

Does Outsourcing Job Search Assistance help Job Seekers Find and Keep Jobs?*

Lionel Cottier[†] Pierre Kempeneers[‡] Yves Flückiger[§] Rafael Lalive[¶]

March 15, 2016

Abstract

Finding jobs for long-term unemployed job seekers is a key challenge. We study how Hestia, a private job search assistance organization, compares to the public employment service in Geneva. In a randomized field experiment, we follow unemployment benefit recipients two years prior to five years after assignment to treatment and find that Hestia's clients enter jobs earlier, especially during the first six months after the start of the experiment. Yet, the initial employment gain dissipates completely, and becomes even negative and significant two years after assignment. Hestia's clients leave jobs quickly immediately after starting a job, and right after the start of the second year. They also earn less on the job than the control group. All of this suggests that Hestia places job seekers faster but into worse jobs. Despite this, we find that the most pessimistic cost-benefit analysis comes out marginally in Hestia's favor.

JEL codes: J64, J68.

Keywords: job placement, long-term unemployment, active labor-market policy, private sector, public sector.

*We would like to thank seminar audiences at the Institute for International Economics in Stockholm, the Einaudi Institute for Economics and Finance in Rome, the Health and Labor Seminar Series in Lausanne, and especially John Hassler, Luigi Guiso, Claudio Michelacci, Arash Nekoei, Jonathan de Quidt, and Guido Tabellini for comments on this paper. We are grateful to the Caisse de Centrale de Compensation in Geneva, specifically to Alex Pavlovic, and to the Swiss Secretariat of Economic Affairs (SECO), especially Damien Yerly, for supporting this project. Special thanks to Ms Daria Kalyaeva for precious ideas and discussions. This project is part of the NCCR LIVES research program on education and employment for vulnerable people. The authors blame each other for the remaining errors.

[†]University of Lausanne; lionel.cottier@unil.ch

[‡]University of Geneva; pierre.kempeneers@unige.ch

[§]University of Geneva; yves.flueckiger@unige.ch

[¶]University of Lausanne, CEPR and IZA; rafael.lalive@unil.ch

1 Introduction

Long-term unemployment is a pervasive problem in today's rapidly evolving economies. As technology has radically changed the professional activity of most people in the last decades, individuals whose skills are not needed or up-to-date anymore are left behind and those who become unemployed see their skills depreciate even faster. The role of active labor market policies as a tool to correct market failures and improve workers' welfare thus becomes salient. More specifically, outsourcing public placement services to private firms has been tried in various contexts with mixed results. The idea that a private firm facing competition might be able to perform better than its public counterpart sounds plausible to many. Yet outsourcing may also entail important agency problems. Private providers of placement services could select only those job seekers who are easy to place, cream-skimming, or not devote any effort into placing hard-to-place individuals, parking.

In this paper, we study the case of outsourcing job placement of long-term job seekers to a private but non-profit firm in the Canton of Geneva, Switzerland. The goal of this experiment was to determine whether using a private placement provider could help lower the high level of long-term unemployment in Geneva. We study the work trajectories of about 890 individuals gathered in a randomized controlled experiment conducted in 2006-2007, in conjunction with social-security data that cover two years before and five after the experiment. This allows us to understand the ex-ante professional trajectories of these job seekers and to evaluate the long-term impacts of the private placement scheme compared to the public one. Our main analyses look at the employment patterns of individuals and their earnings. First, we use the characteristics of our dataset to classify job seekers in three mutually exclusive "labor market states" in order to isolate the effects of the treatment on the different professional situations in which individuals can be. Second, we analyze the dynamics of the more traditional variables that are unemployment benefits received and wages. Third, we study transitions between employment states from a different perspective, namely using a survival analysis approach. Finally, we compare the costs of the two programs in a cost-benefit analysis.

We find that, the private provider significantly improves employment, by about four percentage points, higher in favor of the private provider in the first year after the experiment started. Unemployment to job transitions are higher in the first six months after assignment for Hestia's clients. These dynamics translate into a 20% decrease in the average amount of unemployment benefits received by the privately placed job seekers in the first year following the experiment, which is in line with the goal of the experiment. However, longer-term results suggest that Hestia focused on placing job seekers as fast as possible at the expense of their suitability to the position. The initial employment gain dissipates and turns even significantly negative two years after job seekers were assigned to Hestia. Employment exits are very high immediately after starting the new job, and immediately after finishing the first year on the job. On top of this, we find that privately placed job seekers earn between 10% and 17% less than their public counterparts between two and three years after the experiment. Overall, we find that outsourcing job placement to a private provider affects

employment and earnings dynamics in the short-run but that the effects can be both positive and negative. They cancel out over a five-year period, leading to employment outcomes that are not different from those of the standard public job placement services.

Our paper ties into the general literature on active labor market policies. Specifically, job placement programs have aspects that are similar to job search assistance but they may also have a threat component. In their survey on ALMP, Card *et al.* (2010) show that job search assistance programs generally have fairly positive effects on employment in the short-run. The medium-run effects of those programs are less well explored. Arni (2015) presents a recent randomized evaluation of a job search assistance program for job seekers aged 45 years or older. Research into the threat effects of active labor market programs shows that the threat of training is often more effective than training itself (Black *et al.*, 2003; Rosholm and Svarer, 2008; Graversen and van Ours, 2009). In our own research, we find that unemployment benefit sanctions increase exits from unemployment but sanctioned individuals fare worse after leaving unemployment (Arni *et al.*, 2013).

Four papers are closely related to our work. Bennmarker *et al.* (2013) study the impact of contracting out job placement to private providers using a field experiment that took place in Sweden in 2007. They randomly assign about 5,000 hard-to-place job seekers to either a private placement agency or to a public one in a trial that lasts between three and six months. The target population is mainly immigrants, adolescent and disabled people. The private provider has strong financial incentives as 60% of the full payment was based on successful placement. They find that compliance with assignment to the private provider was low, 28%, but job seekers received better assistance with the private provider. Despite no significant difference in the probability of employment between public and private providers, the privately placed job seekers have higher work income up to 12 months after the experiment. They also highlight the fact that any differences fade away over time and that there is important heterogeneity across different types of job seekers, which is also an important result of our study.

Krug and Stephan (2013) conduct an empirical comparison between two intensified placement services for hard-to-place unemployed job seekers in Germany. Both a private and a public placement providers aim to improve the employability of their job seekers and place them in stable jobs. They follow individuals up to 18 months after the randomized experiment and conclude that public in-house provision reduces accumulated days in unemployment by one to two months but that about two thirds of this effect is attributable to labor market withdrawals.

Rehwald *et al.* (2015) compare the effectiveness of public and private providers of employment services for highly educated job-seekers (unemployed holding a university degree). Private providers deliver more intense services earlier in the unemployment spell but public and private provision of employment services are equally effective and equally costly from a public spending perspective.

Behaghel *et al.* (2014), who study the differences in impact between private and public provision of counseling to job seekers in a large controlled experiment in France. In a similar setup, they compare a public job placement scheme to its private counterparts and study placement outcomes of about 200,000 job seekers. In their case, the public provider offers both standard and intensive counseling tracks, while the private providers offer intensive counseling only. They find that more intensive counseling leads to faster transition to employment for both programs, but that the public program outperforms the private one in all indicators considered. Also, they find that the private providers' behavior is tightly linked to the incentive scheme in place. Our study differs from theirs in various aspects. First, we consider long-term unemployed individuals while they consider job seekers having a high probability of long-term unemployment. Second, participation to the intensive track is voluntary for their subjects, while ours had no choice but to follow the program to which they were attributed. Third, they focus on short-term effects while we follow all individuals from two years before the experiment to five years after the experiment in order to highlight long-term effects.

Our study contributes to the existing literature in at least two ways. First, we find that a simple model of effort choice by the private provider predicts fairly well the pattern of treatment effects in job exits. We hope to build on this framework to explore how to better provide incentives.¹ Second, we follow individuals for five years after the intervention took place. The long follow up period is required to assess the consequences of private placement on the stability of employment. This long-term approach allows us to capture important effects which take time to materialize.

The rest of the article is organized as follows. Section 2 gives an overview of the context in which the experiment has been implemented, describes the experimental setup, and discusses what outcomes can be expected from such a scheme. Section 3 presents the data, provides descriptive statistics and a descriptive analysis of the main outcome variables. Section 4 explains the empirical framework and presents the results. Section 5 explores further some results using a survival analysis approach. Section 6 provides a cost-benefit analysis. Section 7 concludes the study.

2 Experiment

In this section, we present the environment in which the experiment has been conducted, explain the experimental setup and then discuss the incentives faced by the private job placement provider.

¹Behaghel *et al.* (2014) provide extensive theoretical assessments of financial incentives for the private provider and use them to build predictions.

2.1 Background

In what follows, we explain how the unemployment insurance works in Switzerland and provide more specific details about the labor market in Geneva. This section draws extensively on our earlier work on Swiss active labor market policies (Lalive *et al.*, 2008b, 2005; Arni *et al.*, 2013).

Participants in the experiments were job seekers. The rules concerning benefit eligibility are the same all across Switzerland. Job seekers need to fulfil two requirements in order to be eligible for unemployment insurance benefits. First, they must have paid unemployment insurance taxes for at least 12 months in the two years prior to registering at the public employment service (PES). Job seekers entering the labor market are exempted from the contribution requirement if they have been in school, in prison, employed outside of Switzerland or have been taking care of children. Second, job seekers must possess the capability to fulfill the requirements of a regular job – they must be “employable”. During the unemployment spell, job seekers have to fulfill certain job search requirements and participate in active labor market programs in order to remain eligible for benefits.² Job seekers who are ineligible for unemployment insurance can claim social assistance. Social assistance is means tested and replaces roughly 76% of unemployment benefits for a single job seeker with no other sources of earnings (OECD, 1999).

Job seekers are eligible for 18.5 months of benefit payments during a two-year framework period. Job seekers aged 55 years or older who had contributed for at least 18 months prior to entering unemployment are eligible for two full years. The replacement ratio is 80% for low-income workers (earning less than 3,536 Swiss Francs (SFr) before unemployment).³ The replacement rate is 70% for high-income workers (earning more than 4,340 SFr) and is smoothly adjusted between so that there are no discontinuities in the replacement rate.⁴ Job seekers pay all earnings and social insurance taxes except the unemployment insurance tax rate (which stands at about 2%) so the gross replacement rate is similar to the net replacement rate.

Active labor market policy might vary across Switzerland but the main features of the system are homogeneous. Job seekers are in regular contact with a caseworker at one of about 150 PES offices. When individuals register at the PES office, they are assigned to a caseworker on the basis of either previous industry, previous occupation, place of residence, alphabetically or the caseworker’s availability. Job seekers meet at least once a month with the caseworker. Caseworkers monitor job search by checking that job seekers fill in the details of the jobs to which they have applied. Job seekers are typically required to apply to about ten jobs per month. Caseworkers have some discretion to adjust this target. They count the number of new applications in all cases and they may also check up on the applications claimed by job seekers. Participation

²See Gerfin and Lechner (2002) and Lalive *et al.* (2008a) for detailed background information on and an evaluation of the active labor market programs.

³1 SFr \approx 0.95 EUR \approx 1.05 USD in 2015.

⁴Benefits insure monthly earnings up to a top cap. The cap is currently at 10,500 SFr per month. See Eugster (2013) discusses effects of the benefit replacement rate in the Swiss setting.

in a labor market program is monitored by the caseworker because program suppliers only get paid for the actual number of days a job seeker attends the program.

The experiment is set in the canton of Geneva which consists mainly of the city of Geneva. With about 200,000 inhabitants, Geneva is home to a range of international organizations and features a large financial sector. The experimental context is hence typical of large cities, where the occupational dimension of the labor market plays a larger role than the geographical one. Geneva has higher unemployment than the rest of Switzerland. Around the time of the experiment, in April 2006, 7.1% of Geneva's workers were seeking employment, whereas only 3.5% of workers in all of Switzerland were seeking employment. Altogether, the Geneva labor market poses a challenge for workers wanting to re-enter a job. Each canton is in charge of organizing its active labor market policy and Geneva is known to rely less on unemployment benefit sanctions than the nation and somewhat more on training programs.

2.2 Experimental design

The evaluation of the public and private placement schemes is based on a randomized controlled experiment which took place in the canton of Geneva, Switzerland, between 2006 and 2007. The goal of this experiment was to determine whether using a private placement provider could help lower the comparatively high level of long-term unemployment in Geneva.

The private but non-profit firm, called "Les Maisons Hestia" (henceforth Hestia) offered a tailor-made monitoring program which aimed to improve some of the job seekers' skills such as job interview training and resume writing. Hestia's goal was to focus on finding jobs that fit well with their candidates' needs and profiles so as to maximize their chances of success in their new jobs and make sure they reach stable positions. The public placement service offered a more standard program aiming to help and encourage the unemployed in their job search and application process without playing a truly active role in it. As a means of comparison, Hestia had a team of 3 consultants plus two telemarketing operators for a pool of about 100 job seekers, while the public scheme has one or two employees per 100 job seekers on average (Flückiger and Kempeneers, 2008).

The partnership with Hestia was a "pilot project" monitored by the PES in Geneva. It had a fixed duration of one year but could potentially be extended to a contract of indefinite duration in case of positive results, which was Hestia's goal. On top of the long-term incentives linked to the potential contract renewal, Hestia was provided with short-term incentives in the form of a payment scheme. Hestia did not receive money upon placement but instead was given a decreasing monthly flat rate per job seeker enrolled. It received 1000 SFr⁵ a month per job seeker enrolled in their program for six months or less; 500 SFr per job seeker already enrolled in their program for six to 18 months; and 350 SFr per job seeker enrolled for more than 18

⁵ 1 SFr \approx 0.62 EUR \approx 0.8 USD in 2006.

months. Thus, the total amount received by Hestia for a job seeker that stays one full year in the program is 9000 SFr. The idea behind this payment scheme was to encourage Hestia to place job seekers rather fast but also to give money in the medium-run in order to avoid a potential “parking problem”, that is, providing minimal effort to the harder-to-place job seekers.

The experiment was targeted to a specific group of job seekers. The cantonal PES defined a population of reference composed of long-term unemployed people that meet the following criteria: registered at the cantonal unemployment office for at least one year before September 2006, starting month of the experiment period; eligible to unemployment benefits and/or eligible to labor market measures; not on medical leave or maternity leave; and not currently involved in any cantonal program for criminals. These criteria ensured that people in the population of reference were immediately available to participate in the experiment.

Job seekers were then randomly selected each month from the reference population of unemployed and attributed either to the private placement firm Hestia (treated) or to the the standard public placement program (control). Importantly, job seekers did not have the possibility to change their treatment status. Like other field experiments, not all individuals attributed to the treatment group actually got the treatment. We observe that about half the individuals sent to the treatment group did not actually participate to the experiment. One might think that this rather low participation rate indicates a serious problem of compliance, which will bias our results. However, it is actually not due to compliance problems but rather to organizational problems when setting the experiment, screening the potential subjects and splitting them into groups. All individuals leaving the experiment were asked about their reason for not following the treatment and the main reasons were: 22% had subsidized employment; 17% had already found a job; 13% were following another labor-market measure. A myriad of other reasons were also given but refusal to follow the intensive placement program account for only 2% of the total. We show in Section 3.1 that the difference in characteristics between the control and the treated group are not statistically significant. We keep the job seekers assigned to treatment but not following it in all our analyses.

The first cohort entering the experiment saw 355 job seekers being split between the private provider (60%) and the public one (40%), while the following nine cohorts were designed to ensure that Hestia had at least 100 job seekers to take care of. In other words, cohorts two to ten were just “top-ups” to maintain Hestia working at full capacity while ensuring that the allocation between the two programs remains random. The reason why the first cohort exceeds Hestia’s capacity is that the organizers anticipated the potential problem of non-compliance, thus setting a larger first wave. Note also that the upper limit of 100 job seekers was self-imposed by Hestia to ensure they could provide high-quality services.

2.3 Theoretical predictions

In this section, we come back to the incentives facing Hestia, discuss how they could lead to specific behaviors, and what differences in outcomes can be expected between treatment and control groups.

Hestia faces two types of incentives. In the short- and medium-run, it receives money in exchange of providing services to job seekers. In the long-run, it hopes to be able to convert the “pilot project” into a contract of indefinite duration and secure its collaboration with the PES.

Consider first the short- and medium-run incentives. The monthly flat fee per job seeker enrolled is monotonically decreasing. Its highest value is in the first six months of the experiment. However, it never goes down to zero and its minimum value (from 18 months onwards) remains high enough to cover administrative costs and minimal services to job seekers. Such a scheme raises two major concerns. First, a profit-maximizing firm could enroll as many job seekers as possible, and focus exclusively on the easier-to-place ones in order to pocket the high fee in the first months and receive new job seekers to place. Second, it could provide minimal services to the harder-to-place individuals, keeping them for the extra cash that they bring in the long run.

The former problem is similar to that of cream-skimming, where a provider chooses the individuals that it wishes to enroll based on their potential. The latter problem refers to parking, where the private firm offers minimal services to the harder-to-place job seekers. This financial scheme would thus not be very satisfying if it were not accompanied by a condition on the maximum number of job seekers that Hestia can have. Namely Hestia’s capacity must not exceed 100 job seekers at a time. This upper limit has two effects. On the one hand, it ensures that Hestia’s staff is large enough to provide optimal services to all their job seekers. It does not force Hestia to provide the same services across individuals with different placement prospects but it is a good step in the direction of limiting the parking problem. On the other hand, it prevents Hestia from focusing only on easy-to-place job seekers. Assuming that the proportions of easy-to-place job seekers in the population of unemployed is $1/3$ and constant (for the sake of this example), focusing heavily on these individuals is not sustainable in the long-run. Indeed, after placing all easy-to-place individuals of the first cohort, the second cohort will only replace $1/3$ of them with new easy-to-place job seekers, while the rest of the top-up will have lower placement prospects. Thus, the share of harder-to-place individuals enrolled will be growing over time, until the point where Hestia has only hard-to-place individuals. This example is of course extreme but it illustrates well the problem of focusing only on a certain type of individuals while having capacity constraints. This effect should both limit the problem of cream-skimming and parking. Additionally, note that Hestia was not able to fully cream-skim by choosing whom to enroll, since job seekers were randomly attributed to one of the two programs by the PES. Altogether, short-run incentives should then improve short-run return to employment (or alternatively exit from UB) and even out the differences in outcomes (e.g. UB received, employment rate, etc.) across job seekers’ types.

Let us focus now on the long-run incentives. Hestia's goal is to secure its place in the market for services to job seekers in Geneva. To achieve it, it needs to show that it can fulfill its goals and outrun the PES when it comes to providing service to long-run job seekers. This brings us to Hestia's mission, which is to lower the long-term unemployment rate by improving job re-entry and providing job seekers with jobs that fit their needs and are stable over time. Its main force compared to the PES is a much higher staff to job seekers ratio, which allows pro-active services such as calling firms to find unadvertised job vacancies. Overall, long-run incentives should thus materialize in the form of a higher employment rate and more stable jobs over time.

2.4 Incentivized job placement model

The idea behind this section is to provide a simple theoretical model to better understand the specific implications of the financial scheme in place for our experiment. Recall the two key aspects of the financial scheme: (i) decreasing monthly flat rate paid to the private provider (Hestia) for each individual enrolled in its programme; (ii) the private provider has a capacity constraint and each month receives a new inflow corresponding to the number of individuals that it placed in the month before.

Define ω as the “wage” paid to Hestia each period for each job-seeker enrolled, P as the number of job-seekers enrolled in the private programme at any given time (i.e. the capacity constraint), $\delta \in (0, 1)$ as a discount factor applied to the wage so as to mimic the decreasing payment scheme, and $e_t \in (0, 1)$ as the effort level put in by Hestia in order to place its clients, which yields a quadratic cost equal to $\frac{\omega P}{2} e_t^2$. Note that it is assumed that the effort level e_t put in by Hestia corresponds to the placement rate that it achieves for this period. In other words, if it exerts $e_t = 0.3$, it places 30% of its job seekers and it will receive the equivalent number of “fresh” job seekers in the next period. The incentives for Hestia come from the fact that it receives more money per new comer than per job seeker that was already enrolled, due to the decreasing payment. In the first period of the experiment ($t = 0$), Hestia's profit is thus:

$$\pi_0 = \omega P - \frac{\omega P}{2} e_0^2 = \omega P \left[1 - \frac{1}{2} e_0^2 \right] \quad (1)$$

while in the second period of the experiment, it is given by:

$$\pi_1 = \omega P [e_0 + \delta(1 - e_0) - \frac{1}{2} e_1^2] \quad (2)$$

where the fraction of new comers e_0 brings in more money than the fraction of individuals that were not placed $(1 - e_0)$ since δ is smaller than one. In the third time period, the profit becomes:

$$\pi_2 = \omega P [e_1 + \delta(1 - e_1)e_0 + \delta^2(1 - e_1)(1 - e_0) - \frac{1}{2} e_2^2] \quad (3)$$

This logic remains the same for all future time periods, knowing that today's effort level influences tomorrow's profit. To see how the effort levels evolve over time, let's consider a simple version of this game. Assume that it has three time periods, starting with $t = 0$, and ending with $t = 2$. This version can be seen as a representation of the three-stage financial scheme used in the experiment. Hestia's maximization problem is then given by:

$$\begin{aligned} \mathcal{L} = \max_{e_0, e_1, e_2} \sum_{t=0}^2 \beta^t \pi_t &= \omega P[1 - \frac{1}{2}e_0^2] + \beta \omega P[e_0 + \delta(1 - e_0) - \frac{1}{2}e_1^2] \\ &+ \beta^2 \omega P[e_1 + \delta(1 - e_1)e_0 + \delta^2(1 - e_1)(1 - e_0) - \frac{1}{2}e_2^2] \end{aligned} \quad (4)$$

where $\beta \in (0, 1)$ is the standard discount factor for future time periods. The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial e_0} : e_0 = \beta[1 - \delta] + \beta^2 [\delta(1 - e_1) - \delta^2(1 - e_1)] \quad (5)$$

$$\frac{\partial \mathcal{L}}{\partial e_1} : e_1 = \beta [1 - \delta e_0 - \delta^2(1 - e_0)] \quad (6)$$

$$\frac{\partial \mathcal{L}}{\partial e_2} : e_2 = 0 \quad (7)$$

Solving for e_0 , e_1 and e_2 yields the optimal effort levels:

$$e_0 = \frac{1 - \delta^2}{1 + \delta(1 - \delta)} \quad (8)$$

$$e_1 = \frac{1 - \delta^2}{1 + \delta(1 - \delta)} \quad (9)$$

$$e_2 = 0 \quad (10)$$

A few points are worth noting. First, neither the wage paid to Hestia nor the number of job seekers enrolled play any role in this setting. What matters is the speed at which the payment scheme decreases. Second, the effort levels e_0 and e_1 are decreasing in δ :

$$\frac{\partial e_0}{\partial \delta} = \frac{\partial e_1}{\partial \delta} = -\frac{\delta^2 + 1}{(1 + \delta(1 - \delta))^2} < 0 \quad (11)$$

This makes sense as an increase in δ means a lower decrease in the payment scheme. In financial terms, unplaced job seekers are worth more in the future than they used to be, thus reducing the need for a significant effort to place them in the present. If $\delta = 0$, i.e. if the payment scheme is just a flat rate until the end of times, effort levels fall to zero. This corresponds to a situation where the placement provider always receives a fixed amount of money to take care of job seekers but without any incentives to place them. In the Swiss system, this situation corresponds to the Public Employment Services. Third, there is no reason for the

private provider to put in any effort in the last period of the experiment ($t=T$) since it is costly and it will not reap the future benefits of it. Let us now consider a more general version of this game. After normalizing ωP to one, the general form of Hestia's profit function is:

$$\pi_t = \begin{cases} 1 - \frac{1}{2}e_0^2 & \text{for } t = 0 \\ \sum_{H=1}^t \left(\delta^{H-1} \frac{e_{t-H}}{1-e_{t-H}} \prod_{h=1}^H (1 - e_{t-h}) \right) + \delta^t \prod_{h=1}^t (1 - e_{t-h}) - \frac{1}{2}e_t^2 & \text{for } t > 0 \end{cases} \quad (12)$$

Hestia's maximization problem can then be written as:

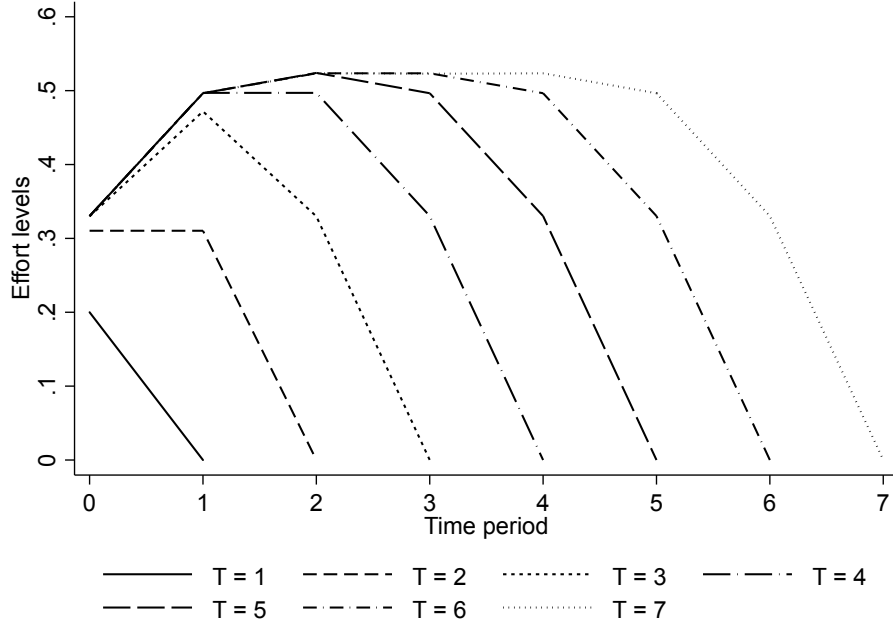
$$\mathcal{L} = \max_{e_t, \dots, e_T} \sum_{t=0}^T \beta^t \pi_t \quad (13)$$

where $\beta \in (0, 1)$ is the standard discount factor for future time periods. The maximization problem yields the following first order conditions:

$$e_t = \begin{cases} \beta(1 - \delta) \left[1 + \sum_{H=1}^{T-t-1} \left(\beta^H \delta^H \prod_{h=1}^H (1 - e_{t+h}) \right) \right] & \text{for } t = 0 \\ \beta \left[1 + \sum_{H=1}^{T-t-1} \left(\beta^H \delta^H \prod_{h=1}^H (1 - e_{t+h}) \right) \right] \times \\ \left[1 - \sum_{H=1}^t \left(\delta^H \frac{e_{t-H}}{1-e_{t-H}} \prod_{h=1}^H (1 - e_{t-h}) \right) - \delta^{t+1} \prod_{h=1}^t (1 - e_{t-h}) \right] & \text{for } t = 1, \dots, T-2 \\ \beta \left[1 - \sum_{H=1}^t \left(\delta^H \frac{e_{t-H}}{1-e_{t-H}} \prod_{h=1}^H (1 - e_{t-h}) \right) - \delta^{t+1} \prod_{h=1}^t (1 - e_{t-h}) \right] & \text{for } t = T-1 \\ 0 & \text{for } t = T \end{cases} \quad (14)$$

In other words, Hestia's effort level is a function of the two discounts rates, as well as past and future effort levels. In the initial period ($t = 0$), it is a function of future efforts only, while in the second last period ($T - 2$), it is a function of past efforts. The closed form solutions for optimal effort levels becomes exponentially difficult to compute as the number of time periods increase and does not add much to the mechanisms already observed in the three period model. However, we provide numerical solutions calculated for a given δ and various T in Figure 1.

Figure 1: Simulation of optimal effort levels, $\delta = 0.8$



Notes: Figure 1 shows the values of the optimal effort levels for a fixed δ and a different number of time periods. $T = 1$ means that the private provider's contract is not renewed after time period 1, and thus that the game only lasts two period ($t = 0$ and $t = 1$).

For each duration of the game, we observe that the optimal effort levels are increasing until the median period and then decreasing in a symmetrical fashion. As already mentioned above, the private provider does not exert any effort in the last period of the game as it would just reduce its immediate profit without improving future benefits. This implies that cohorts which entered the program in the second half of it are worse off compared to those which were enrolled at the start or in the middle. In Section 5.3, we get back to the model and show how empirical estimates compare to the theoretical predictions.

3 Data

For the purpose of this study, we use two different types of data. The first type are the data collected in the experiment described in Section 2.2, while the second type are administrative data about the individuals who took part in the experiment.

The data specific to the experiment contains various socio-demographic and job related variables, such as gender, marital status, education, age, residence permit, and placement prospects. The variable “Placement prospects” is an indicator created by a PES caseworker when a new job seeker enters the unemployment database. It groups the job seekers into four categories according to their personal and professional background: excellent placement prospects means that the job seekers does not need any help for finding a new

job; good placement prospects indicates that the job seekers needs very little support; average placement prospects means that the job seeker’s background is not as good as the first two categories and/or that the individual suffers from lower than average professional qualification; and poor placement prospects, meaning that on top of having a weaker background, the individual may lack professional qualifications and/or even base qualification (e.g. poor education). The creation of this variable relies on objective measures linked to the job seekers abilities and experience, but can also incorporate soft-information gathered by the consultant in charge of this person, such as past placement of similar job seekers, current labor market situation, observed motivation, etc. Note that this variables had been created before allocating job seekers into the control and treatment group, and is thus available for everyone in our sample.

We use administrative data from the Social Security Administration (SSA) to track labor market histories of job seekers in the experiment. The main purpose of this data is to keep track of labor market participation in order to assess old age or disability pension eligibility which depends on social security contributions. Both firms and unemployment insurance agencies inform the SSA every year about total earnings and start and end month of a spell of employment or unemployment because the SSA levies taxes on earnings from employment and unemployment. From this raw data, we construct a detailed monthly calendar that spans two years prior to the experiment and up to five years after it. In each month, we have information on whether the individual receives any earnings from employment (regardless of whether this is self-employment or salaried employment) or from unemployment benefits. We use this information below to characterize whether someone is employed without unemployment benefits, receiving unemployment benefits, or neither of the two. We centered individuals’ administrative record data around the time when they enter into the experiment. In all table and figures, $t = 0$ corresponds to the month when the individual was assigned to Hestia’s placement service or to the control placement service.

3.1 Descriptive statistics

Table 1 shows descriptive statistics on the variables specific to the experiment. The table also indicates the number of individuals allocated to the control group and to the treatment group (Hestia).

We have roughly 50% of men and women in both groups and about half of the job seekers are married. 25% of individuals have between one and three year of work experience, while more than 50% benefit from more than three years of experience. Around 10% of the sample is below 25 years old, 15% is above 55, and the rest is evenly spread in between. The highest education achieved is compulsory schooling for 40% of job seekers, secondary education (e.g. high-school) for 36%, and tertiary education (e.g. university level) for 20%. Half the sample has a Swiss citizenship and a third has a permanent residence permit. Finally, around 60% of job seekers are reported to have good or excellent placement prospects, 20% average placement prospects, and 20% poor placement prospects. All characteristics are equally balanced in both groups,

Table 1: Summary statistics

Variable	Control Group		Hestia Group		Difference (%)	t-stat
	Mean (%)	s.e.	Mean (%)	s.e.		
Women	50.5	0.03	48.8	0.02	1.7	0.50
Marital status						
Single	34.9	0.02	33.2	0.02	1.7	0.53
Married	50.8	0.03	53.5	0.02	-2.7	-0.80
Widower	0.5	0.00	0.4	0.00	0.1	0.30
Divorced	13.8	0.02	12.9	0.01	0.9	0.37
Experience						
None	1.3	0.01	2.5	0.01	-1.2	-1.33
Less than 1 year	7.1	0.01	8.6	0.01	-1.5	-0.80
1-3 years	25.4	0.02	26.0	0.02	-0.6	-0.20
More than 3 years	55.8	0.03	50.8	0.02	5.0	1.49
Age						
17-24	8.2	0.01	10.4	0.01	-2.2	-1.10
25-34	27.5	0.02	31.2	0.02	-3.7	-1.21
35-44	27.5	0.02	27.0	0.02	0.6	0.19
45-54	21.4	0.02	18.9	0.02	2.5	0.91
55-64	15.3	0.02	12.1	0.01	3.2	1.38
Schooling						
Compulsory	40.2	0.03	42.6	0.02	-2.4	-0.71
High-school level	36.0	0.02	35.9	0.02	0.0	0.01
University level	19.8	0.02	18.4	0.02	1.5	0.55
Workers						
Swiss	52.1	0.03	49.8	0.02	2.3	0.68
C permit	30.2	0.02	29.3	0.02	0.9	0.28
Other	17.7	0.02	20.9	0.02	-3.2	-1.19
Placement prospects						
Excellent	4.5	0.01	6.2	0.01	-1.8	-1.16
Good	55.0	0.03	52.9	0.02	2.1	0.62
Average	23.0	0.02	20.9	0.02	2.1	0.75
Poor	17.5	0.02	19.9	0.02	-2.5	-0.93
Number of observations	378	-	512	-	-	-

Notes: Table 1 presents summary statistics on the variables specific to the experiment. Some categories may not add up to 100% due to missing observations. The sixth column calculates the difference between control and treatment groups, defined as control minus treatment. The seventh column reports two-sided t-statistics on the differences.

which reflects the random allocation of individuals. This is confirmed by the sixth column, which indicates the difference between the control and treatment groups, and the seventh column, where we report t-tests on the differences.

Table 2 describes the same data as the one used for Table 1 but we now split up the treatment group (Hestia Group) into individuals who received the treatment and those who did not. We observe some important differences with respect to work experience and age. Treated individuals had more work experience (56% with three years or more compared to 46 % in the non-treated group) than non-treated individuals. Treated individuals were also older than the non-treated individuals (17 % treated individuals aged 55-64 years vs 7% in the non-treated group). This is not surprising as the pilot program carefully screened job seekers to focus only on those that were immediately ready to start a job.

In what follows, we will keep anyone allocated to the treatment in that group and report intention-to-treat (ITT) effects. ITT analysis includes all individuals previously assigned to the treatment. It ignores withdrawals, protocol non-compliance, or more generally everything that happens after the random allocation. ITT results are usually more conservative because of the dilution of the treatment effects due to withdrawals and non-compliance.

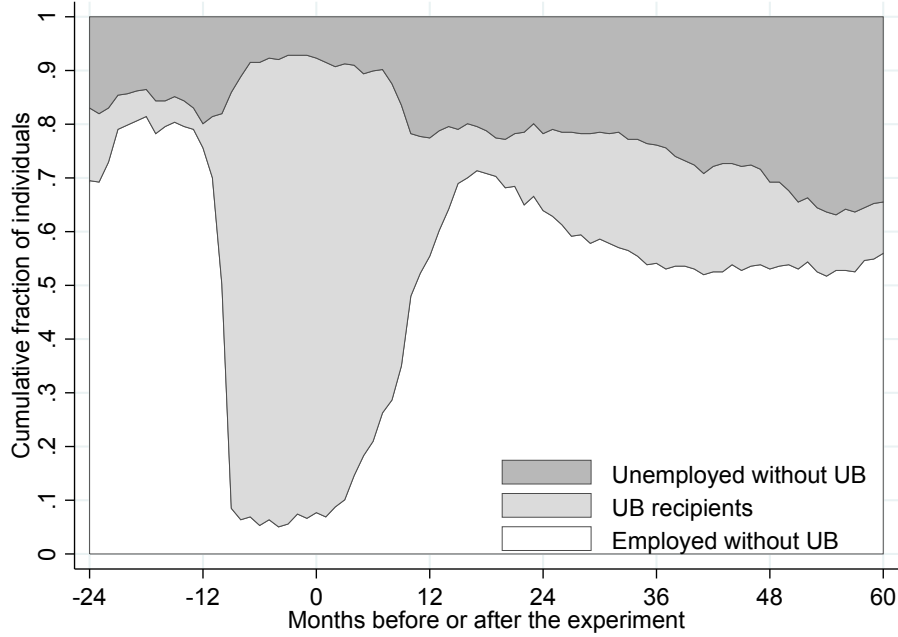
Our main analyses look at the monthly employment patterns of individuals. We classify job seekers in three mutually exclusive “labor market states” in order to isolate the different effects of the treatment. The states are: employed, defined as having a positive income from work, and receiving no unemployment benefits; receiving unemployment benefits (with and without work income); and unemployed but without unemployment benefits (i.e. dependent on social assistance or another insurance scheme). These states represent any situation in which an individual can be at a given point in time. They can be seen as a snapshot of the employment situation of all the individuals in our sample for a given month. Figure 2 presents the dynamics of these three outcome variables for the whole sample.

Table 2: Summary statistics

Variable	Hestia Group									
	Control Group (1)		Treated Group (2)		Non-treated Group (3)		Difference between groups			
	Mean (%)	s.e.	Mean (%)	s.e.	Mean (%)	s.e.	(1)-(2)	t-stat	(1)-(3)	t-stat
Women	50.5	0.03	49.6	0.03	48.0	0.03	0.9	0.23	2.5	0.62
Marital status										
Single	34.9	0.02	30.8	0.03	35.7	0.03	4.2	1.10	-0.8	-0.20
Married	50.8	0.03	55.0	0.03	52.0	0.03	-4.2	-1.05	-1.2	-0.29
Widower	0.5	0.00	0.4	0.00	0.4	0.00	0.1	0.27	0.1	0.24
Divorced	13.8	0.02	13.8	0.02	11.9	0.02	-0.1	-0.03	1.9	0.68
Experience										
None	1.3	0.01	1.2	0.01	4.0	0.01	0.2	0.19	-2.6	-1.94
Less than 1 year	7.1	0.01	6.9	0.02	10.3	0.02	0.2	0.11	-3.2	-1.36
1-3 years	25.4	0.02	23.8	0.03	28.2	0.03	1.6	0.45	-2.8	-0.77
More than 3 years	55.8	0.03	55.8	0.03	45.6	0.03	0.1	0.01	10.2	2.51
Age										
17-24	8.2	0.01	10.0	0.02	10.7	0.02	-1.8	-0.77	-2.5	-1.04
25-34	27.5	0.02	31.5	0.03	31.0	0.03	-4.0	-1.09	-3.4	-0.93
35-44	27.5	0.02	21.9	0.03	32.1	0.03	5.6	1.62	-4.6	-1.24
45-54	21.4	0.02	19.6	0.02	18.3	0.02	1.8	0.56	3.2	0.98
55-64	15.3	0.02	16.9	0.02	7.1	0.02	-1.6	-0.53	8.2	3.32
Schooling										
Compulsory	40.2	0.03	44.6	0.03	40.5	0.03	-4.4	-1.10	-0.3	-0.07
High-school level	36.0	0.02	33.1	0.03	38.9	0.03	2.9	0.76	-2.9	-0.74
University level	19.8	0.02	18.5	0.02	18.3	0.02	1.4	0.44	1.6	0.50
Workers										
Swiss	52.1	0.03	50.8	0.03	48.8	0.03	1.3	0.33	3.3	0.81
C permit	30.2	0.02	31.2	0.03	27.4	0.03	-1.0	-0.27	2.8	0.76
Other	17.7	0.02	18.1	0.02	23.8	0.03	-0.4	-0.11	-6.1	-1.83
Placement prospects										
Excellent	4.5	0.01	8.1	0.02	4.4	0.01	-3.6	-1.79	0.1	0.08
Good	55.0	0.03	53.5	0.03	52.4	0.03	1.6	0.39	2.6	0.65
Average	23.0	0.02	19.6	0.02	22.2	0.03	3.4	1.04	0.8	0.23
Poor	17.5	0.02	18.8	0.02	21.0	0.03	-1.4	-0.44	-3.6	-1.11
Number of observations	378	-	260	-	252	-	-	-	-	-

Notes: Table 2 presents a breakdown of the summary statistics on the variables specific to the experiment. Some categories may not add up to 100% due to missing observations. Columns seven to twelve calculate the difference between control and treatment groups, defined as control minus treatment and report two-sided t-statistics on the differences.

Figure 2: Labor market states



Notes: Figure 2 shows the labor market states in which individuals can be at a given point in time. These three states are mutually exclusive. The figure can be seen as a snapshot of the employment situation of all the individuals in our sample for a given month.

We observe that about 80% of the sample has a paid job and no unemployment benefits one year before the start of the experiment. The remaining 20% are mostly neither getting an income from work nor unemployment benefits, which means that these individuals are most likely covered either by the disability insurance, or by social assistance, or in full-time education. We see that most individual lose their job twelve months before the experiment. This is due to the requirement (when setting the experiment) that the target group is composed of individuals that have been unemployed for at least one year. The organizers chose to select mainly “newly” long-term unemployed to avoid heterogeneity in unemployment histories. One year after the start of the experiment, 70% of individuals have found a job but almost 25% have left the labor force. This share keeps increasing as time goes by, to reach around 40% five years after the experiment.

3.2 Descriptive analysis

In what follows, we analyze work trajectories of the treatment (Hestia) and the control groups. We start by analyzing the effects on labor market participation and then effects on income from two sources: unemployment insurance or paid employment.

3.2.1 Labor market states

Figure 3 shows the fraction of individuals who have a work income and no unemployment benefits (top), who receive UB (middle), and who have no work income but do not receive unemployment benefit either (bottom). These three categories are meant to describe the standard cases of a regular employee who works for her living, that of an individual receiving UB because she is unemployed or in a subsidized job, and that of an individual who has exited the labor force and cannot (or does not) claim UB. The three figures on the left-hand side plot averages for control and treatment groups, while the figures on the right-hand side plot the difference between them, defined as treatment minus control.

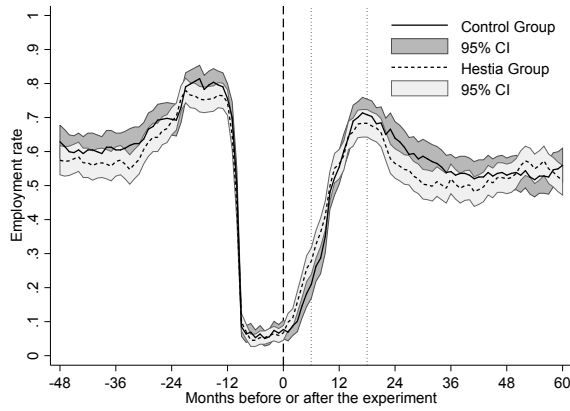
Figure 3a shows that Hestia places significantly faster its job seekers from the start of the experiment to around nine months after it but it seems that an important fraction of its job seekers lose their job again after 18 months. This happens less to job seekers who went through the public scheme and yield a statistically significant difference between the two programs.

Figure 3c shows the fraction of individuals who receive UB. We observe that Hestia pushes job seekers out of benefit rolls at a very high rate in the first six months. This rate is significantly higher than the public program, as confirmed by Figure 3d. The first six months also correspond to the period when Hestia is paid the highest rate per job seeker enrolled, which suggests that Hestia responds very strongly to the financial incentives in place. From six to twelve months, Hestia's performance tends to decrease and the two programs yield very similar outcomes after two years. Note also that around 15% of individuals in both programs are back on the benefit rolls 18 months after the experiment.

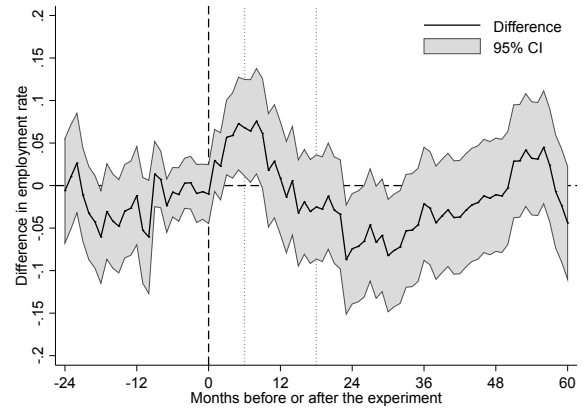
Finally, Figure 3e shows that fraction of individuals that have no work income and no unemployment benefits. Our current data does not allow us to precisely know where these people get their revenue from, but it is very likely to be either from the disability insurance or from social assistance. This fraction is increasing over time as the job seekers who did not find a job gradually run out of unemployment benefits. The effect is again similar for both the treatment and the control group though a higher fraction Hestia's job seekers (compared to public job seekers) fall into this category between 12 and 36 months after the experiment.

Combining the results that we have so far yields the following conclusions about the impacts of the experiment: (i) the private placement program responded strongly to the financial incentives in place and placed its job seekers faster than the public program in the first months; (ii) these job seekers left the unemployment benefit rolls to join the active workforce, which leads to a significant positive difference in employment rate compared to the public program; (iii) however, a large share of job seekers placed by the private provider lose their job again around two years after the experiment.

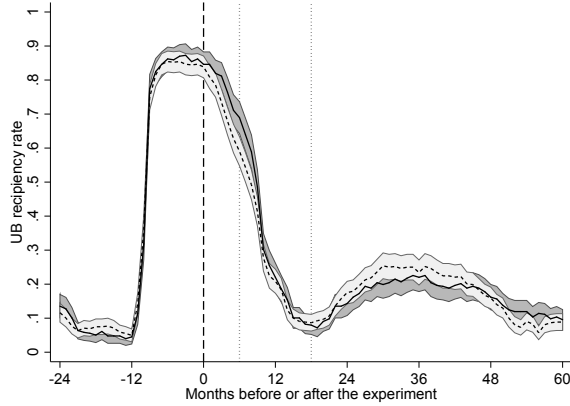
Figure 3: Comparison of labor market states



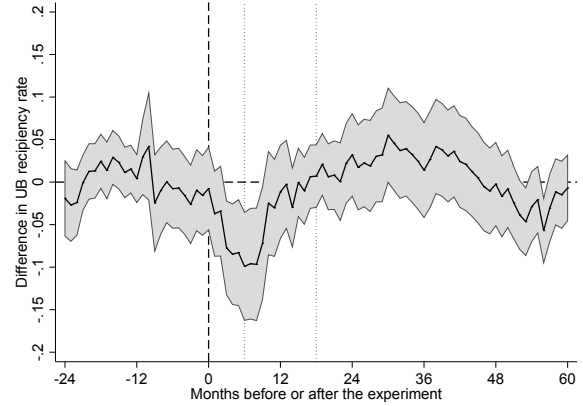
(a) Employed without UB



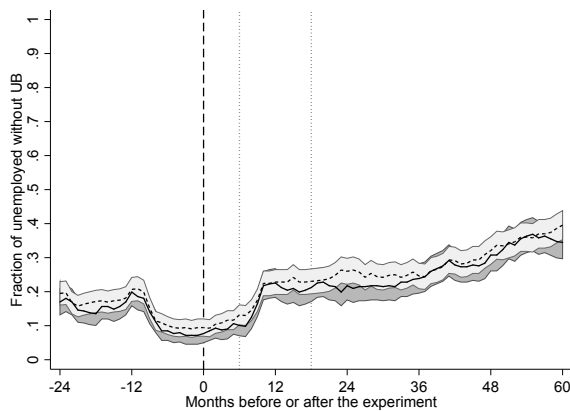
(b) Employed without UB (difference)



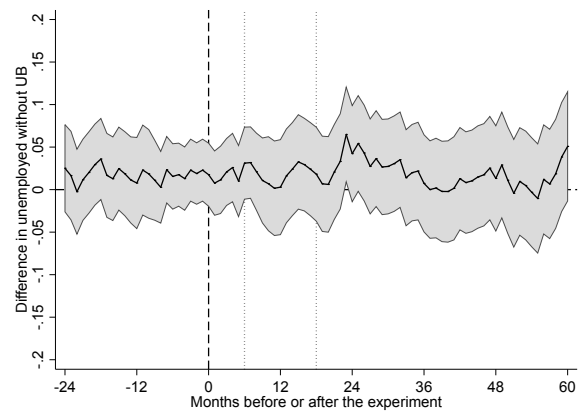
(c) UB recipients



(d) UB recipients (difference)



(e) Unemployed without UB



(f) Unemployed without UB (difference)

Notes: Figure 3 plots the fraction of individuals who are employed without UB (top), the fraction of individuals who receive UB (middle), and the fraction of unemployed individuals who do not receive UB (bottom). For the three categories, we also report the difference between the two groups, defined as treatment minus control. The two vertical dotted lines represent the changes in the payment scheme to Hestia after 6 and 18 months.

3.2.2 Unemployment benefits and wages

So far we have separately studied three mutually exclusive labor market states in which job seekers can be so as to isolate the different effects of the experiment. To get the big picture, we now consider two other variables that are useful when conducting policy and welfare evaluations: the average amount of UB received; and work income.

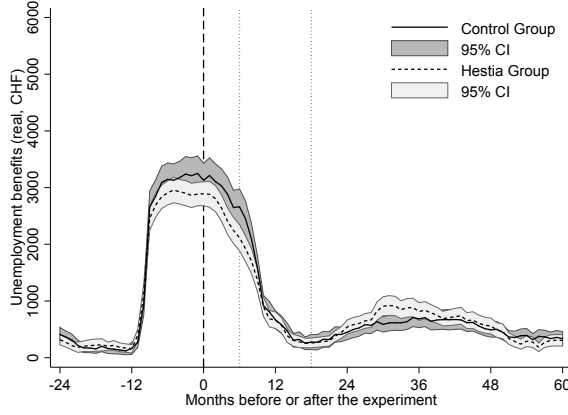
Figure 4 plots the average amount of UB benefits received (top) and the average work income over the whole sample (middle). Note that average work income can increase for two reasons: either work income increases for workers, a true work income effect, or more individuals are receiving income from work, an employment effect. To disentangle the two, we also report average income of those who work, i.e. individuals with a strictly positive work income (bottom).

Figure 4a show the evolution of the average amount of unemployment benefits received. The pattern here is similar to the one analyzed when looking at the three labor market states. Hestia gets individuals out of the unemployment benefit rolls faster, but they come back within three years. Both differences appear to be significant, as shown in Figure 4b.

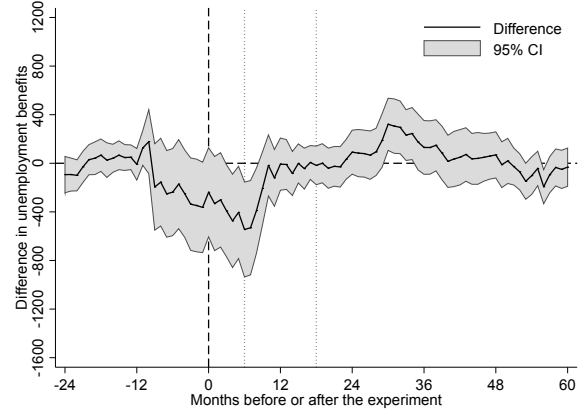
Figure 4c shows the average income from work. We see the abrupt fall due to job losses one year before the experiment. As the placement programs start, the income level start increasing again but never reach its ex-ante level. It remains slightly lower even five years after the experiment. Looking at differences in income from work, we see that Hestia's job seekers have a lower wage between two and three years after the experiment. The reason for this decrease is that some individuals loose their job again (i.e. more than those who went through the public scheme), as we saw previously. Note that the two averages are calculated over the whole sample, which implies that individuals getting no work income and/or no unemployment benefits lower the average.

Thus, another interesting variable to analyze is the average income of those who work. The dynamics of the income of those who work are of great interest since we have seen previously that Hestia placed its job seekers faster than the public provider. A way to do so would be by pushing its job seekers towards jobs that pay less than what the individuals in the public scheme would accept. Figures 4e and 4f show that job seekers in both programs get equally paid jobs, ruling out this possibility. However, the usual concern about sample selection applies. We have seen previously that the placement records of the two programs differ for certain time periods. It could be the case that Hestia focused on easy-to-place job seekers while the PES provided even services across job seekers. Thus, it is hard to attribute changes in work income to the treatment without considering further subgroup analysis. This is what we do next.

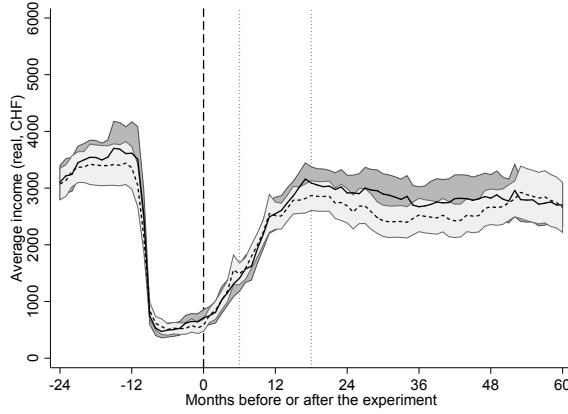
Figure 4: Sources of revenues



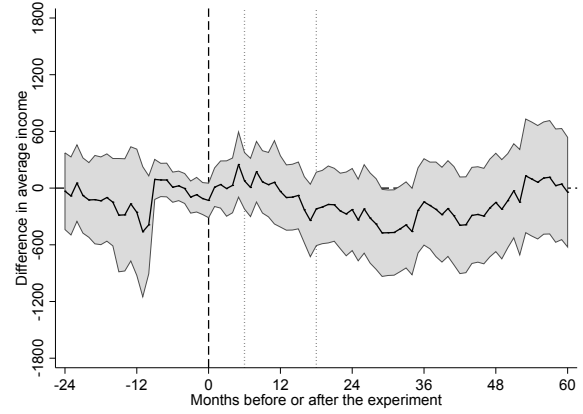
(a) Unemployment benefits



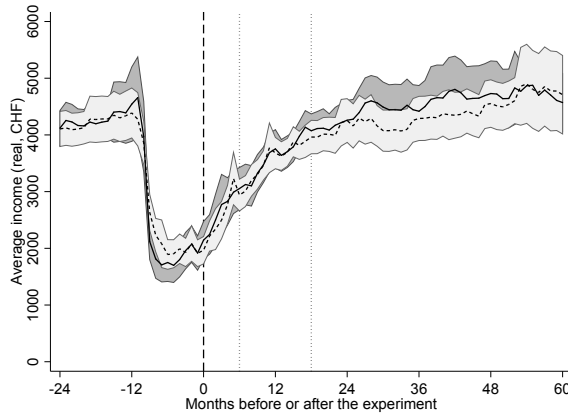
(b) Unemployment benefits (difference)



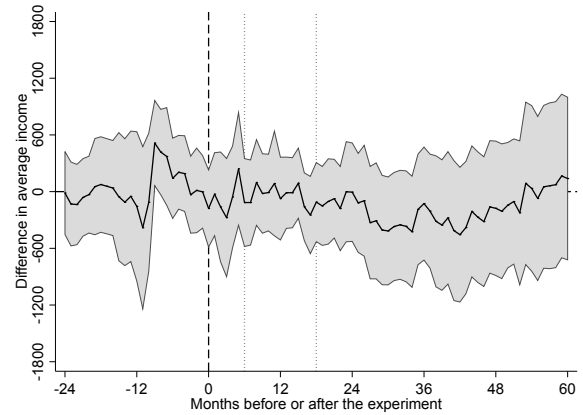
(c) Income from work



(d) Income from work (difference)



(e) Income if employed



(f) Income if employed (difference)

Notes: Figure 4 plots the average amount of UB received (top), the average income from work (middle), and the average income from work of those who work, that is have a strictly positive work income (bottom). The first two variables are averages are calculated over the whole sample. For all three variables, we also report the difference between the two groups, defined as treatment minus control. The two vertical dotted lines represent the changes in the payment scheme to Hestia after 6 and 18 months.

3.2.3 Breakdown by job seekers' characteristics

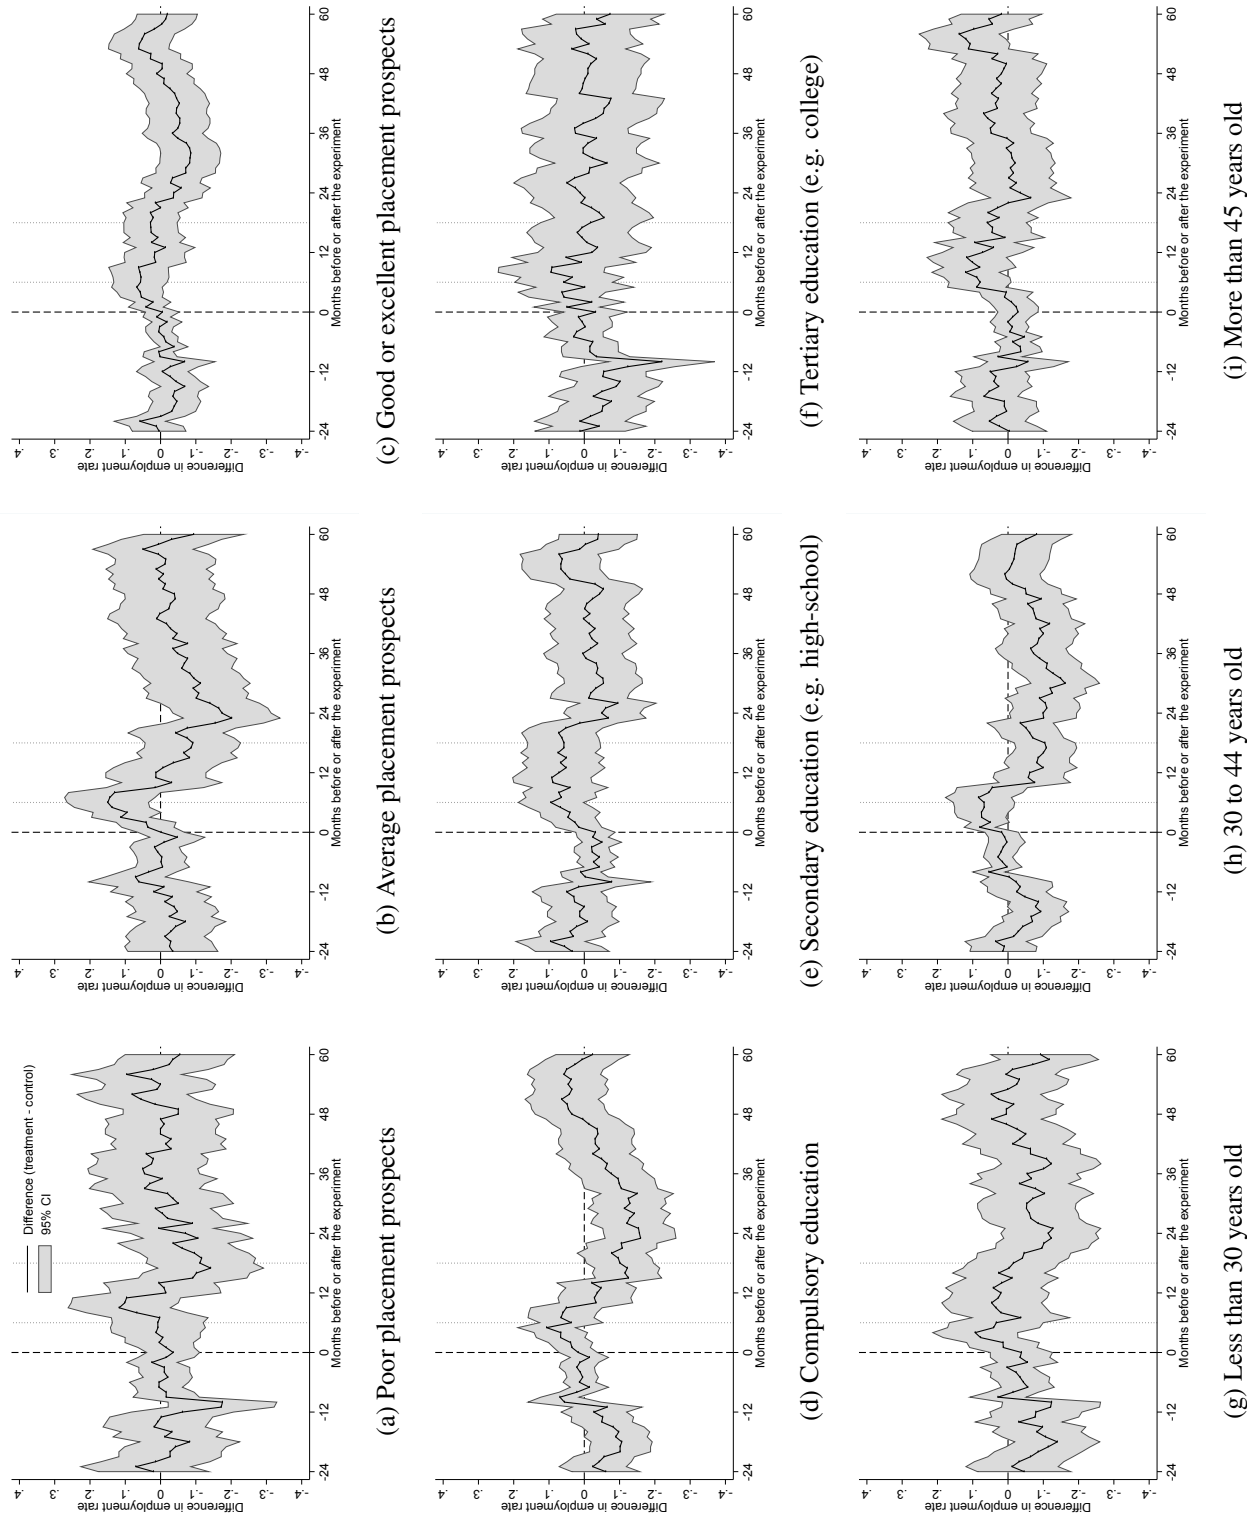
Earlier, we discussed that Hestia might engage in strategic targeting of placement efforts. We now look at the effects of Hestia on employment, namely the fraction of individuals with strictly positive work income and without UB. Figure 5 shows how the monthly difference between the treatment and the control groups vary over time across individuals with different placement prospects, age, and education. All three are correlated with labor market chances. The difference is defined as treatment minus control. Thus, a positive difference means that the private provider outperforms the public one for the given month.

We observe strong heterogeneity in the treatment effects. Consider first the breakdown by placement prospects, which is reported in the first row. From Figures 5a and 5c, we see that Hestia and the public scheme yield very similar results for individuals at the two ends of the placement prospects spectrum. Yet, there are significant differences for job seekers that have average placement prospects as seen in Figure 5c. Hestia places job seekers very fast in the first year following the experiment, outperforming the public placement provider, but a significant number of individuals lose their job again two years after the experiment. Around 2 years after the experiment, the share of employed individuals without UB placed by public program is about 5 percentage point higher than that of the private program. From three years on, we observe a convergence and the two programs eventually yield to similar outcomes.

Second, let us look at the breakdown by education level, which is reported in the second row. Figure 5d highlights significant differences between treatment and control groups over time. Hestia places individuals with compulsory schooling as highest education achieved slightly faster than the public program in the short-run. However, it is outrun by the public program between one and three years after the experiment. This pattern is not observed in Figures 5e and 5f, where the fraction of employed individuals without UB is equal in both programs for all time periods after the experiment.

Third, consider the breakdown by age, reported in the third row. Figure 5g show that there is no major differences between treatment and control groups for individuals aged 30 or less. However, Figure 5h shows a very different picture. The pattern is similar to what we observe in the other cases. Namely, Hestia places individuals between 30 and 44 years old slightly faster than the public program in the short-run but is outrun by the public program between one and four years after the experiment. Finally, Figure 5i shows that Hestia places people above 45 years old faster in the short-run but that there are no difference on a longer horizon.

Figure 5: Share of employed individuals without UB (difference)



Notes: Figure 5 shows the monthly difference in the fraction of employed individuals without UB between the treatment and the control groups (defined as treatment minus control). A positive difference means that the private provider outperforms the public one for the given month. The first row shows how this difference varies across individuals with different placement prospects, while the second row does the same with education, and the third row with age. The two vertical dotted lines represent the changes in the payment scheme to Hestia after 6 and 18 months.

4 Results

In this section, we introduce the econometric approach and then present our empirical results.

4.1 Econometric framework

The model that we estimate is given by:

$$Y_{it} = \alpha + X_i' \beta + \mathcal{T}_t' \gamma + D_i * \mathcal{T}_t' \delta + u_{it}$$

where Y_{it} is the dependent variable for individual i at process time t , X_i is a vector of individual-specific, time-constant controls, D_i is a dummy taking the value one if a job seeker is assigned to the treatment (Hestia), and zero otherwise. \mathcal{T}_t is a vector of time period dummies that span two years before and five years after assignment. Specifically, we distinguish six time periods: 24 to 13 months before, 12 to 1 month before, 0 to 11 months after, 12 to 23 months after, 24 to 35 months after, and 36+ months after assignment. $D_i * \mathcal{T}_t'$ is a vector of interaction terms between the treatment dummy and the time dummies, and u_{it} is the error term.

The parameter γ measures the detailed outcome dynamics for each outcome. The parameter δ measures the intention-to-treat effect (ITT). We report the components of δ in two separate groups. The first group contains the effects during before assignment to treatment has taken place so we can assess whether outcomes are balanced at baseline. To correctly identify the ITT, we need to make sure that the randomization worked. In other words, both groups should have similar outcomes had the treatment not been given. We can test balance of outcomes before the treatment with this set parameters. We call these effects “Randomization” parameters. The second group contains the effects after assignment to treatment has taken place. These parameters provide evidence on the effects of assignment to treatment. We call them “Treatment effects”.

4.2 Empirical results

This section presents the main results. We discuss the estimation results for the six variables discussed previously, starting with the three labor market states and then discussing UB and wages. We control both for socio-demographic and work related factors using the following variables: gender, age, marital status, schooling, nationality,⁶ mother tongue, residence permit, professional qualifications, placement prospects, job code,⁷ and cohort number. All regression tables report standard errors clustered at an individual level.

⁶For confidentiality reasons, we do not have access to the exact nationality of individuals. Instead, we use a variable that indicates whether an individual comes from specific geographic area such as Eastern Europe, South America, etc.

⁷We use one-digit PLASTA codes.

4.2.1 Labor market states

Table 3 presents the results of OLS regressions on the three mutually exclusive labor market states in which job seekers can be: (i) employed, defined as having a positive income from work, and receiving no unemployment benefits; (ii) receiving UB; and (iii) unemployed but without unemployment benefits (i.e. dependent on social assistance, on another insurance scheme, or in full-time education). All three states are continuous variables ranging between 0 and 1. Coefficients can thus be interpreted as percentage point changes with respect to the base category. We estimate the same model twice for each dependent variable. First without control variables (baseline model), and then adding control variables (main model). Coefficients are grouped in two categories. The first part highlights treatment effects on the treatment group (Hestia) on a five-year period after the experiment (A). The second part tests whether treatment and control groups differed before the experiment (B).

Consider first the treatment effects on the fraction of employed individuals with no UB (part A, columns 1-2). We observe a positive effect of Hestia in the first 12 months after the experiment. The fraction of employed individuals is around 4 percentage points higher than the public service and is statistically significant. However, this positive impact disappears after 12 months and even reverts after 24 months. The effect of Hestia on the fraction of employed individuals between two and three years after the experiment is significantly negative. This fraction is around 7 percentage points lower than for publicly placed individuals. The effect again vanishes over time.

Second, consider the treatment effects on the fraction of individuals receiving UB (part A, columns 3-4). Not very surprisingly, we observe the opposite pattern to the one for employed individuals. Hestia decreases the fraction of individuals receiving UB by 6 percentage points in the first 12 months following the start of the experiment. The effect is significant at a one percent level. In the second year after it, the effect vanishes and reverses in the third year, yet with a weak statistical significance.

Finally, consider the treatment effects on the fraction of unemployed individuals without UB (part A, columns 5-6). All coefficients are positive, which suggests that a higher share of Hestia's job seekers tend to fully exit the labor force compared to their public counterparts. However, the effect is not statistically significant for any time period.

The randomization (part B) has worked well. We do not see any significant differences between the control and the treatment group before the experiment for any of the dependent variables, which supports our empirical approach.

Table 3: Treatment effects on the three labor market states

	Employed, no UB		UB recipients		Unemployed, no UB	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Treatment Effects						
Hestia*0-11 m. after	0.046** (0.02)	0.035* (0.02)	-0.062*** (0.02)	-0.060*** (0.02)	0.016 (0.02)	0.025 (0.02)
Hestia*12-23 m. after	-0.025 (0.03)	-0.029 (0.03)	0.001 (0.02)	0.001 (0.02)	0.023 (0.02)	0.028 (0.03)
Hestia*24-35 m. after	-0.064** (0.03)	-0.075*** (0.03)	0.032 (0.02)	0.039* (0.02)	0.032 (0.03)	0.036 (0.03)
Hestia*36+ m. after	-0.014 (0.03)	-0.015 (0.03)	-0.006 (0.01)	-0.005 (0.01)	0.020 (0.03)	0.020 (0.03)
B. Randomization						
Hestia*24-11 m. before	-0.024 (0.02)	-0.039 (0.03)	0.006 (0.01)	0.010 (0.01)	0.018 (0.02)	0.029 (0.02)
Hestia*12-1 m. before	-0.013 (0.01)	-0.019 (0.01)	-0.003 (0.02)	-0.007 (0.02)	0.016 (0.02)	0.026 (0.02)
Control variables	No	Yes	No	Yes	No	Yes
Adjusted R^2	0.120	0.164	0.259	0.273	0.043	0.102
Individuals	885	849	885	849	885	849

Notes: Table 3 reports point estimates of OLS regressions on the three labor market states. All three states are continuous variables ranging between 0 and 1. The constant is included in the regressions but not reported here. Control variables include: gender, age, marital status, schooling, nationality, mother tongue, residence permit, professional qualifications, placement prospects, job code, and cohort number. Standard errors clustered at an individual level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.2.2 Unemployment benefits, wages and employment rate

We now look at the effects of Hestia on income from two sources: unemployment benefits and income from paid employment. The variables are measured in Swiss Francs. “UB” and “work income” are averages over the whole sample, whereas “income if employed” only considers individuals who have a strictly positive work income. Table 4 presents the estimation results.

Consider first the treatment effects on the amount of UB received (part A, columns 1-2). We see the short-term positive effect of sending job seekers to the private provider Hestia. They receive (or rather claim) significantly less UB (around 350 SFr) in the first 12 months after the start of the experiment. Compared to the average received by the control group during the same period (around 2000 SFr, see Figure 4a), this means a 17% difference in favor of the treatment. However, the difference vanishes in the year that follows and reverts between two and three years, which looks very similar to the trend in the share of employed individuals observed in Table 3. Between two and three years after assignment, Hestia’s job seekers are back into unemployment and receive UB for 200 SFr more than the publicly placed group. This represents a difference of about 30% in favor of the public program this time. The effect vanishes when looking at a longer horizon.

Second, let us focus on the impact of the treatment on work income (part A, columns 3-4). We see that there are no differences between privately and publicly placed job seekers in the 12 months that follow the experiment. After that, Hestia’s job seekers earn significantly less. Once we include control variables, they earn around 300 SFr less between one and two years after the experiment, and more than 500 SFr less between two and three years after it. From Figure 4, we know that individuals in the control group enjoy an average wage of around 3,000 SFr between one and three years after the experiment. This means that the difference between treatment and control group in the same period ranges between 10% and 17%. However, we see from the randomization (part B) that Hestia’s job seekers already had a lower work income before the assignment. Thus, the effects shown in the table might not be the result of the experiment only. After three years, the work income of both groups converges and the difference vanishes.

Finally, consider the treatment effects on the average income of employed individuals (part A, columns 7-8). follows closely that of work income discussed above. Hestia’s job seekers have a lower salary than their publicly placed counterparts but the difference is statistically significant only between two and three years after the start of the experiment. It then vanishes.

The randomization (part B) has worked rather well again. We do not see major differences between the control and the treatment group before the experiment. As already mentioned, there are small differences in work income and employment rate in the 12 months before the experiment.

Table 4: Treatment effects on UB and wages

	UB		Work income		Income if employed	
	(1)	(2)	(3)	(4)	(5)	(6)
A. Treatment Effects						
Hestia*0-11 m. after	-329.3** (146.72)	-352.5** (143.63)	50.9 (116.56)	-83.8 (122.33)	-27.7 (190.49)	-274.2 (170.76)
Hestia*12-23 m. after	-18.2 (68.20)	-37.0 (69.60)	-180.5 (172.45)	-301.0* (157.74)	-89.3 (187.43)	-316.0* (163.05)
Hestia*24-35 m. after	181.6** (86.57)	197.2** (86.62)	-368.0* (216.31)	-514.8*** (194.28)	-278.7 (271.26)	-534.4*** (205.73)
Hestia*36+ m. after	-10.4 (51.67)	-24.2 (56.09)	-114.5 (242.61)	-233.3 (216.01)	-78.7 (336.53)	-252.7 (250.37)
B. Randomization						
Hestia*24-11 m. before	4.4 (45.89)	2.9 (49.87)	-125.8 (217.02)	-272.1 (207.96)	-31.8 (237.56)	-279.3 (203.77)
Hestia*12-1 m. before	-167.7 (131.62)	-206.2 (130.06)	-90.7 (92.75)	-222.0* (118.00)	33.6 (206.13)	-284.9 (218.03)
Control variables	No	Yes	No	Yes	No	Yes
Adjusted R^2	0.172	0.196	0.039	0.130	0.024	0.213
Individuals	885	849	885	849	866	830

Notes: Table 4 reports point estimates of OLS regressions on three different dependent variables: “UB” and “work income” are averages over the whole sample; and “income if employed” only considers individuals who have a strictly positive work income. The constant is included in the regressions but not reported here. Control variables include: gender, age, marital status, schooling, nationality, mother tongue, residence permit, professional qualifications, placement prospects, job code, and cohort number. Standard errors clustered at an individual level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 Transition analysis

So far, we have considered in our analysis the treatment effects on the level of several economic variables, such as labor market states, UB, and wages. However, the nature of our data suggests that we could also evaluate the treatment effects on transitions from one state to another. In what follows, we explore further two results presented in the previous sections using a different approach, namely survival analysis. First, we show that the private provider improved job entry even after accounting for censoring. Second, in the same context, we demonstrate that Hestia’s placed job seekers did reach less stable positions, which can explain the negative treatment effect after one year. Finally, we compare the estimation results with the theoretical predictions of our model.

5.1 Kaplan-Meier estimates

A natural point to start from is to look at how much time job seekers in each program need to find a new job, and how long they keep it.⁸ Following Kaplan and Meier (1958), the nonparametric estimate of the survivor function is given by:

$$\hat{S}(t) = \prod_{j|t_j \leq t} \frac{n_j - d_j}{n_j} \quad (15)$$

where n_j is the number of individuals at risk of failure before time t_j , t_j represents the time at which failure occurs, and d_j is the number of failures at time t_j . This function estimates how long individuals “survive” in a given state. We present below the estimates of the survivor functions for unemployment (i.e. how long it takes for job seekers to find a new job) and for employment (i.e. how long those who have found a new job stay in their new position).

5.1.1 Transition from unemployment to employment

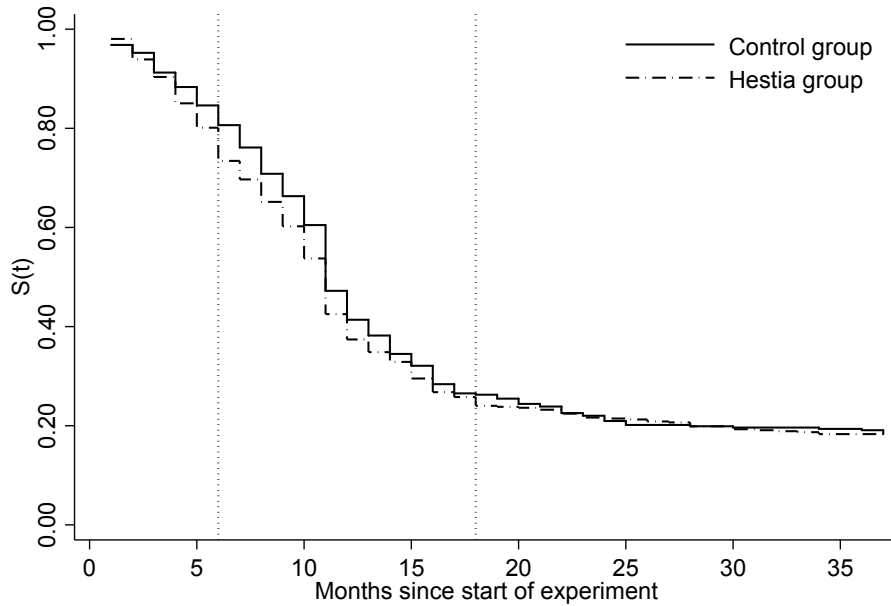
Consider first the transition from unemployment, the initial state for job seekers participating in the experiment, to employment. Figure 6 plots Kaplan-Meier survivor estimates for the two treatment groups. The origin is defined as the start of the experiment, while the failure is the entry into a new job. The two vertical dotted lines at 6 and 18 months represent the changes in the payment scheme to Hestia.

When the experiment starts, all individuals are unemployed and thus the survivor function $\hat{S}(t)$ is equal to one. As time passes, some individuals find a new job and leave the state of unemployment, making the survivor function drop. In this situation, a faster drop of the survivor function is a positive result as it means that fewer individuals are (still) in the initial state of unemployment.

⁸Note that in this section we do not consider for the analysis all individuals who started the experiment being already employed.

As we can see, the survivor curve of the treatment group is almost always below that of the control group. This suggests that the private provider improved the transition from unemployment to employment. Similar to the previous results, the difference between the two groups is greatest around six months after the start of the experiment, and slowly vanishes afterwards.

Figure 6: Kaplan-Meier survivor estimates for unemployment



Notes: Figure 6 shows Kaplan-Meier survivor estimates for unemployment. The origin is defined as the start of the experiment, while the failure is the entry into a new job. The two vertical dotted lines at 6 and 18 months represent the changes in the payment scheme to Hestia.

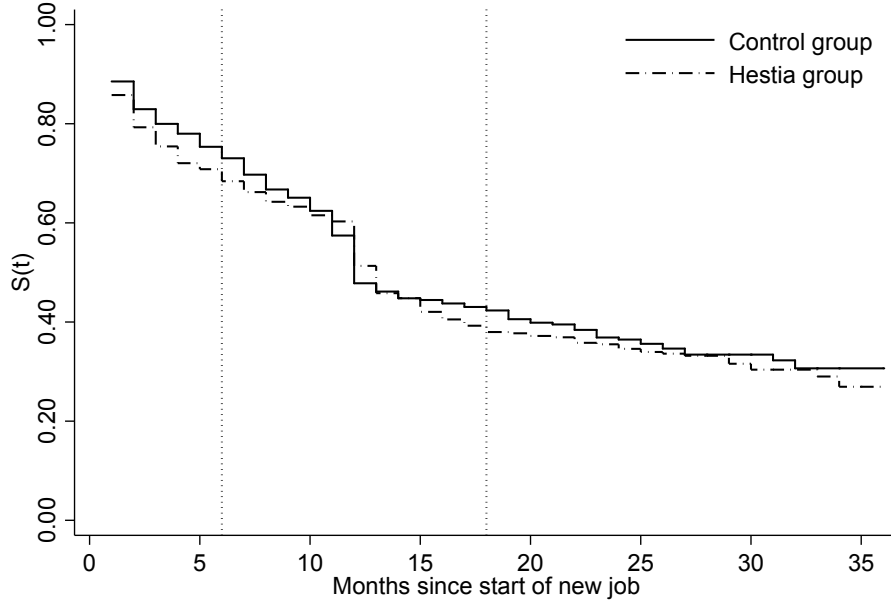
5.1.2 Transition from employment back to unemployment

Consider now the transition from employment back to unemployment. In other words, how long newly employed individuals keep their job. There is no prior as to which group should enjoy a better situation. However, previous results suggest that privately placed individuals might reach less stable positions. Two points must be noted before analyzing the results. First, only individuals who have found a job after the start of the experiment are considered here. It is therefore a subset of the individuals considered in Figure 6. Second, the interpretation of the graph is reversed this time. A high survivor function is a positive outcome as it means that more people are still employed.

We observe that the survivor curve of the treatment group is again below that of the control group, suggesting that privately placed individuals lose their job faster than the publicly placed ones. It is a negative outcome of the experiment which contrasts with the positive result discussed previously, namely faster job

re-entry. Another fact worth noting is the significant drop of the survivor function of both groups at 12 months. Many individuals lost their job exactly one year after the start of the experiment. One year fixed duration work contracts could explain this phenomenon.

Figure 7: Kaplan-Meier survival estimates for employment



Notes: Figure 7 shows Kaplan-Meier survivor estimates for employment. The origin is the beginning of a new employment spell after the start of the experiment, while failure is the loss of the job. The two vertical dotted lines at 6 and 18 months represent the changes in the payment scheme to Hestia.

5.2 Cox regressions

Kaplan-Meier estimates show that the treatment group found jobs faster but reached less stable positions. We can push the analysis further by estimating the rates at which the transitions take place using Cox regressions (Cox, 1972). This will allow us to get an estimate of the hazard rate while controlling for other factors. The hazard rate at time t is given by:

$$\lambda(t|X) = \lambda_0(t) \exp(X\beta') \quad (16)$$

where $\lambda_0(t)$ is the (unspecified) baseline hazard function, and X a vector of covariates. Note that two assumptions are required for the model to be valid. First, censoring must be non-informative.⁹ In other words, cases of censoring (e.g. failure unobserved for some individuals) must not be related to the probability of an

⁹Note that this assumption is also needed for the Kaplan-Meier estimates to be valid.

event occurring. This assumption is satisfied by design in our study as we follow individuals using administrative data and sample attrition is very low. The second assumption is that of proportional hazards. In our context, this means that the survival curves for the control and the treatment groups must have hazard functions that are proportional over time after controlling for other factors. Following Grambsch and Therneau (1994), this assumption can be formally tested using scaled Schoenfeld residuals. In our case, we cannot reject the proportional-hazard assumption both for the estimations on unemployment (chi-square p-value of 0.37 for the whole model) and on employment (chi-square p-value of 0.40 for the whole model).

Table 5 presents the results. Column (1) reports the treatment effects on the transition rate from unemployment to employment, while column (2) shows the effects on the transition rate from employment back to unemployment. Recall that positive coefficients in column (1) are positive results since it means that the transition from unemployment to employment increases. Alternatively, positive coefficients in column (2) are a negative outcome, since it means that people leave the state of employment at a higher rate.

Consider first the results on unemployment. The private provider seems to have a positive short term effect. Between four and six months, it significantly increases transition to employment. However, this positive effect vanishes and becomes negative after a 18 months, and significantly so for the period 19-21 months. In total, 849 individuals started the experiment unemployed and 719 found a job during the observation period, which represents almost 85% of the subjects.

Second, consider the effects on the transition from employment to unemployment. As before, only individuals who have found a job after the start of the experiment are considered here, which explains why the number of subjects in column (2) is equal to the number of failures in column (1). Privately placed individuals leave employment significantly faster in the first three months after the start of the experiment. This could suggest that these individuals were not satisfied with the positions that they reached and left when they first got the chance. In Swiss labor law, many contracts come with a two-month probation period, after which both sides (employer and employee) can break the contract unilaterally. The negative treatment effect suggests that this could have happened. There is a positive treatment effect between 10 and 12 months after the start of the experiment. As stated before, this could be due to one-year fixed duration contracts. However, the effect is short-lived as it is followed by a significant negative effect in the next time period.

Overall, the analysis of transitions support the results of Section 4.2.1, namely that Hestia improved placement rate in the short-run but had a negative effect on job stability in the medium-run.

5.3 Comparison with theoretical predictions

We now link the theoretical predictions of our model with the empirical estimates from the experiment. To do so, we estimate the treatment effect on the transition rate from unemployment to employment using a Cox regression similar to the one presented in Figure 5 column (1). The only difference is that we fit it with

Table 5: Cox regressions on transitions

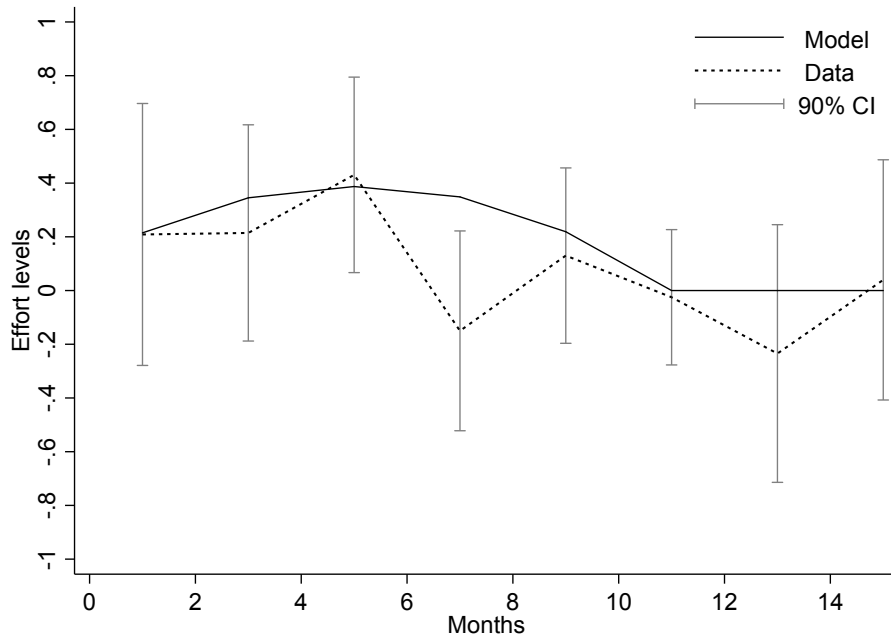
	(1) To Job	(2) To Unemployment
Hestia*1-3 months	0.050 (0.23)	0.270* (0.16)
Hestia*4-6 months	0.475** (0.19)	-0.036 (0.28)
Hestia*7-9 months	-0.052 (0.18)	-0.280 (0.28)
Hestia*10-12 months	0.018 (0.13)	-0.458** (0.20)
Hestia*13-15 months	0.019 (0.24)	0.774** (0.34)
Hestia*16-18 months	-0.066 (0.29)	0.529 (0.45)
Hestia*19-21 months	-1.056* (0.60)	-0.638 (0.58)
Hestia*22-24 months	-0.494 (0.45)	-0.422 (0.43)
Hestia*25+ months	-0.104 (0.29)	-0.162 (0.19)
Control variables	Yes	Yes
Subjects	849	719
Failures	719	566

Notes: Table 5 reports point estimates of Cox regressions on transitions to job and to unemployment. For the transition from unemployment to job, the origin is defined as the start of the experiment, while failure is the entry into a new job. For the transition from job back into unemployment, the origin is the beginning of a new employment spell after the start of the experiment, and failure is the loss of the job. Control variables include: gender, age, marital status, schooling, nationality, mother tongue, residence permit, professional qualifications, placement prospects, job code, and cohort number. Standard errors clustered at an individual level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

time-dummy intervals of two months instead of three so as to make the comparison with the model easier. The assumption here is that each period in the model will represent two months in real life. We report the estimated coefficients in Figure 8 (dashed line), along with their 90% confidence intervals. For example, the coefficient on the interaction term for the first two months after the start of the experiment (0.2088) is reported above month one in the graph. Similarly, the coefficient on the interaction term for months three and four (0.2145) is reported above month three.

Then, we add to the plot the expected effort levels predicted by our model (full line). As stated before, the six-period specification assumes that each period counts for two months. In order to be as close as possible to the financial scheme used in the experiment, we solve our model with $\delta = 0.9$. Recall that in case of a non-decreasing payment scheme (i.e. a flat rate until the end of time), the model would predict effort levels of zero for all period. Thus, the predictions of our model represent the difference in efforts between a provider without payment incentives (e.g. PES) and one with financial incentives (e.g. Hestia), which is also what the Cox regression is meant to capture in our context.

Figure 8: Model vs data



Notes: Figure 8 plots the optimal effort levels predicted by our model with six time periods and $\delta = 0.9$, and the estimates of a Cox regression on the transition between unemployment and employment. For the Cox regression, we only report the estimates on the interactions terms between treatment and two-months time dummies, along with their 90% confidence interval. Control variables in the Cox regression include: gender, age, marital status, schooling, nationality, mother tongue, residence permit, professional qualifications, placement prospects, job code, and cohort number.

We observe that the predicted effort levels and the realized ones are close for most time periods. The private provider exerts significant effort levels in the first six months of the experiment. However, there is a

drop in realized efforts after six months, which is not predicted by the model. Two explanations are possible.

First, it could be that Hestia was facing an even shorter horizon than the 12 months assumed here. A look back at Figure 1 suggests that the realized effort levels could be consistent with the model's predictions if we consider fewer time periods. This could be the case for example if Hestia doubted that job seekers would stay in their program even after running out of unemployment benefits, which occurred between six and twelve months for most individuals.

Second, it could be that the model does not capture well the PES' incentives. The PES' predicted effort levels should be zero for all periods but we observe some negative realized effort levels. This implies that the PES' effort level was greater than Hestia's for certain time periods, pointing towards a missing ingredient in our model. For example, it is likely that caseworkers at PES are not only motivated by financial incentives but also have an intrinsic motivation to do their job well.¹⁰ Such a scenario would explain why PES' effort levels are greater than zero for some time periods, and why the fit of our model is not accurate.

6 Cost-benefit analysis

Our evaluation of the placement scheme would not be complete without a cost-benefit analysis. From the State's perspective, there are two types of costs, and one type of benefit. The program costs represent the amount paid to Hestia and PES for their job placement programs. On top of it, there are indirect costs in the form of unemployment benefits paid to job seekers enrolled. Finally, the benefits represent the amount saved on unemployment benefits thanks to shorter unemployment duration.

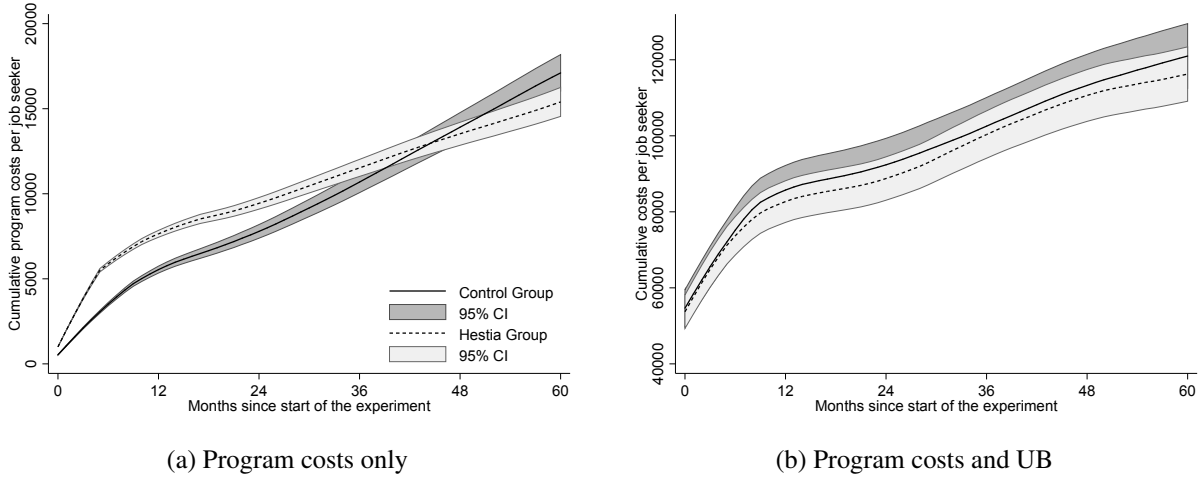
Flückiger and Kempeneers (2008) have estimated the average monthly program costs for the public track to be equal to SFr. 573. It has been done by summing up two budgets: the PES budget; and a budget for additional ALMP in Geneva; and then dividing the amount by the total number of job seekers that benefited from it in the previous year (2006). Program costs for the private provider are more straightforward. From Section 2.2, we know that a job seeker enrolled with Hestia costs 1000 SFr. a month in the first six months; 500 SFr a month for six to 18 months; and 350 SFr. a month after it. As soon as this person is placed, the cost drops to zero. If this individual falls back into unemployment during the observation period, then her costs will be the same as for the public track, namely 573 SFr. per month.

Figure 9 plots the average cumulative cost per individual enrolled in one of the two programs, first without accounting for unemployment benefits, and then accounting for variations in unemployment benefits. Note that we do not account here for effects on taxes (e.g. more or less taxes paid depending on changes in earnings, which would be a benefit from the State's perspective) and for non-monetary benefits for job

¹⁰One way to account for this would be to postulate a different utility maximization problem for the PES. Since the PES budget does not depend on case-load and that there are no financial incentives, one could express the case-workers' utility as $U = \alpha e - \frac{1}{2}e^2$ where $\alpha \in (0, 1)$ is the PES efficiency of effort. A utility maximizing case-worker would then select $\alpha = e$. This would yield again a flat effort level but it could be greater or less than the private provider's effort.

seekers (e.g. improved well-being due to shorter unemployment period).

Figure 9: Costs per job seeker



Notes: Figure 9a plots the average cumulative cost per individual enrolled in one of the two programs without accounting for unemployment benefits, while Figure 9b does the same but also accounts for unemployment benefits.

Results are not surprising. By design, the private track is much more costly in the first six months than the public one, which is reflected in Figure 9a. Without accounting for UB, faster job entry in the short run does not compensate for a monitoring cost 75% more expensive than the PES. However, the PES catches up in the long-run, as Hestia's payments are being reduced.

The most interesting result can be seen in Figure 9b. Once the endogenous effect on UB is accounted for, outsourcing job placement is self-financing. The amount saved on UB even in the short-run is large enough to compensate for the difference in program costs, which makes the private program cheaper than the public one for all time periods.

7 Conclusion

In this paper, we are interested the effects of outsourcing job search assistance for long-term unemployed. More specifically, we look at whether a private placement provider can help lowering the level of long-term unemployment. We study the work trajectories of about 890 individuals gathered in a randomized controlled experiment conducted in 2006-2007, in conjunction with social-security data that cover two years before and five after the experiment. Focusing on the monthly employment patterns of individuals, we find that results change dramatically depending on the time horizon considered.

In the short-run, the private provider significantly improves job seekers re-entry into the labor force, with a difference of around 4 percentage points compared to the public one. This also lowers by around 17% the average amount of unemployment benefits received by the same group. In the medium-run though,

these positive impacts vanish and both the public and the private providers have a similar performance until approximately two years after the experiment. Then, the patterns revert. Privately placed job seekers lose their job again, earn less when they are employed and claim more unemployment benefits than their public counterparts. The difference is significant up to three years after the experiment and finally disappears when we looking at a longer horizon. These results suggest that the private provider focuses on placing job seekers as fast as possible at the expense of their suitability to the position. Indeed, finding a job faster does not mean it is better.

We also observe a strong heterogeneity in job seekers' types. Namely, it seems that the private provider struggles to place well "average" job seekers. It yields comparable performance to the public placement scheme on easy-to-place and hard-to-place individuals but does worse for individuals that are in between. This could be explained by its capacity constraint, which forces it to allocate its resources across all job seekers' types to survive in the long-run.

Finally, a cost-benefit analysis suggests that despite its high monitoring costs, outsourcing job search assistance is self-financed through a decline in unemployment benefits claimed.

References

- Arni, P. (2015). Do older job seekers find more jobs when being coached? a field experiment on search behavior. mimeo, Institute for the Study of Labor (IZA).
- Arni, P., Lalive, R., and Van Ours, J. C. (2013). How effective are unemployment benefit sanctions? looking beyond unemployment exit. *Journal of Applied Econometrics*, **28**(7), 1153–1178.
- Behaghel, L., Crépon, B., and Gurgand, M. (2014). Private and Public Provision of Counseling to Job Seekers: Evidence from a Large Controlled Experiment. *American Economic Journal: Applied Economics*, **6**(4), 142–74.
- Bennmarker, H., Grönqvist, E., and Öckert, B. (2013). Effects of contracting out employment services: Evidence from a randomized experiment. *Journal of Public Economics*, **98**(0), 68 – 84.
- Black, D. A., Smith, J. A., Berger, M., and Noel, B. (2003). Is the Threat of Reemployment Services More Effective Than the Services Themselves? Evidence from Random Assignment in the UI System. *American Economic Review*, **93**(4), 1313–1327.
- Card, D., Kluve, J., and Weber, A. (2010). Active Labour Market Policy Evaluations: A Meta-Analysis. *Economic Journal*, **120**(548), F452–F477.
- Cox, D. (1972). Regression models and life tables (with discussion). *J Royal Stat Soc Ser B*, **34**, 187–220.

- Eugster, B. (2013). Effects of a higher replacement rate on unemployment durations, employment, and earnings. Economics Working Paper Series 1320, University of St. Gallen, School of Economics and Political Science.
- Flückiger, Y. and Kempeneers, P. (2008). Evaluation de l'impact économique, social et financier des programmes de retour en emploi proposés par les *Maisons Hestia* à Genève. Technical report, Observatoire Universitaire de l'Emploi (OUE) Geneva.
- Gerfin, M. and Lechner, M. (2002). A microeconomic evaluation of the active labour market policy in Switzerland. *The Economic Journal*, **112**(482), 854–893.
- Grambsch, P. M. and Therneau, T. M. (1994). Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika*, **81**, 515–526.
- Graversen, B. K. and van Ours, J. C. (2009). How a Mandatory Activation Program Reduces Unemployment Durations: The Effects of Distance. IZA Discussion Papers 4079, Institute for the Study of Labor (IZA).
- Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, **53**(282), 457 – 481.
- Krug, G. and Stephan, G. (2013). Is the contracting-out of intensive placement services more effective than provision by the pes? evidence from a randomized field experiment. IZA Discussion Paper 7403, IZA.
- Lalive, R., van Ours, J. C., and Zweimüller, J. (2005). The Effect of Benefit Sanctions on the Duration of Unemployment. *Journal of the European Economic Association*, **3**(6), 1–32.
- Lalive, R., Van Ours, J. C., and Zweimüller, J. (2008a). The impact of active labour market programmes on the duration of unemployment in Switzerland*. *The Economic Journal*, **118**(525), 235–257.
- Lalive, R., van Ours, J. C., and Zweimüller, J. (2008b). The Impact of Active Labor Market Programs on the Duration of Unemployment. *The Economic Journal*, **118**(2008), 235–257.
- Rehwal, K., Rosholm, M., and Svarer, M. (2015). Are public or private providers of employment services more effective? evidence from a randomized experiment. IZA Discussion Paper 9365, IZA.
- Rosholm, M. and Svarer, M. (2008). The threat effect of active labour market programmes. *Scandinavian Journal of Economics*, **110**(2), 385–401.