$ reduces to the harmonic-mean job-finding probability among survivors, $1/\mathbb{E}_d[1/P_d]$. The arithmetic-mean statement in part (ii) approximates this exact statistic and is the empirically more convenient version.

A.2 Additional Figures and Tables

Figure A1: Distribution of π and τ and Duration Dependence in Job Finding in Calibrated Model

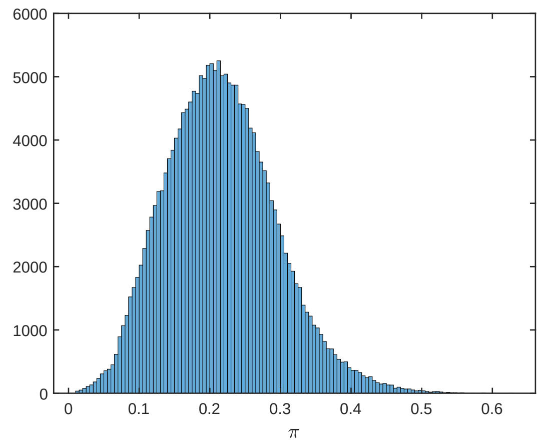

(a) Histogram of π

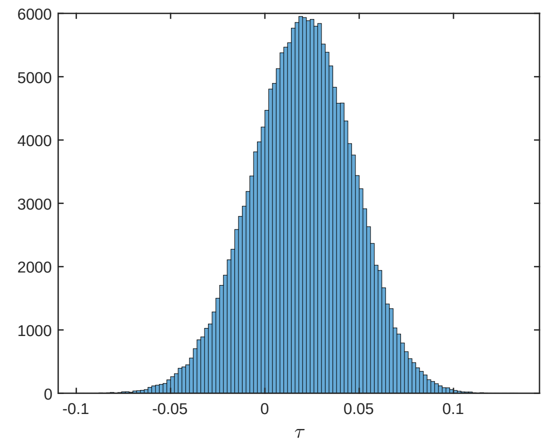

(b) Histogram of τ

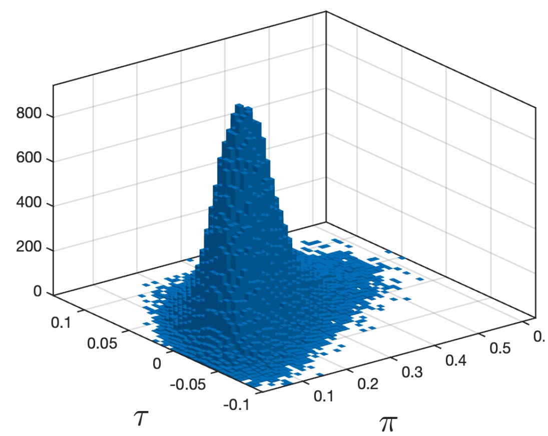

(c) Histogram of π and τ

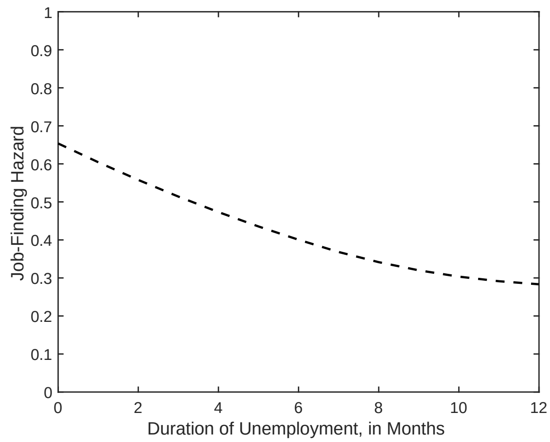

(d) Duration Dependence in Job Finding

Notes: The histograms shows the distributions of the 200,000 individuals sampled for our model calibration.

Table A1: Shapley Decomposition of the Variance of the Targeting Index τ/π

$\mathrm{Var}(\tau/\pi)$	0.0316
$V_\tau = \mathrm{Var}(\tau \cdot \mathrm{E}[1/\pi])$	0.0217
$V_\pi = \mathrm{Var}(\mathrm{E}[\tau] \cdot (1/\pi))$	0.0039
Shapley contribution φ_τ	0.0247
Shapley contribution φ_π	0.0069
Share φ_τ	0.781
Share φ_π	0.219

Notes: Based on model simulations with 200,000 individuals from the baseline calibration, using raw treatment effects (no floor or sample restriction). π is the one-month job-finding probability at the start of the unemployment spell absent treatment (observed part), and τ is the treatment effect of the observed type: the increase in the one-month job-finding probability in the month of treatment, with treatment at the start of the spell. V_τ fixes $1/\pi$ at its mean so that only τ varies, and V_π fixes τ at its mean so that only π varies. The Shapley contributions $\varphi_\tau = (V_\tau + \mathrm{Var}(\tau/\pi) - V_\pi)/2$ and $\varphi_\pi = (V_\pi + \mathrm{Var}(\tau/\pi) - V_\tau)/2$ sum exactly to $\mathrm{Var}(\tau/\pi)$; the Share rows divide them by $\mathrm{Var}(\tau/\pi)$.

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

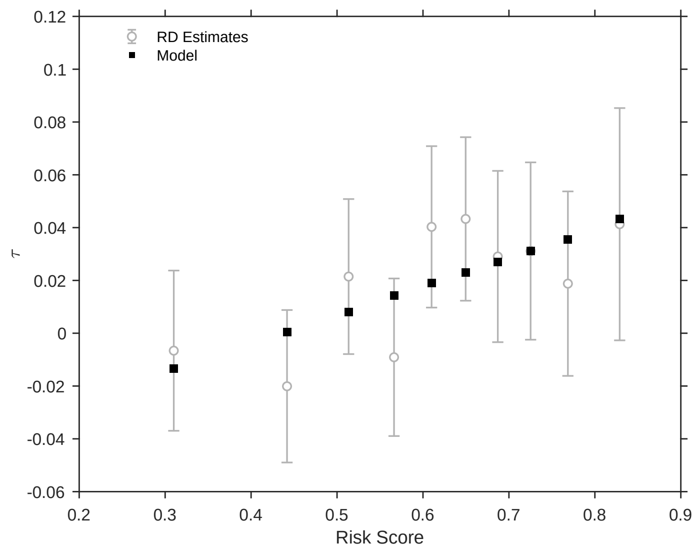

Note: The outcome variable for the RD estimates is the probability of working within the next four weeks.

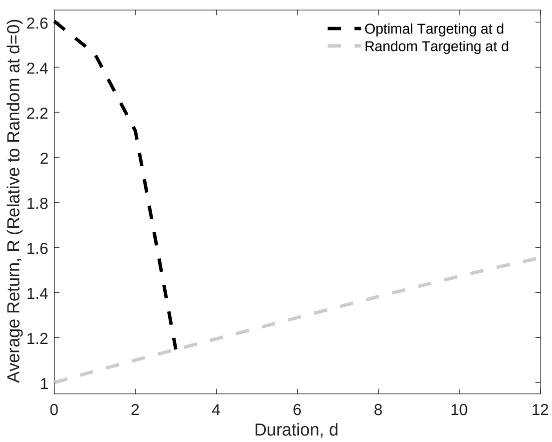

(a) Optimal Targeting at d vs. Random at $d = 0$

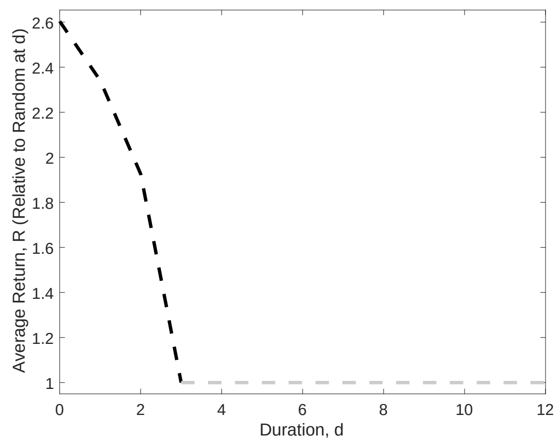

(b) Optimal Targeting at d vs. Random at d

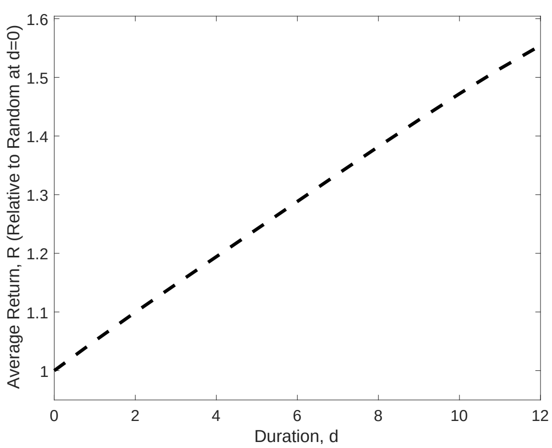

(c) Random Targeting at d vs. Random at $d = 0$

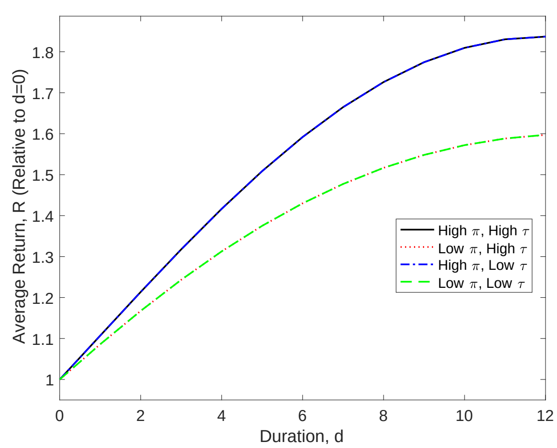

(d) Targeting at d vs. $d = 0$, By Type

Figure A3: Welfare Gains of Targeting, By Duration

A.3 Algorithmic Bias

To assess algorithmic bias and its effect on the correlation between risk scores (RS) and treatment effects (τ), we use the treatment effects from our baseline calibration to the IV estimates and treat 53% of the sample. For this exercise, we assume that the observed risk score (RS) incorporates treatment effects. We then infer the underlying baseline job-finding probability π and construct a counterfactual risk score ($\widetilde{RS}$) as the 6-month job-finding probability in the absence of treatment.

Table A2 and Figures A4 and A5 report the simulation results. In Panel A of the table and in Figure A4, treatment occurs in period 2 of the unemployment spell, and we report various moments for the baseline risk score RS and the treatment-exclusive risk score $\widetilde{RS}$. The last two columns of the table are the most relevant: they report the bivariate regression coefficient of τ on the risk score, which is the moment to which we calibrate the model. The results confirm the presence of a bias: the coefficient falls from 0.73 when using RS to 0.50 when using $\widetilde{RS}$. A similar bias arises in a version of the model with no correlation between the observed risk score and the treatment effects. The binned scatter plots in Figure A4 illustrate these patterns and show that, despite the bias, the positive relationship between the counterfactual risk score and the treatment effect remains strong in the baseline calibration.

Panel B of the table and Figure A5 confirm that these patterns are robust to treating individuals with $RS < 0.65$ in month 2 and those with $RS \geq 0.65$ in month 5, as in the Belgian policy. The results are very similar to those in Panel B and Figure A4.

Table A2: Algorithmic Bias Based on Model Simulations

	Cov(RS, τ)	Cov($\widetilde{RS}, \tau$)	Var(RS)	Var($\widetilde{RS}$)	Var(τ)	$\frac{\text{Cov}(RS,\tau)}{\text{Var}(RS)}$	$\frac{\text{Cov}(\widetilde{RS},\tau)}{\text{Var}(\widetilde{RS})}$
Panel A: Treatment at 2 months							
Baseline Calibration	0.0164	0.0105	0.0226	0.0210	0.0236	0.7255	0.5019
Calibration with cov(RS,τ)=0	0.0001	-0.0036	0.0226	0.0299	0.0118	0.0023	-0.1220
Panel B: Treatment at 2 months for RS < 0.65 *and at 5 months for RS* ≥ 0.65							
Baseline Calibration	0.0164	0.0108	0.0226	0.0212	0.0236	0.7255	0.5100
Calibration with cov(RS,τ)=0	0.0001	-0.0036	0.0226	0.0301	0.0118	0.0023	-0.1186

Notes: Based on model simulations with 200,000 individuals. RS denotes the baseline risk score. $\widetilde{RS}$ denotes the risk score without treatment at a given duration. For the baseline estimates, we assume that everyone is treated at a given duration, whereas for the IV estimates we assume that only 53% of the sample (chosen at random) is treated at a given duration.

Figure A4: Algorithmic Bias Simulations: Binned Scatter Plots (Treatment in Month 2)

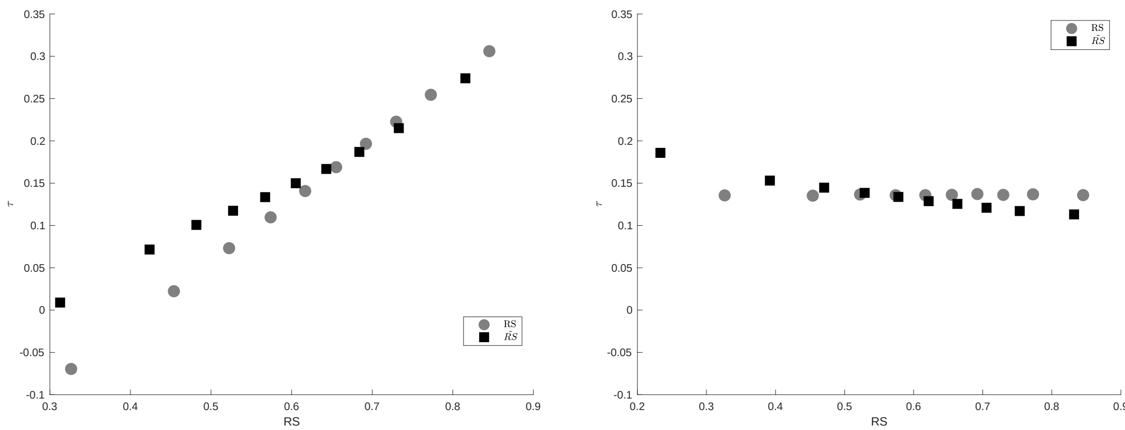

(a) Baseline Calibration (b) Calibration with cov(RS,τ)=0

Notes: Based on model simulations with 200,000 individuals. RS denotes the baseline risk score. $\widetilde{RS}$ denotes the risk score without treatment.

Figure A5: Algorithmic Bias Simulations: Binned Scatter Plots (Treatment in 2 months for $RS <$ 0.65 and in 5 months for $RS \geq 0.65$)

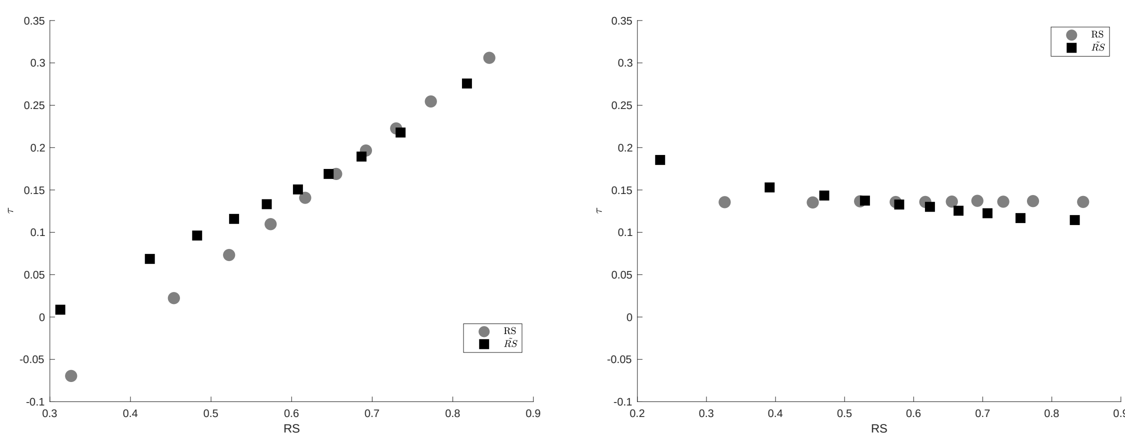

(a) Baseline Calibration (b) Calibration with cov(RS,τ)=0

Notes: Based on model simulations with 200,000 individuals. RS denotes the baseline risk score. $\widetilde{RS}$ denotes the risk score without treatment.

A.4 Optimal Targeting With Targeting for Two Consecutive Periods

Figure A6: Optimal Targeting when Targeting in Two Consecutive Periods

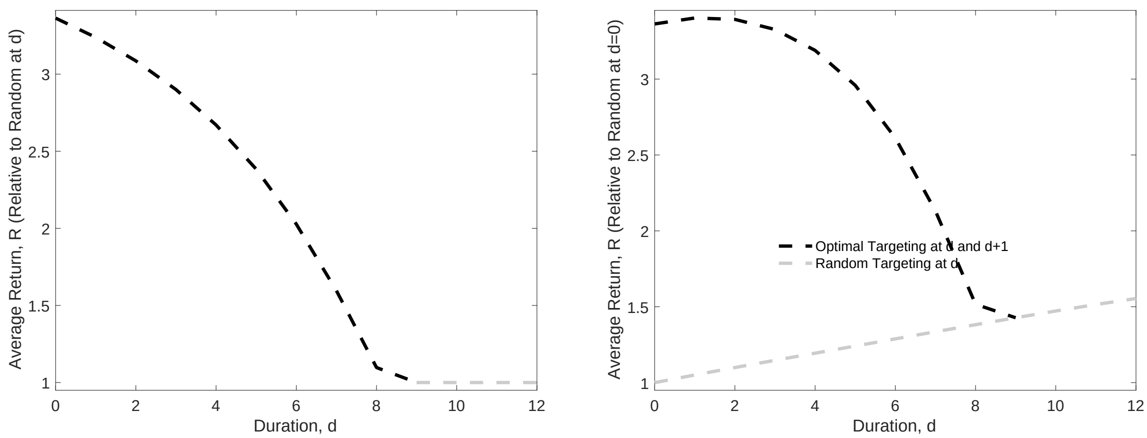

(a) Welfare Gain of Optimal Targeting at d and $d+1$ relative to Average Targeting at d

(b) Welfare Gain of Optimal Targeting at d and $d+1$ relative to Average Targeting at $d = 1$

Figure A6 shows that allowing the policymaker to target individuals in two consecutive periods—rather than just one—increases the welfare gains from optimal targeting, holding the budget fixed (see Figure A3 for comparison for the optimal single period targeting). The reason is that the unemployment agency can now reach the highest-return individuals multiple times.

B Data

B.1 Descriptive Statistics

	Full Sample	Baseline RD	Het. RD	Dyn. RD
Share female	0.46	0.45	0.40	0.44
Share of Belgian origin	0.67	0.73	0.57	0.69
Share tertiary educated	0.25	0.26	0.14	0.23
Mean age	34.7	30.6	35.1	31.2
Share with very good Dutch ability	0.70	0.80	0.64	0.78
Mean prev. unempl. duration | prev. unempl.	255	195	364	213
Mean days unemployed in last 5 years	342	218	710	250
Mean pred. job-finding rate (week 1)	0.615	0.637	0.621	0.626
N	781,944	79,416	129,398	40,080

B.2 Risk Score Accuracy

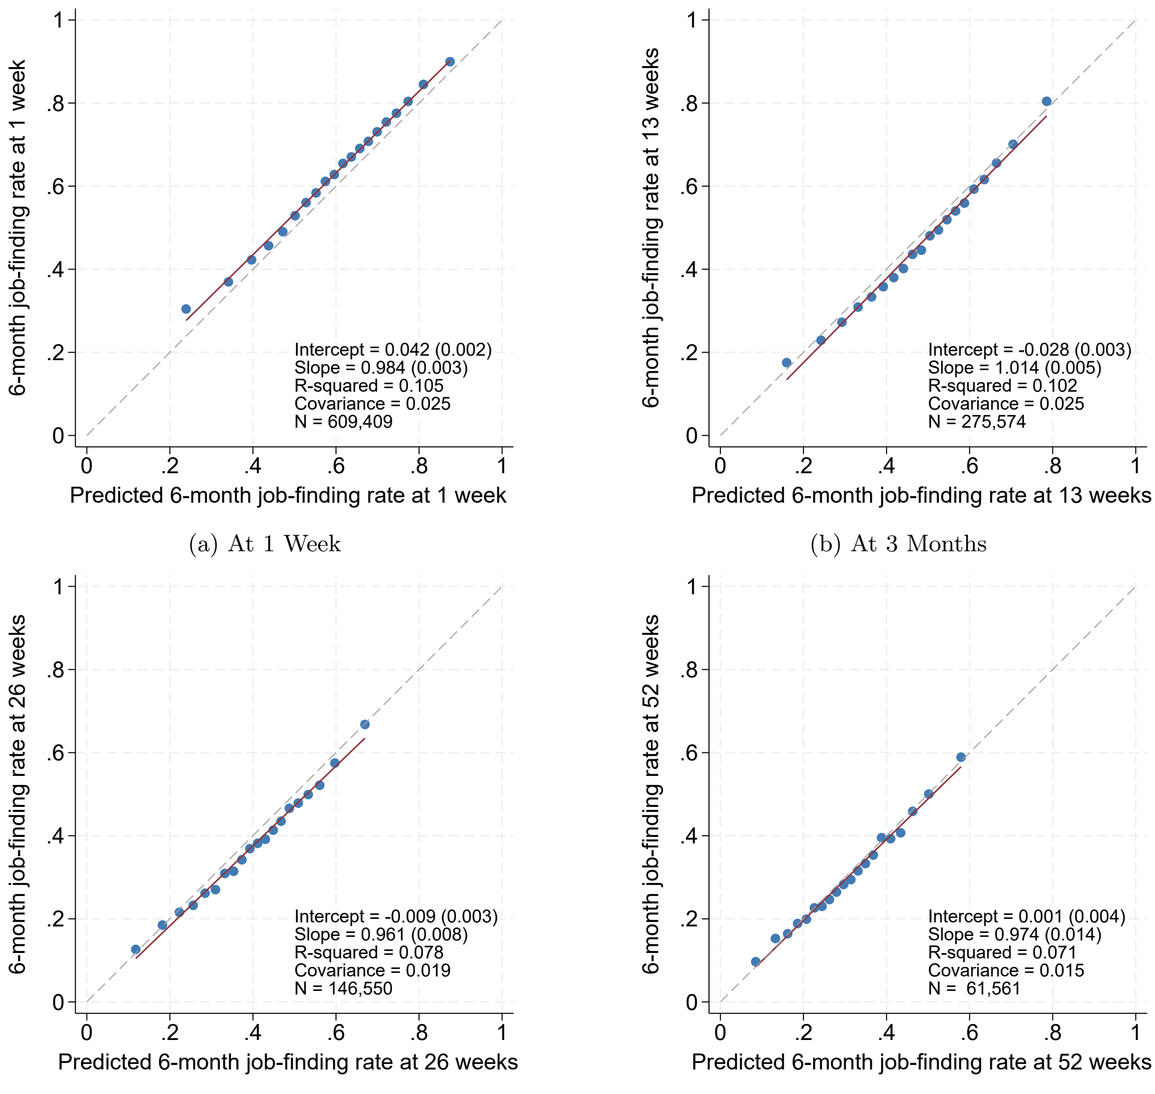

(a) At 1 Week

(b) At 3 Months

(c) At 6 Months

(d) At 12 Months

Figure A7: 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 Baseline RD: Bunching

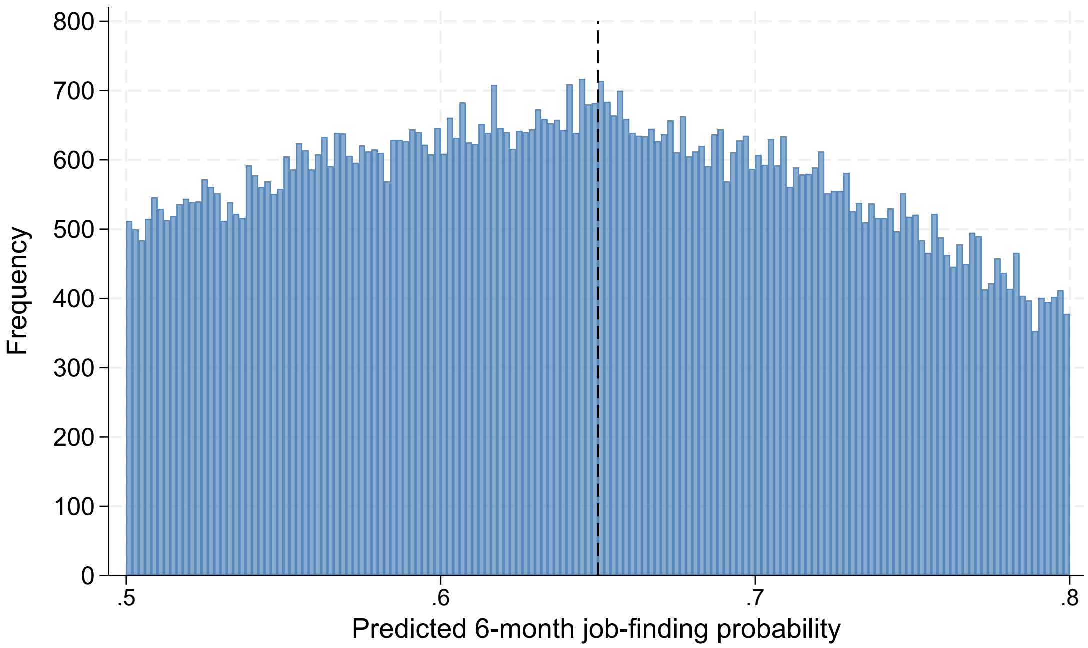

Figure A8: Distribution of Risk Scores

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

C.2 Baseline RD: Kernel and Bandwidth

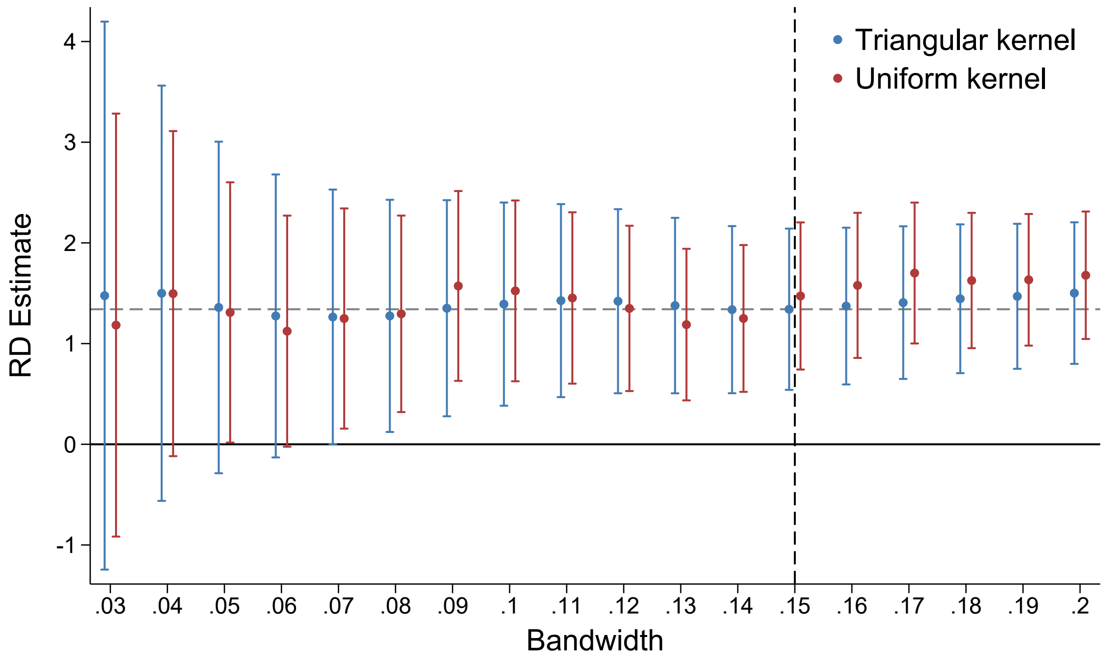

Figure A9: Estimates by Kernel and Bandwidth

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 Baseline RD: Balance Tests

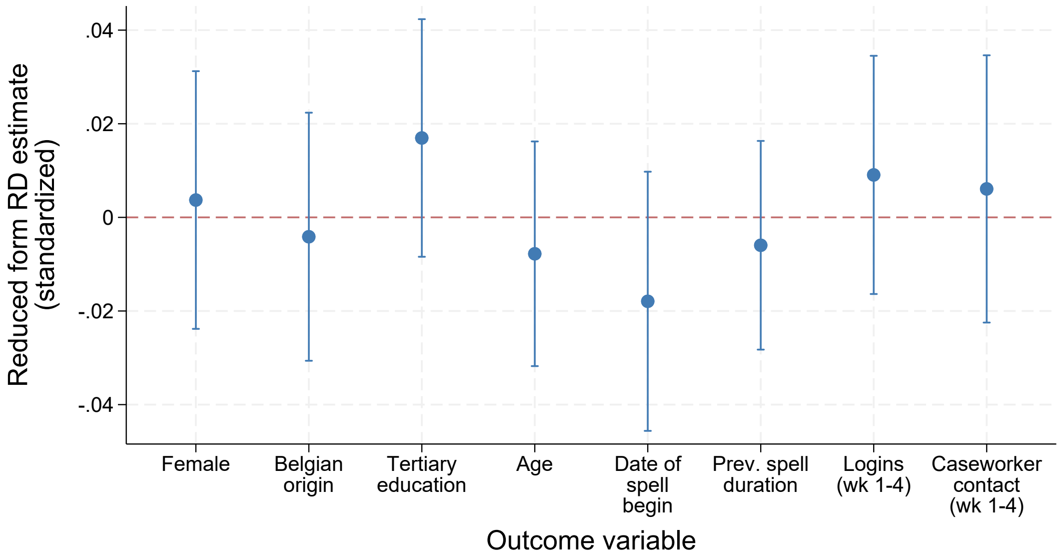

Figure A10: Balance Test

Note: This figure plots reduced form estimates of the effect of being below the cut-off on predetermined covariates. The covariates are all standardized within the estimation sample. We display 95% confidence intervals for the estimated effects.

C.4 Baseline RD: Additional Outcomes

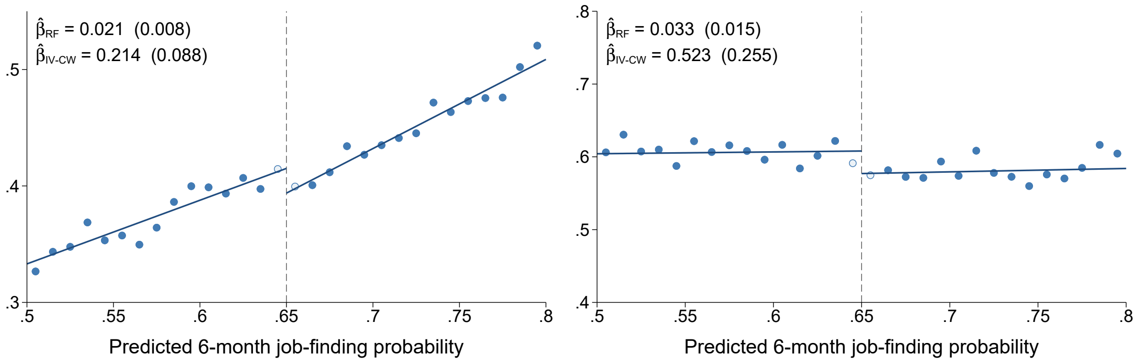

(a) Employment (b) Employment Tenure

Figure A11: Other Employment Outcomes

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 proportion 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

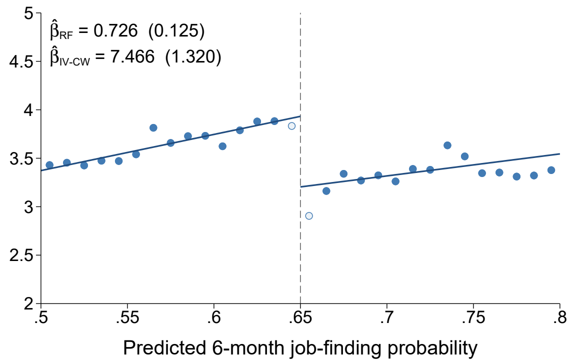

(a) Logins

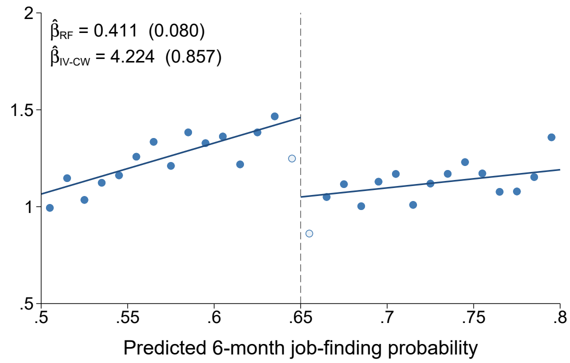

(b) Self-Reported Job Applications

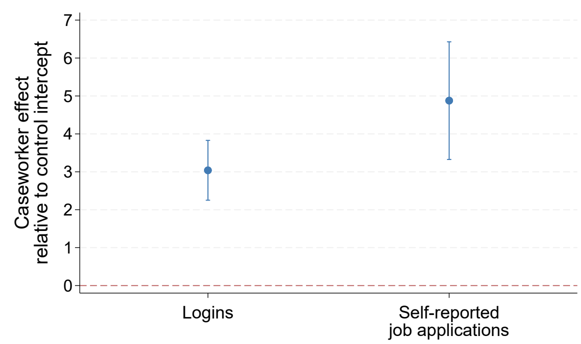

(c) Comparison

Figure A12: Search Behavior

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 average number of logins to the VDAB website in weeks 6–17. Panel (b) plots the average number of job applications reported by job seekers 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. For comparison in panel (c), estimates are divided by the threshold intercept of the local linear fit on the right side of the cut-off. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses. We report 95% confidence intervals.

C.5 Baseline RD: Program Participation

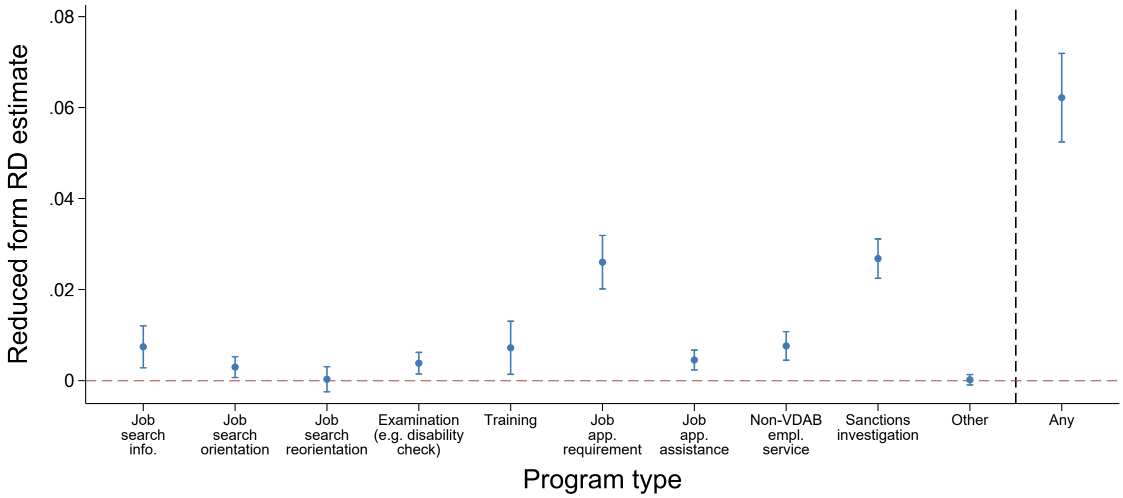

Figure A13: Program Participation (Weeks 6–17)

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. The figure shows the reduced form effect of being below the cut-off on participation in various programs in weeks 6–17 of unemployment. The same controls are used as in the baseline regression. We report 95% confidence intervals.

D Heterogeneity in Effects

D.1 Heterogeneity RD: Bunching

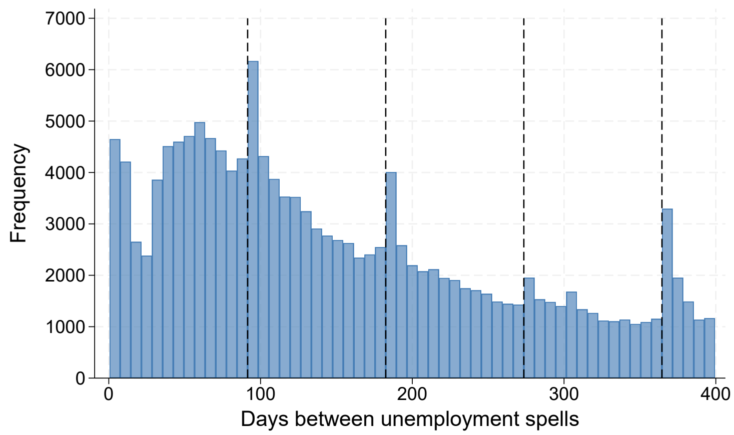

Figure A14: Frequency Distribution of Days Between Unemployment Spells (0–400 Days)

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

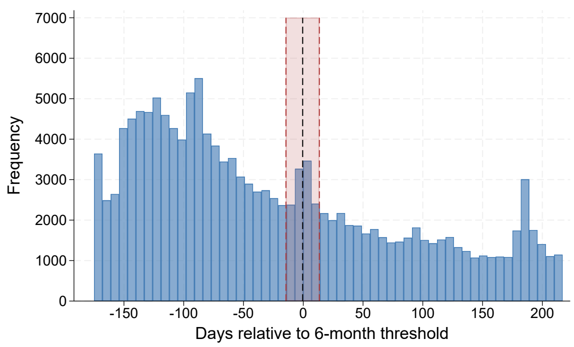

(a) 31 Weeks from the Threshold

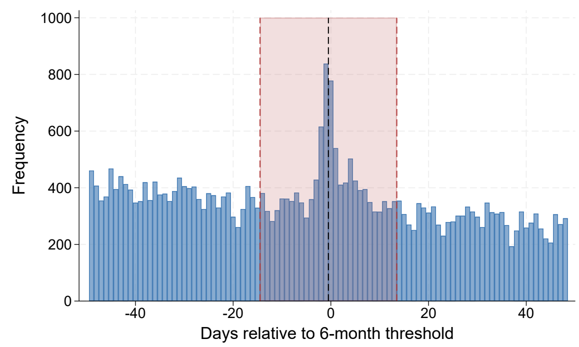

(b) 7 Weeks from the Threshold

Figure A15: Frequency Distribution of Days To 6-Month Threshold

Note: These figures plot the distribution of unemployment spell start dates relative to the day that is 6 months and 1 day after the end of the previous unemployment spell, applying the sample restrictions used in the heterogeneity RD. Panels (a) and (b) apply bin widths of 7 days and 1 day, respectively. The highlighted region (within 2 weeks of the threshold) is the excluded region in the baseline specification.

D.2 Heterogeneity RD: Donut Robustness

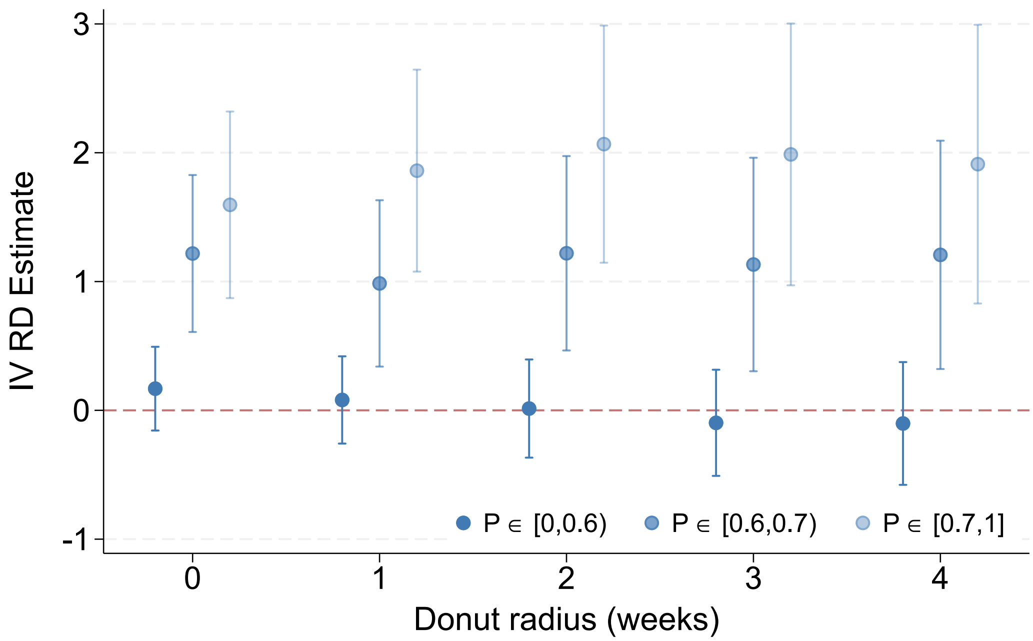

Figure A16: Estimate Robustness to Different Donuts

Note: These panels plot IV estimates of the effect of caseworkers on days worked per week estimated using the heterogeneity RD over weeks 1-4 since the start of unemployment. A two week radius is used for the baseline specification. Robustness is shown to 0–4 week radii. We display 95% confidence intervals. Across all donut specifications, we reject the null hypotheses that the effect is equal in the lowest and highest risk groups with a p-value of at most 0.0009.

D.3 Heterogeneity RD: Kernel and Bandwidth

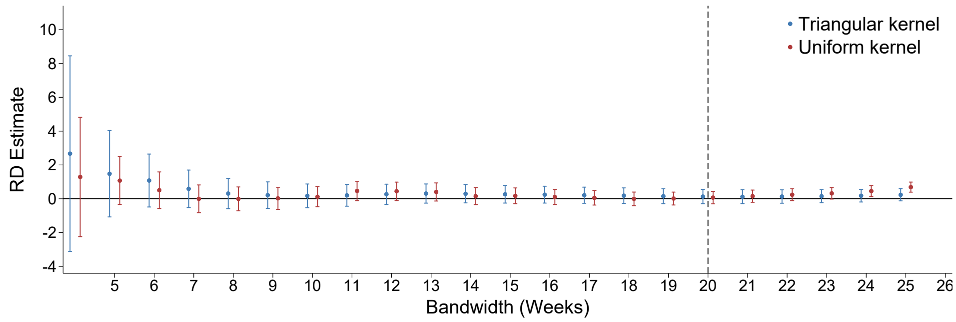

(a) $P \in [0, 0.6)$

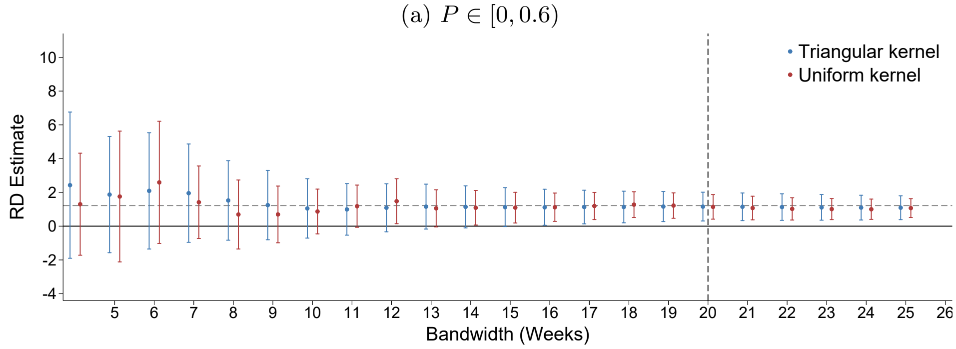

(b) $P \in [0.6, 0.7)$

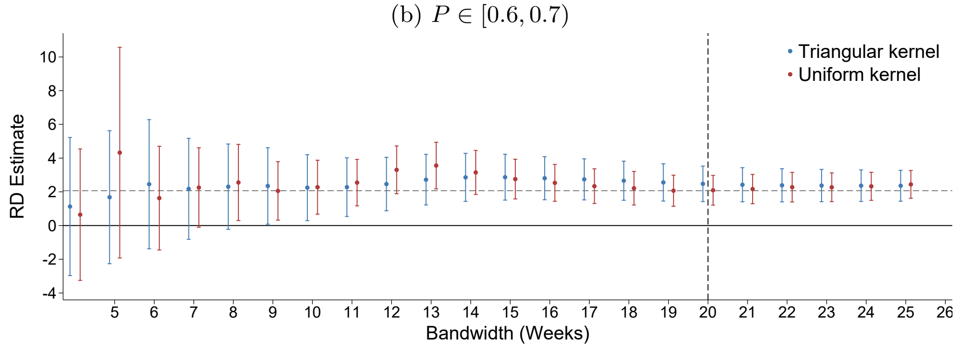

(c) $P \in [0.7, 1]$

Figure A17: Estimates by Kernel and Bandwidth

Note: These panels plot IV estimates of the effect of caseworkers on days worked per week estimated using the heterogeneity 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.

D.4 Heterogeneity RD: First Stage

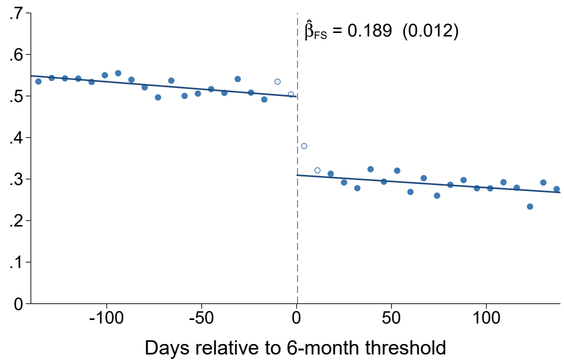

(a) $P \in [0, 0.6)$

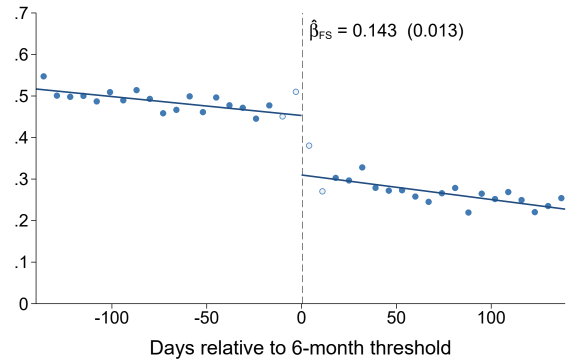

(b) $P \in [0.6, 0.7)$

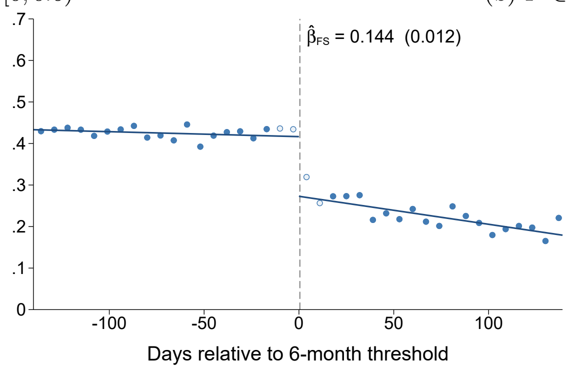

(c) $P \in [0.7, 1]$

Figure A18: 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) plot the proportion of job seekers who have contact with a caseworker in weeks 1–4 of unemployment for each risk group. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by the number of days to the 6-month threshold. 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

D.5 Heterogeneity RD: By Risk Ventiles

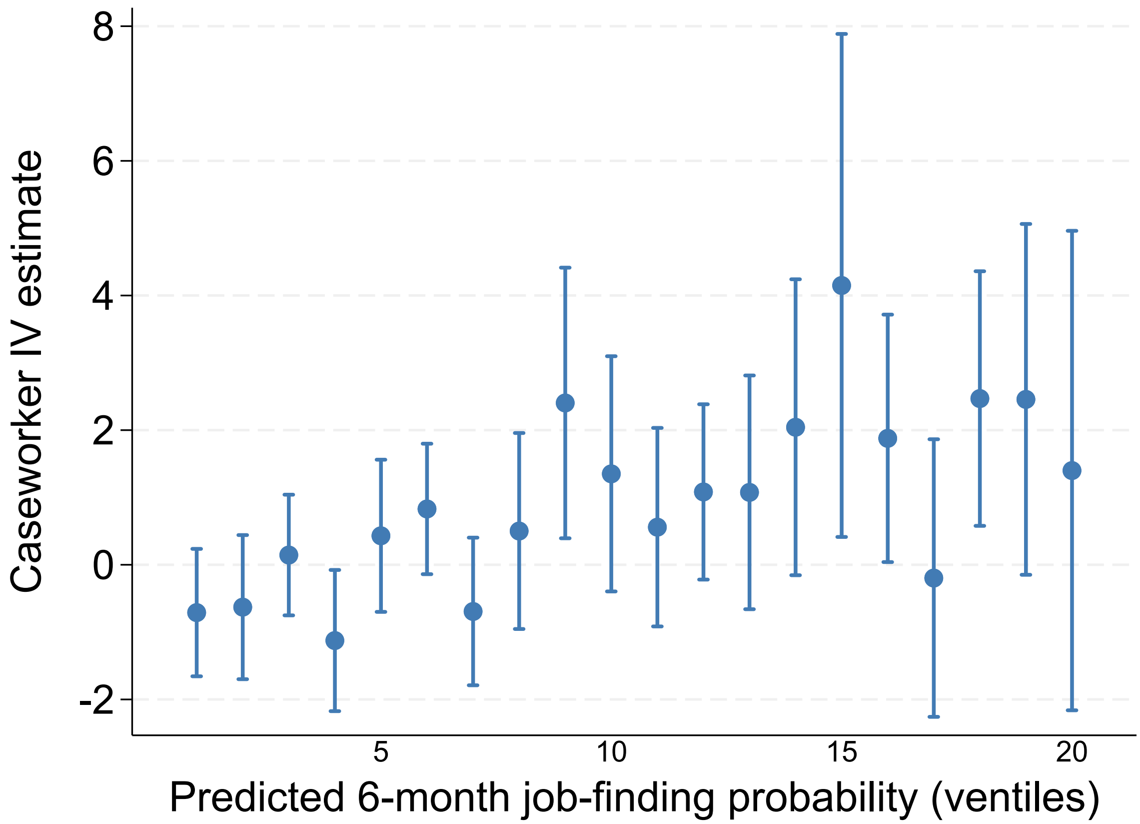

Figure A19: $P \in [0.7, 1]$

Figure A20: Caseworker IV Estimate by Risk Ventile

Note: The sample is restricted to job seekers who had contact with a caseworker in their previous spell. Job seekers are then split into 20 ventiles in the probability of finding a job within 6 months P that was predicted for them after 1 week of unemployment. The plot shows caseworker IV estimates of the effect of contact with a caseworker on days worked per week in weeks 1–4. We report 95% confidence intervals.

D.6 Heterogeneity RD: Other Outcomes

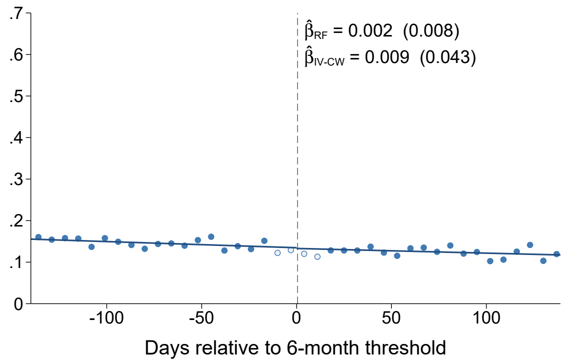

(a) $P \in [0, 0.6)$

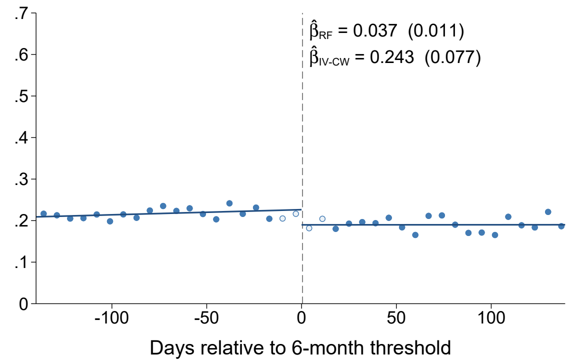

(b) $P \in [0.6, 0.7)$

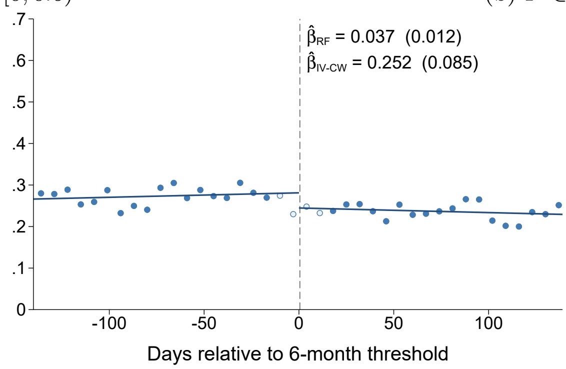

(c) $P \in [0.7, 1]$

Figure A21: 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) plot the proportion of job seekers who are employed four weeks after the start of their unemployment spell. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by the number of days to the 6-month threshold. 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

D.7 Heterogeneity RD: Dynamics

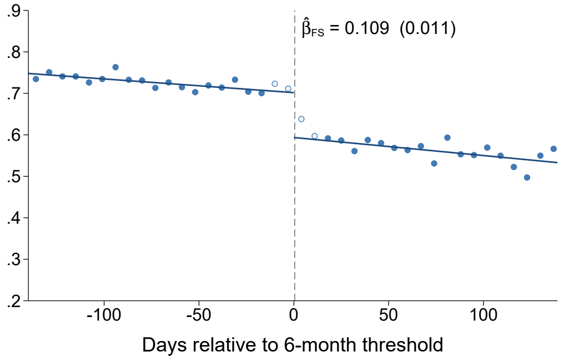

(a) $P \in [0, 0.6)$

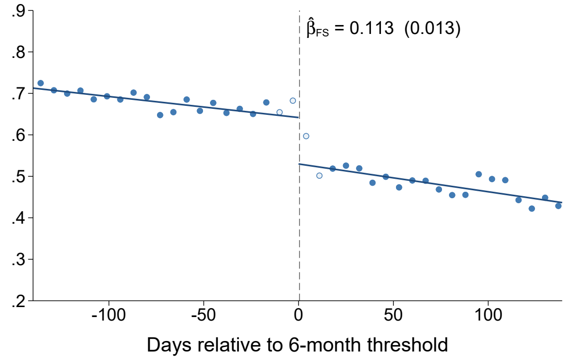

(b) $P \in [0.6, 0.7)$

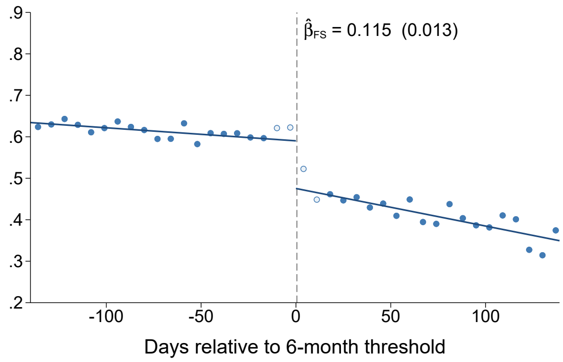

(c) $P \in [0.7, 1]$

Figure A22: 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) plot the proportion of job seekers who have contact with a caseworker in weeks 1–16 of unemployment for each risk group. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by the number of days to the 6-month threshold. 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

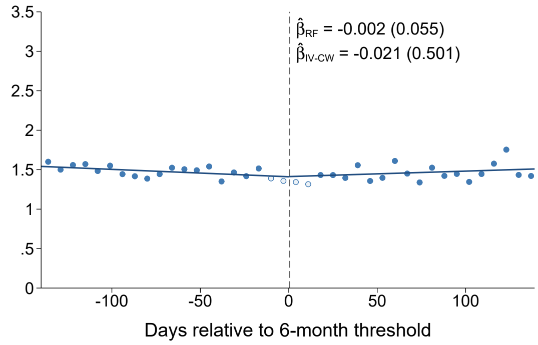

(a) $P \in [0, 0.6)$

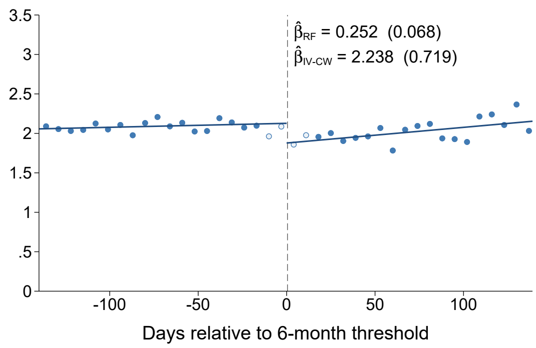

(b) $P \in [0.6, 0.7)$

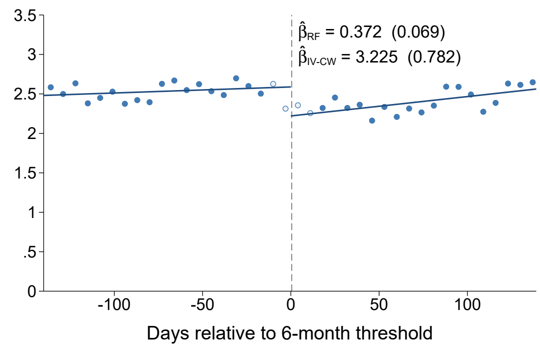

(c) $P \in [0.7, 1]$

Figure A23: 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.8 Heterogeneity RD: Placebo

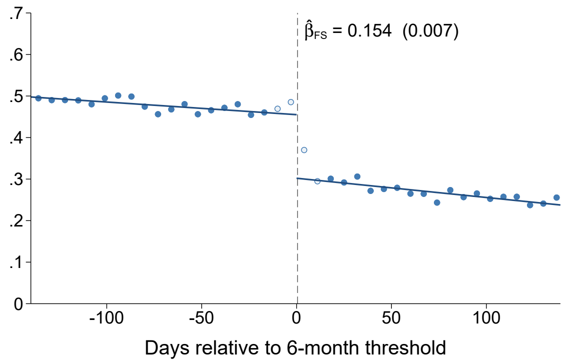

(a) Caseworker Contact (Baseline)

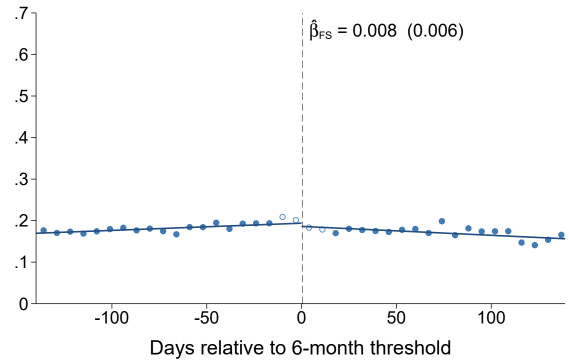

(b) Caseworker Contact (Placebo)

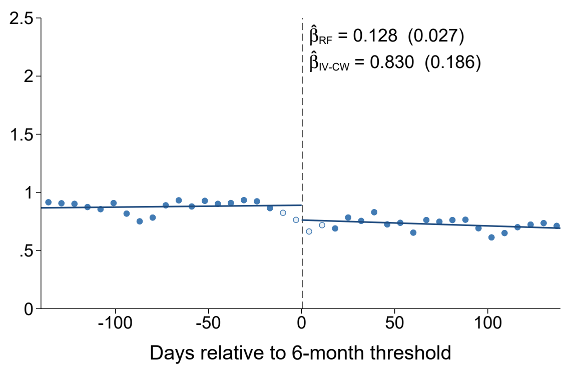

(c) Days Worked (Baseline)

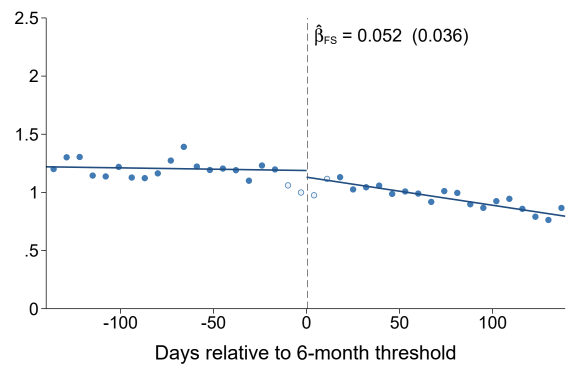

(d) Days Worked (Placebo)

Figure A24: Placebo without Heterogeneity

Note: The baseline sample, used in panels (a) and (c), is restricted to job seekers who had contact with a caseworker in their previous spell. The placebo sample, used in panel (b) and (d), is restricted to job seekers who have not had contact with a caseworker in any spell ending less than a year before the start of the current spell. The outcome variable in panels (a) and (b) is contact with a caseworker in weeks 1–4 of unemployment, while the outcome variable in panels (c) and (d) is days worked per week in weeks 1–4 of unemployment. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by the number of days to the 6-month threshold. 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

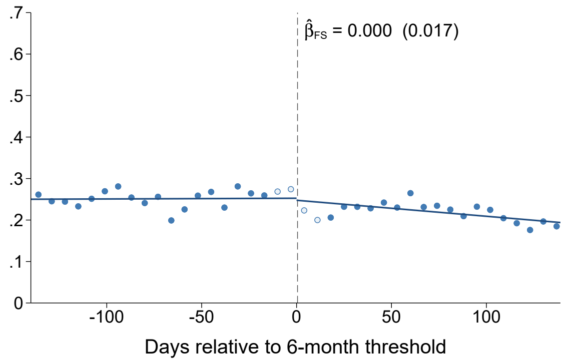

(a) $P \in [0, 0.6)$

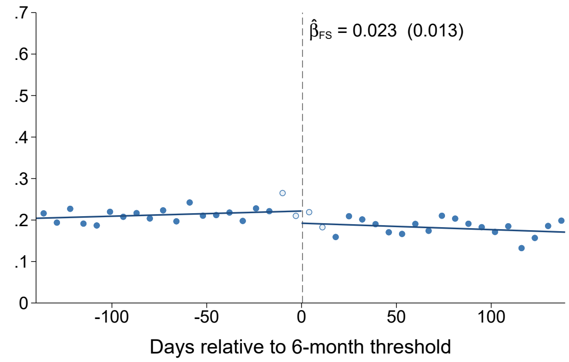

(b) $P \in [0.6, 0.7)$

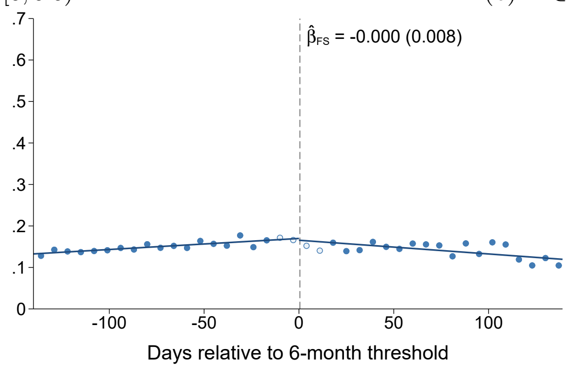

(c) $P \in [0.7, 1]$

Figure A25: Placebo Caseworker Contact (Weeks 1–4)

Note: The placebo sample is restricted to job seekers who have not had contact with a caseworker in any spell ending less than a year before the start of the current 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) plot the proportion of job seekers who have contact with a caseworker in weeks 1–4 of unemployment for each risk group. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by the number of days to the 6-month threshold. 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

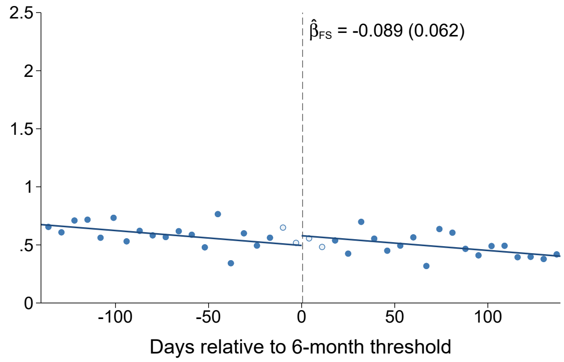

(a) $P \in [0, 0.6)$

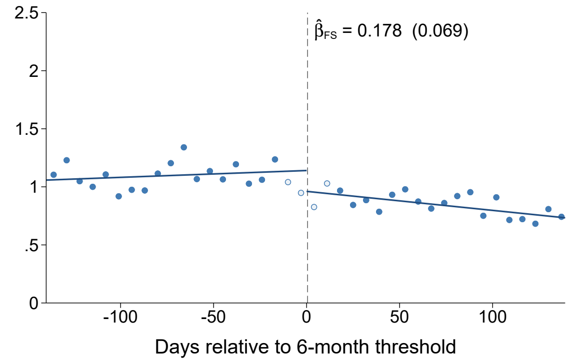

(b) $P \in [0.6, 0.7)$

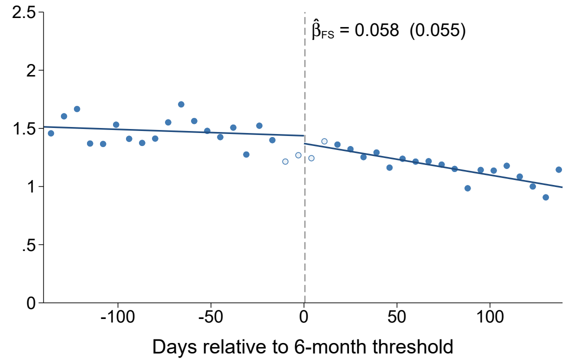

(c) $P \in [0.7, 1]$

Figure A26: Placebo Days Worked per Week (Weeks 1–4)

Note: The sample is restricted to job seekers who have not had contact with a caseworker in any spell ending less than a year before the start of the current 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. The outcome variable is days worked per week in weeks 1–4 of unemployment. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by the number of days to the 6-month threshold. 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. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

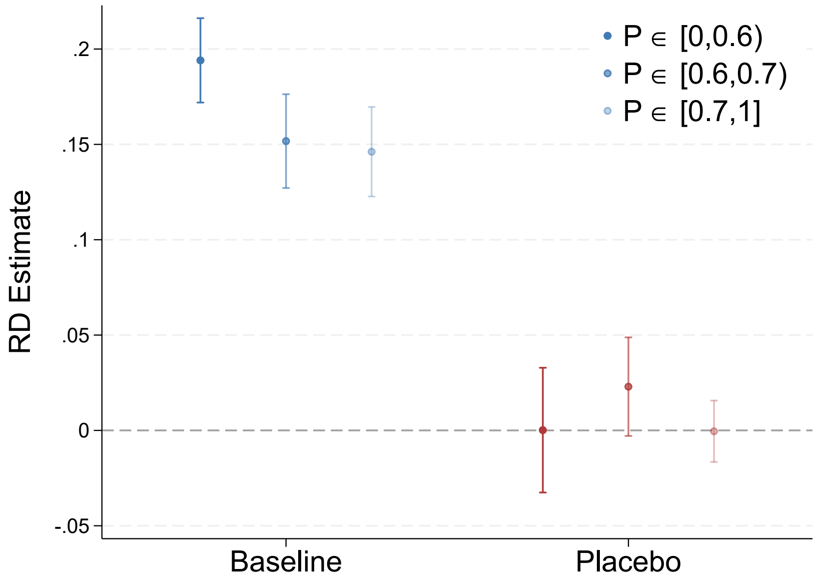

Figure A27: Placebo with Heterogeneity (Caseworker Contact)

Note: The placebo sample is restricted to job seekers who have not had contact with a caseworker in any spell ending less than a year before the start of the current spell. The plot shows the reduced form RD estimate of being below the cut-off on contact with a caseworker in weeks 1–4 of unemployment, with controls 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.

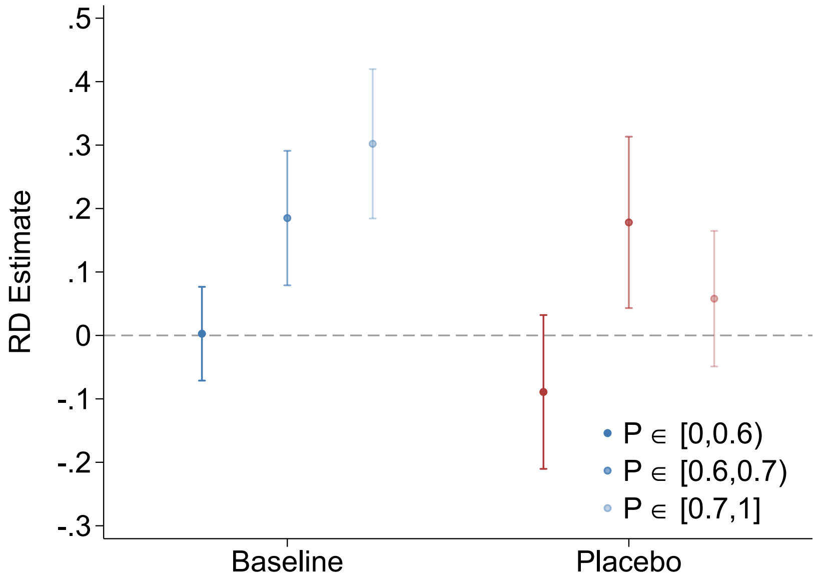

Figure A28: Placebo with Heterogeneity (Days Worked per Week)

Note: The placebo sample is restricted to job seekers who have not had contact with a caseworker in any spell ending less than a year before the start of the current spell. The plot shows the reduced form RD estimate of being below the cut-off on days worked per week in weeks 1–4 of unemployment, with controls 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.

D.9 Heterogeneity RD: Program Participation

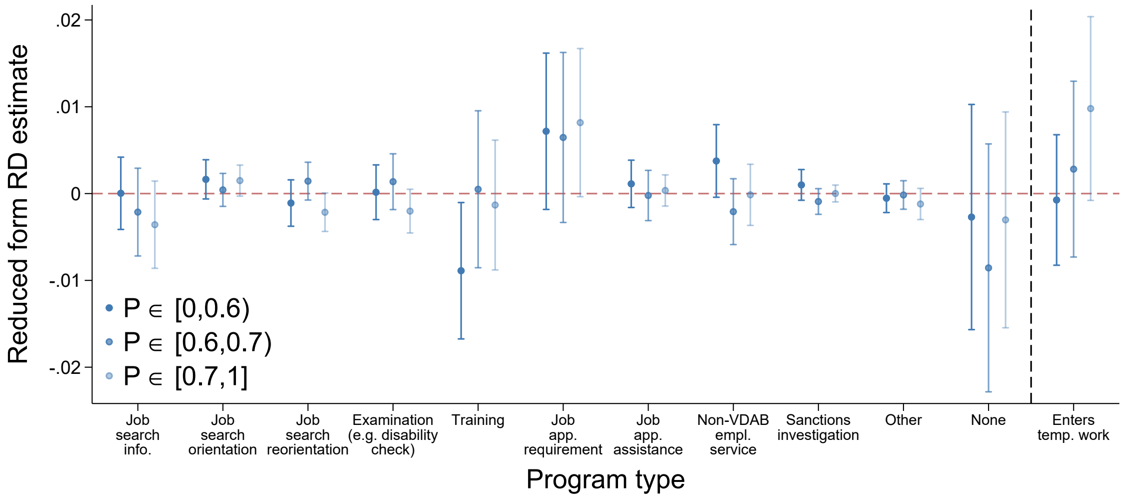

Figure A29: Program Participation

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.

D.10 Descriptive Statistics by Risk Score and CATE

Table A3: Descriptive Statistics by Quadrant

	$P < P_{\text{med}}$		$P \geq P_{\text{med}}$	
	$\hat{\tau} \geq \hat{\tau}_{\text{med}}$	$\hat{\tau} < \hat{\tau}_{\text{med}}$	$\hat{\tau} \geq \hat{\tau}_{\text{med}}$	$\hat{\tau} < \hat{\tau}_{\text{med}}$
Female	32%	46%	32%	48%
Age $\leq$ 25	13%	25%	18%	48%
Age 26–44	44%	58%	54%	48%
Age 45+	43%	17%	28%	5%
Belgian-origin	46%	53%	55%	76%
Tertiary education	17%	13%	14%	16%
Advanced Dutch proficiency	49%	64%	59%	86%
Work disability	10%	11%	9%	8%
Low-employment occupation	16%	38%	20%	32%
Mean duration of prev. spells	368	376	233	166
N	6,984	15,189	15,189	6,984

Note: The sample used to estimate the causal forest is divided into quadrants based on the median risk score and the median CATE. This table reports statistics for each quandrant. "Low-employment occupation" refers to the proportion with an occupational employability index in the lowest quartile. The occupational employability index is calculated in two steps. First, for each occupation, the average 3-month job-finding rate among individuals who list that occupation as a "preferred occupation" is calculated. Then, for each individual, the average is taken of the first step averages of that individual's preferred occupations. "Prev. spell duration" is the mean duration of previous unemployment spells, counting at most five.

E Dynamics

E.1 Dynamics RD: RD Plot

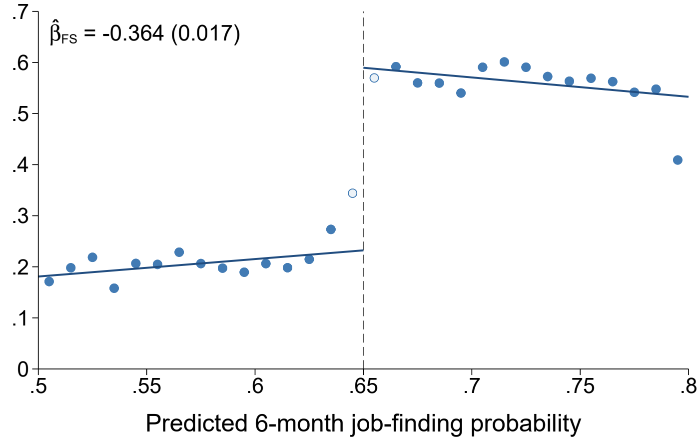

(a) Contact with the service line

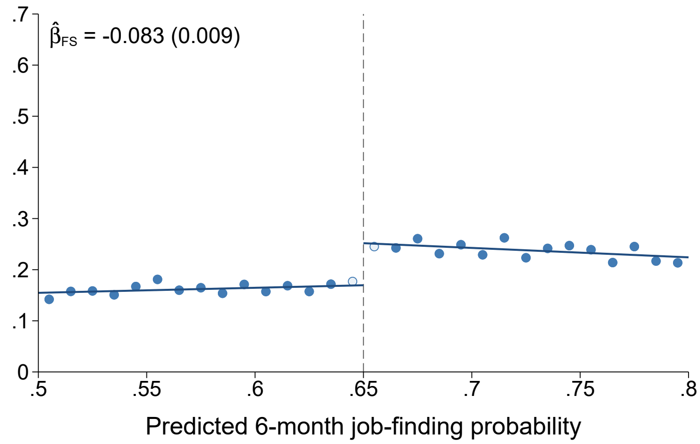

(b) First contact with a caseworker

Figure A30: First stage (weeks 20–31)

Note: The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment and who go on to complete at least 19 weeks of unemployment. The plots show the proportion of job seekers with contact in weeks 20–31 without controls. The reported RD estimates control for the baseline set of controls. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

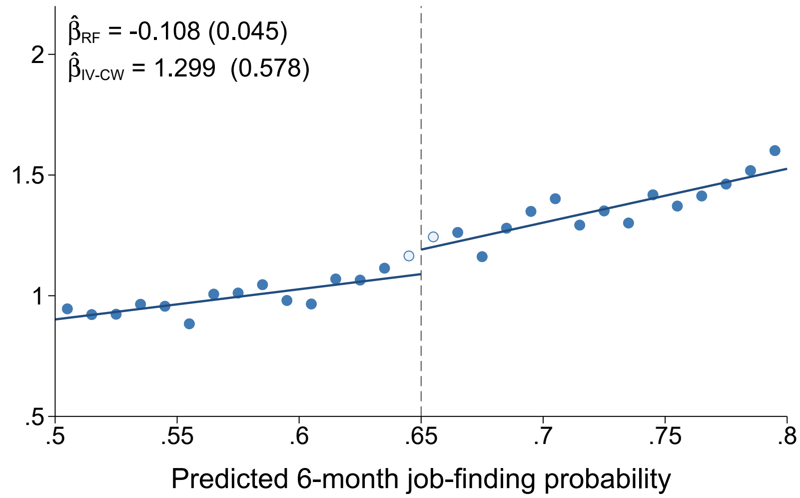

Figure A31: Days worked per week (weeks 20–31)

Note: The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment and who go on to complete at least 19 weeks of unemployment. We plot the proportion of job seekers who have contact with a caseworker in weeks 20–31 of unemployment. For the purpose of plotting, the residuals of the outcome regressed on a set of controls are added to the mean value of the outcome. The mean of the resulting variable is then plotted by risk score. We use the same set of controls as used in the baseline RD. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

E.2 Dynamics RD: Kernel and Bandwidth

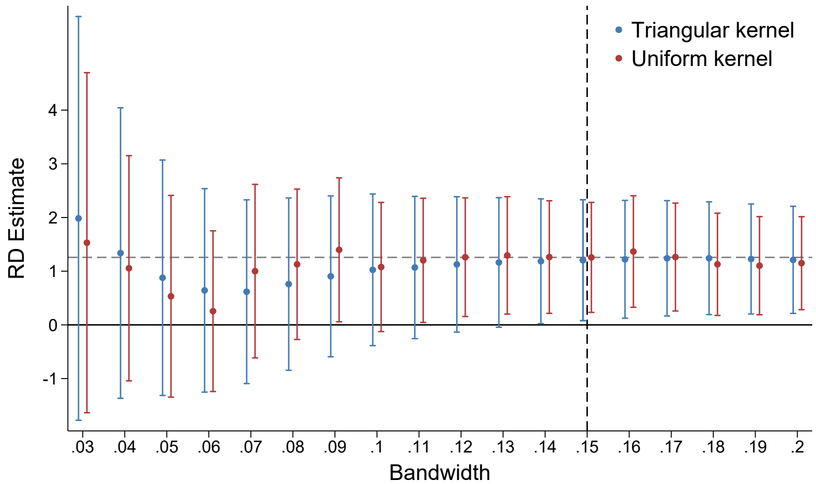

Figure A32: Estimates by kernel and bandwidth

Note: This figure plots IV estimates of the effect of caseworkers on days worked per week estimated using the dynamics RD over weeks 20–31 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.

E.3 Baseline and Dynamics RDs: Alternative Horizons

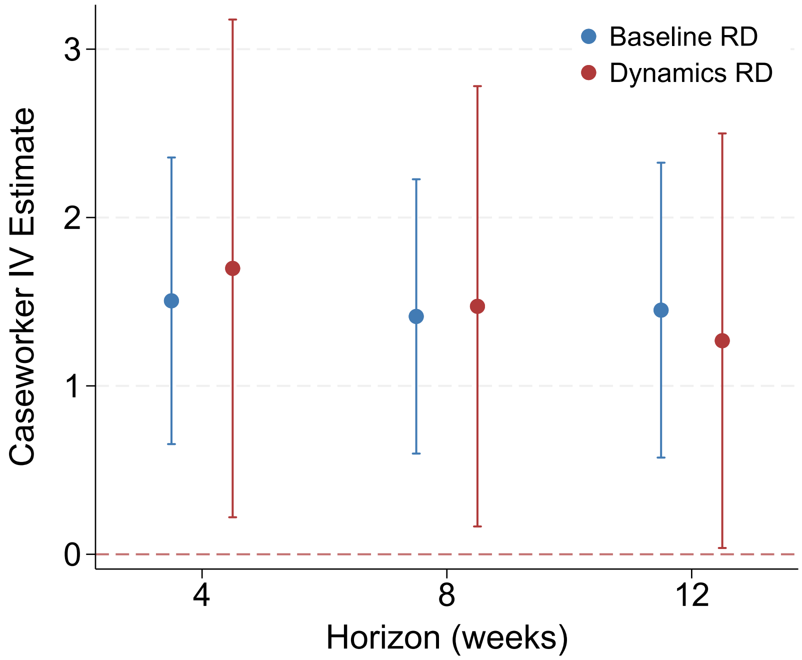

Figure A33: Estimates by horizon

Note: This figure plots IV estimates of the effect of a caseworker on days worked per week estimated using the baseline RD and the dynamics RD. The effects are estimated at three horizons: 4 weeks, 8 weeks and 12 weeks (the baseline). The same specification and control variables are used as for the baseline estimates. We display 95% confidence intervals for the estimated effects.

E.4 Dynamic RD: Balance Tests

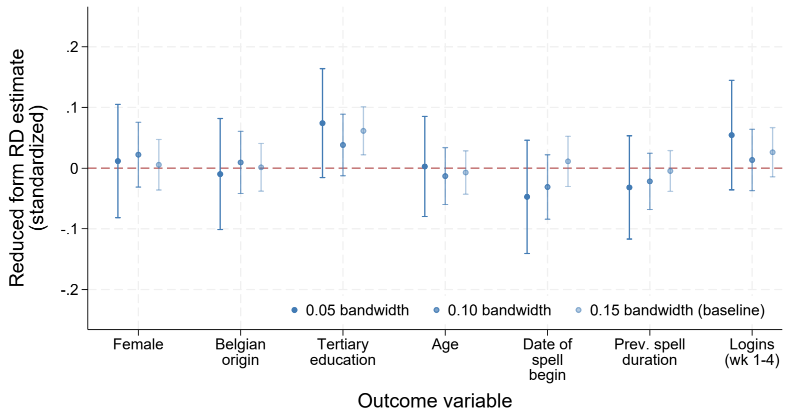

(a) Full Sample

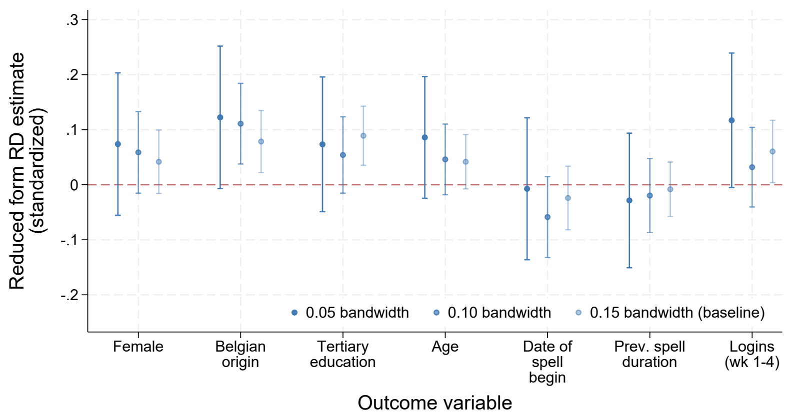

(b) Untreated at Week 20

Figure A34: Balance Test

Note: The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment and who go on to complete at least 19 weeks of unemployment. In panel (b), the sample is additionally restricted to individuals who did not meet with a caseworker in the first 19 weeks. The plots show RD estimates of the effect of being below the cut-off on days worked per week alongside a set of pre-determined covariates. All outcomes are standardized. We show estimates using three bandwidth sizes and plot 95% confidence intervals.

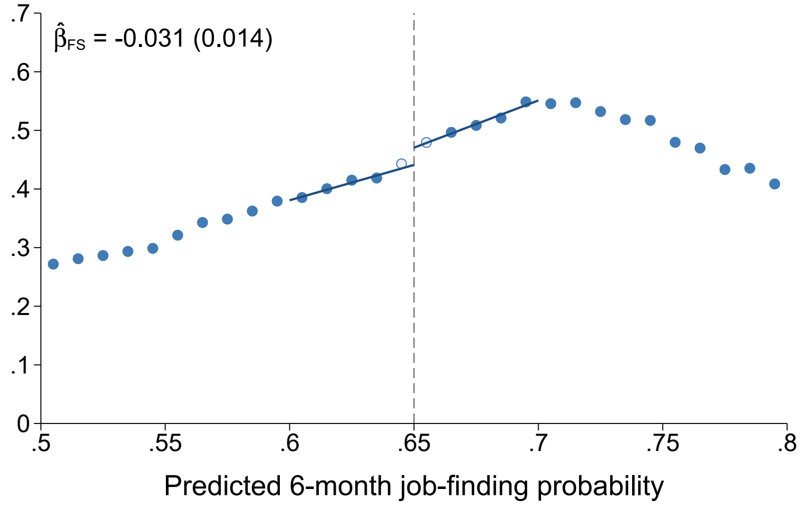

(a) Full Sample

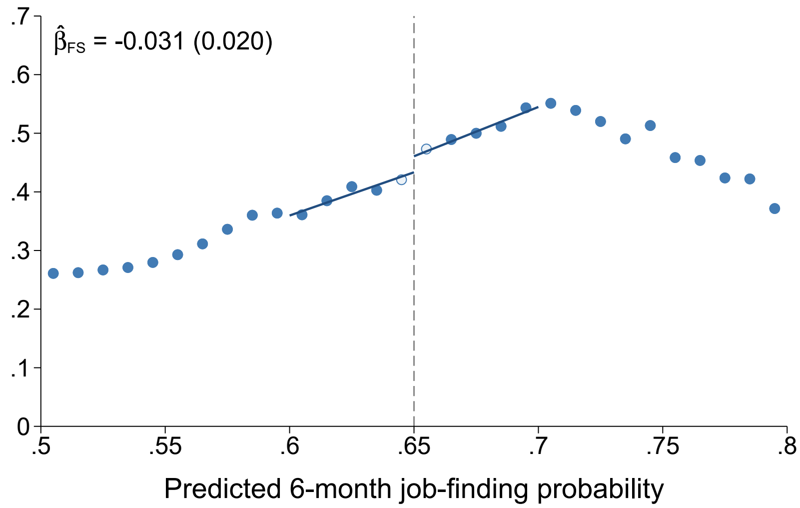

(b) Untreated at Week 20

Figure A35: Balance Test (CATEs)

Note: The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment and who go on to complete at least 19 weeks of unemployment. In panel (b), the sample is additionally restricted to individuals who did not meet with a caseworker in the first 19 weeks. The average predicted CATE is plotted by risk score and RD estimates for the effect of being below the cut-off are reported. We control for the same set of variables as in the baseline regressions. A 0.05 bandwidth is used rather than 0.15 due to clear non-linearities. Standard errors clustered at the individual level are calculated using the nearest-neighbor variance estimator and reported in parentheses.

E.5 Reverse Risk Score RD with Sample Restrictions

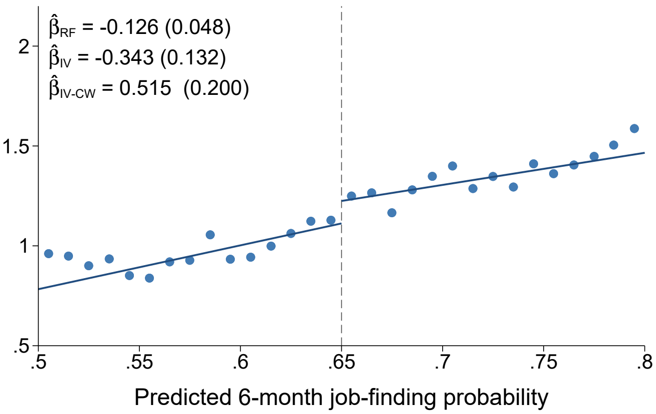

(a) Only Treated Below Cut-off

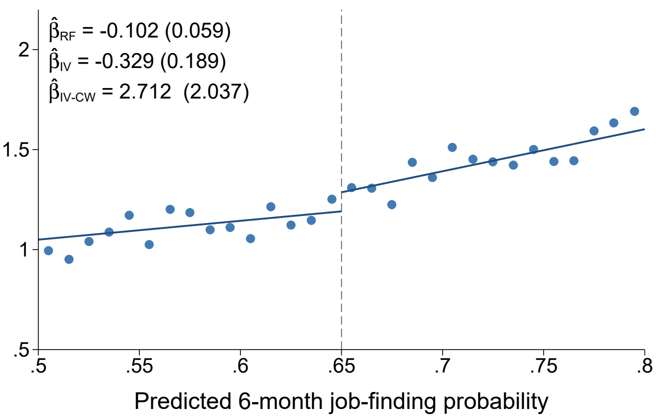

(b) Only Untreated Below & Above Cut-off

Figure A36

Note: The sample is restricted to job seekers for whom the service line receives a risk score after 5 weeks of unemployment and who go on to complete at least 19 weeks of unemployment. In panel (a), an additional restriction is made below the cut-off, keeping only those who did met with a caseworker in the first 19 weeks. In panel (b), the sample is restricted to individuals who did not meet with a caseworker in the first 19 weeks, both below and above the cut-off. The figures then plot the mean days worked per week in weeks 20-31, 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.

E.6 Search Behavior

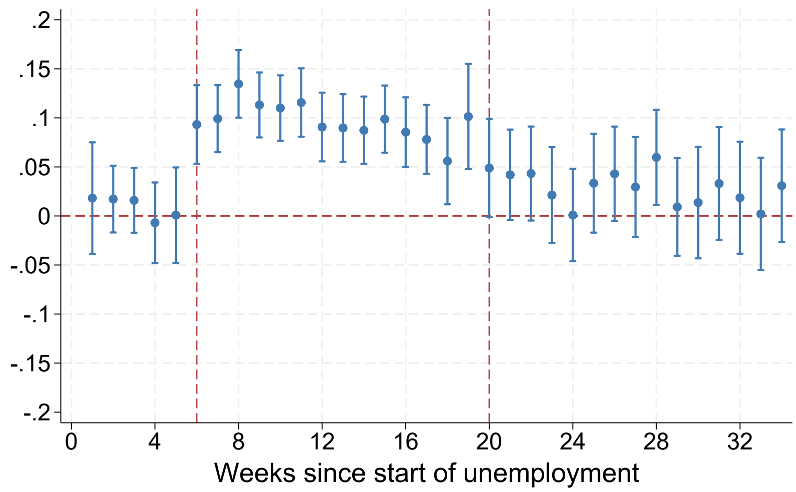

(a) Effect on Logins (Weekly)

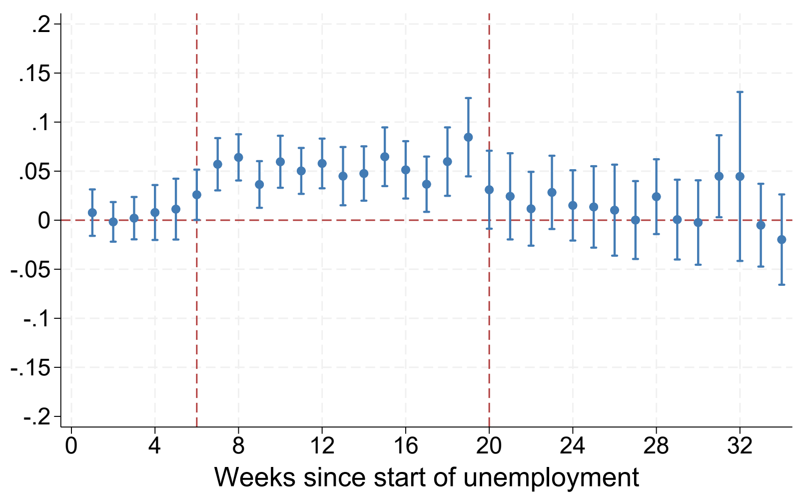

(b) Effect on Self-Reported Job Applications (Weekly)

Figure A37: Search Behavior Dynamics

Note: All estimates use the baseline specification used in the risk score RD. Panels (a) and (b) plot RD estimates of the effect of being below $P = 0.65$ on the number of logins and the number of self-reported job applications in a week, respectively, conditional on being unemployed at the beginning of that week. We display 95% confidence intervals for the estimated effects.