

IZA DP No. 9786

Job-Search Periods for Welfare Applicants: Evidence from a Randomized Experiment

Jonneke Bolhaar
Nadine Ketel
Bas van der Klaauw

March 2016

Job-Search Periods for Welfare Applicants: Evidence from a Randomized Experiment

Jonneke Bolhaar

*Netherlands Bureau for Economic Policy Analysis (CPB)
and VU University Amsterdam*

Nadine Ketel

*University of Gothenburg
and IZA*

Bas van der Klaauw

*VU University Amsterdam, Tinbergen Institute,
CEPR and IZA*

Discussion Paper No. 9786
March 2016

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Job-Search Periods for Welfare Applicants: Evidence from a Randomized Experiment*

This paper studies mandatory job-search periods for welfare applicants. During this period the benefits application is put on hold and the applicant is obliged to make job applications. We combine a randomized experiment with detailed administrative data to investigate the effects of imposing a job-search period. We find strong and persistent effects on the probability to collect welfare benefits. The reduced benefits are fully compensated by increased earnings from work. Furthermore, we do not find evidence of adverse consequences for the most vulnerable applicants. Our results therefore suggest that a job-search period is an effective instrument for targeting welfare-benefits applicants.

JEL Classification: C21, C93, I38, J64, J08

Keywords: job search, welfare-to-work, active labor-market policies, randomized experiment

Corresponding author:

Bas van der Klaauw
Department of Economics
VU University Amsterdam
De Boelelaan 1105
1081 HV Amsterdam
The Netherlands
E-mail: b.vander.klaauw@vu.nl

* We gratefully acknowledge valuable comments from Hessel Oosterbeek, Randi Hjalmarsson and seminar participants in Amsterdam, Gothenburg, London, Aarhus, Munich, Rotterdam, Louvain-la-Neuve and Vienna. We thank the Dienst Werk en Inkomen Amsterdam and all the caseworkers for their cooperation.

1 Introduction

Many countries have some welfare system which provides benefits to low-income households. Welfare then acts as a safety net guaranteeing a minimum level of income. In the United States, welfare is mostly used to support single-parent households. In European countries, it also supports long-term unemployed workers, who are no longer entitled to social-insurance benefits. Whereas welfare aims at households with a low capability of generating sufficient income, governments often have imperfect information about the income-generating potential of a household. This may induce moral hazard, which causes excess spendings on welfare benefits. Currently, there is a tendency among policy makers in many countries to restrict access to benefit schemes and to be more strict on job-search requirements. Causal evidence on the effect on take up of more stringent entry requirements is, however, scarce (Currie, 2006).

This paper evaluates mandatory job-search periods for individuals applying for welfare benefits. During such a four-week period the application is put on hold, and applicants are supposed to make many job applications. The application for welfare benefits is only activated if the applicant returns to the welfare agency after the job-search period. Benefits are then paid from the moment of the initial application. Therefore, the mandatory job-search period does not change eligibility or the amount of benefits, but only postpones the first payment. To evaluate such job-search periods we conducted a randomized experiment between April 2012 and March 2013, incorporating the full population of welfare applicants with a potential to work in the city of Amsterdam. Combining various data sources we construct a very detailed administrative dataset describing the participants in the experiment. This allows us to not only look at the implications of a job-search period on welfare benefits receipt, but also study the effect on alternative sources of income. In addition, we can establish whether a job-search period increases the likelihood to engage in criminal activities.

A job-search period can affect labor-market outcomes in several ways. First, the job-search requirement can increase the likelihood to find a job and thereby reduce the likelihood to receive welfare benefits. During the job-search period workers have to search

very actively, but they can themselves decide to which jobs to apply. Once starting collecting welfare benefits workers should accept all jobs offered to them. Black et al. (2003) suggest that announcing obligations may stimulate benefits recipients to find work. Second, a job-search period makes the application process for welfare benefits more complex and increases the costs of applying. After the job-search period, applicants have to pay a second visit to the welfare office to confirm their application for welfare benefits. This can decrease the likelihood to receive benefits even in absence of an effect on job finding. Both mechanisms can serve as a self-selection or self-screening device (Parsons, 1991), but possibly effect a different part of the population of applicants. An increase in job finding reduces take up of applicants with relatively good labor-market prospects, while increased complexity discourages applicants that do not find a job but cannot deal with the complexity of the application process.¹

Several studies document a decrease in take up of means-tested welfare benefits as a result of increased application costs or complexity of the application process (e.g. Bhargava and Manoli (2015), Currie and Grogger (2001), Krueger (1990), Bitler et al. (2003) and Brien and Swann (1999)). These increased application costs include requiring more frequent visits to the welfare office, reduced re-certification intervals or requiring extensive income documentation. Kleven and Kopczuk (2011) model complexity as an instrument used by program administrators to extract a better signal of true eligibility, which is chosen jointly with benefit levels and eligibility rules in the design of a program. In all these cases the question remains whether (non-financial) barriers target the desired population. The mentioned studies report a decrease in take up, but do not observe the source of income of non-recipients. Without this it cannot be established whether the intended goal is accomplished. We exploit the availability of extensive administrative data in the Netherlands to follow participants in our experiment. This allows us to combine the advantages of field experiments and administrative data.

¹A third mechanism could run via time-inconsistent preferences (O'Donoghue and Rabin, 1999). Many costs of applying for welfare are borne immediately, whereas benefits are received later. Hence, a job-search period might discourage a person who puts more weight on the present than on the future to carry through the application for benefits, even though it would be utility maximizing to receive benefits at a later date.

The requirement in the job-search period relates our study to the literature on the effectiveness of active labor-market policies. The majority of this literature concentrates on programs aimed at recipients of unemployment insurance benefits with relatively good employment prospects (for an overview, see Card et al. (2010)). Welfare recipients are often more disadvantaged and at risk of leaving the labor force permanently, so the potential gains of effective programs for welfare applicants can be large. Furthermore, for credit-constrained welfare applicants small financial incentives can matter a lot. For example, Van der Klaauw and Van Ours (2013) find that imposing benefit sanctions substantially increases the individual transition rate from welfare to work in the Netherlands. Card and Robins (1998) show that the financial incentives in the Canadian self-sufficiency program induced welfare recipients to work more. However, randomized experiments with welfare applicants are scarce. Given the vulnerability of the population, authorities are not very likely to agree with a randomized experiment. This is particularly true when evaluating existing policies (in contrast to the evaluation of additional policies or resources).

The setup of our randomized experiment is similar to a so-called encouragement design (Duflo et al., 2007). However, whereas in the usual encouragement design treatment is encouraged on a randomly selected group of subjects and there are no costs of noncompliance, we impose a stronger encouragement causing higher compliance rates. We randomize treatments over caseworkers, who receive the instruction to apply one particular treatment, a default option, to all their new clients in a given period. In case the default option is really not appropriate, caseworkers are allowed to deviate but should provide a motivation. The possibility to deviate increases support for the experiment among caseworkers. The design exploits the random assignment of applicants to caseworkers within each local welfare office.² Our empirical strategy is similar to Maestas et al. (2013), who exploit variation in examiners' allowance rates as an instrument for disability benefit receipt, using that applicants are randomly assigned to disability examiners.³ Also, Crépon et al.

²Within a local office all caseworkers have the same target concerning exit to work. The random assignment of applicants to caseworkers ensures that all caseworkers have the same fair chance to meet the target. This also allows the welfare agency to benchmark caseworkers.

³We can not exploit the natural variation per caseworker, because aside from the decision about a job-search period caseworkers are also responsible for the further guidance of the applicant towards a job. And there may be differences between caseworkers in how successful they are in providing guidance.

(2013) used an encouragement design to evaluate active labor-market programs in France. Compared to these two studies, the variation that we can exploit is somewhat higher. Our study has a compliance rate to the random treatment assignment of 46 percent compared to 35 percent in Crépon et al. (2013), while Maestas et al. (2013) exploit natural variation of 23 percent.

We find a strong and persistent negative effect of a job-search period on the likelihood to receive welfare benefits. Imposing a job-search period reduces the likelihood to receive benefits by 20 percentage points. The effect is significant up to six months after registration, and during these months total welfare benefits payments are about 25 percent lower. There is no spillover to other benefits schemes and the lower income from welfare benefits is fully compensated (112 percent) by higher earnings. A job-search period does not increase the likelihood to engage in criminal activities. The fact that the reduced income from benefits is fully compensated by higher earnings suggests that a job-search period targets the desired population of applicants, which is confirmed by a subgroup analysis. The effect of a job-search period increases with education. For applicants with at least a bachelor degree the likelihood to receive benefits decreases with 50 percent. We find no evidence of negative side effects of a job-search period for the most vulnerable applicants. Only for higher educated applicants the job-search period increases the probability to have a very low income. Finally, we find that the estimated local average treatment is not sensitive to changes in the sample and the group of compliers. This implies that caseworkers do not succeed in targeting job-search periods to those applicants for which the effects are largest.

The remainder of the paper is structured as follows. The next section provides details about the Dutch benefit system, explains the experimental design and provides evidence on compliance rates. Section 3 describes the data and provides evidence on the random assignment. In section 4 we discuss the empirical strategy and identification. Section 5 presents the main results, while section 6 assesses the heterogeneity of treatment effects. Section 7 discusses the generalizability of our results by interpreting the estimated local average treatment effect. Section 8 concludes.

2 Setting and experimental design

Welfare in the Netherlands

In the Netherlands, welfare serves as a safety net and provides households that have no or not enough means of living with a minimum income level. Welfare benefits are means tested (on both income and wealth) and the benefits level only depends on the composition of the household. Rules about eligibility and the level of benefits are determined nationwide, but the responsibility for the implementation is at the municipality level. Welfare is paid from general taxation and municipalities receive from the Ministry of Social Affairs a fixed annual budget for managing their welfare system, of which any unused excess may be kept.⁴

The regular benefits range from 70 percent of the minimum wage for a single-person household to 100 percent for a couple with children. During our observation period the (nationwide) net minimum wage was about 1200 euros per month. Municipalities can give additional benefits on top of the regular benefits. Amsterdam pays a housing allowance to welfare benefits recipients of 133 euro per month. Furthermore, all Dutch households with low income are entitled to housing subsidies, health-insurance subsidies and child subsidies.⁵ Unlike in the US, there is no maximum to the time period that a household can receive welfare benefits (or subsidies). If a welfare recipient finds part-time employment or has part-time employment with earnings below the welfare benefits level, earnings have a marginal tax rate of 100 percent. For the health-insurance and housing subsidies other marginal tax rates apply (child subsidies are not income related).

Individuals apply for benefits at the local welfare office in their city district. The applicant is next invited for an intake meeting with a caseworker, during which the rights and obligations of receiving welfare benefits are explained. Applicants have to bring extensive proof of their (past) income, bank accounts, housing etc. to this meeting. Based on this information, supplemented with information from administrative sources, it is determined

⁴In 2012, the municipality of Amsterdam spent 103 percent of its budget for welfare which was also the average for all municipalities in the Netherlands.

⁵The maximum monthly amounts of these subsidies are 309 euro (for housing subsidies), 70 euro (health-insurance subsidies) and per child 84 euro (child subsidies).

whether an individual is entitled to benefits. Welfare recipients have to comply with job-search requirements, and are obliged to accept all jobs, irrespective of the match with their education or work experience. These obligations are set by the national government, but the municipality has discretion in deciding about how welfare recipients are supported in their job search and how the number of welfare recipients is kept under control.

Setting of experiment

The sample for our experiment consists of individuals that applied for welfare benefits in Amsterdam between April 2012 and March 2013 and who should be able to find regular employment within six months. The latter is determined through a computerized program that profiles all workers based on their characteristics. These characteristics include, among others, work history, age, education, language and computer skills, recent detention and psychological problems. Based on this profile the applicant is classified in one of four classes. The type and intensity of guidance given to the applicant and the required job-search effort varies over these classes.

Mandatory job-search periods are only applied to the class of applicants that should be able to find regular employment within six months, so the population participating in our experiment. In addition, we restrict the experiment to individuals that are at least 27 years old, as different rules apply to welfare recipients under age 27. These selection criteria apply to approximately 25 percent of the total inflow (2860 applicants). The welfare recipients in Amsterdam are divided over five welfare offices serving workers in different districts.

Amsterdam has a relatively large population of welfare recipients. In January 2012, 6.4 percent of the population between 20 and 65 years old received welfare benefits in Amsterdam (i.e. 34,550 individuals) compared to 3.1 percent in the Netherlands. During 2012, the inflow into welfare in Amsterdam consisted of 11,706 individuals while in the same year 8,944 individuals stopped receiving benefits. The larger inflow is a direct result of the economic conditions. Figure B1 in the appendix shows GDP growth for the Netherlands and inflow and outflow into welfare benefits in Amsterdam from 2008 until

2014. At the time of the experiment, in 2012 and the first half of 2013, the Dutch economy was suffering from a second downturn as a consequence of the financial crisis. Worsening economic conditions have both a direct and a lagged effect on inflow into welfare benefits. If unemployment increases, it will directly increase the inflow into welfare benefits for individuals that have no or limited entitlement to unemployment insurance benefits. From individuals that can first deplete their unemployment insurance benefits there will be a delayed inflow into welfare benefits.

During the experiment we asked caseworkers to fill in a form with information about the applicant in each intake meeting. Table A1 in the appendix gives a description of the population that took part in the experiment, based on these forms. The caseworkers report that 26 percent of the applicants have a bad financial situation. The applicants are considered to be relatively independent, which is probably related to the fact that they are classified as being able to find employment within six months. 24 percent of the applicants applied for benefits directly after losing (self-)employment, while the remainder either depleted their unemployment insurance benefits or applied for another reason. Other reasons to apply for welfare include exhaustion of savings, divorce, less hours at an existing job such that the total earnings drop below the welfare level, etc. Finally, 56 percent of the applicants are non-western immigrants and 47 percent of the applicants have received welfare benefits before.

Intervention: Mandatory job-search period

Our experiment focuses on the job-search period, which was introduced by the welfare agency in 2011. A job-search period postpones the application for welfare benefits with at most four weeks, during which the individual has to actively search for work. The application for benefits will only be activated if the applicant returns to the welfare agency after the job-search period. If the welfare application is activated and processed, the applicant will (retrospectively) receive benefits starting at the date of the initial registration. A job-search period thus only delays the first payment of benefits, it does not reduce the amount of benefits that an individual is entitled to. Applicants who find employment dur-

ing a job-search period can file a request for receiving the welfare benefits for the period between the date of registration and the starting date of new employment. This requires completing the application process and is not actively promoted by the welfare agency. So most individuals that find a job during the job-search period do not use this possibility. Figure 1 gives a schematic representation of the application process for welfare benefits in Amsterdam, including the job-search period. Irrespective of the job-search period, the welfare application needs to be processed within eight weeks after the day of registration.

The decision to apply a job-search period is made during the intake meeting and, before the experiment, was left to the caseworkers' discretion. However, job-search periods should not be applied if an applicant has severe financial problems or can prove that she has been very active in applying for jobs prior to registration at the welfare agency. An applicant cannot refuse a job-search period. When imposing a job-search period the caseworker specifies a minimum number of job applications that the applicant should make within the job-search period. The caseworker also stresses that during the job-search period the applicant can still choose which jobs to apply to, but as soon as the applicant starts receiving welfare benefits it is mandatory to accept all jobs. If the applicant returns from a job-search period, the caseworker generally checks whether the applicant has complied with the job-search requirement, and can impose a sanction if this is not the case. This sanction is generally a 30 percent reduction in benefits for the duration of one month. In practice, these sanctions are almost never applied.

Experimental design

To isolate the causal effect of a job-search period we conducted a randomized experiment in which we manipulated the assignment of job-search periods.⁶ Applicants are not informed about the experiment, to prevent that this knowledge would influence their behavior. Instead of randomizing the treatments over individuals, we randomized the treatments over caseworkers. Caseworkers received the instruction to apply one particular treatment to all their new clients during a three-month period. This particular treatment we call their

⁶The original research design, including a power analysis, can be found at <http://personal.vu.nl/b.vander.klaauw/OnderzoeksOpzetDWI.pdf> (in Dutch).

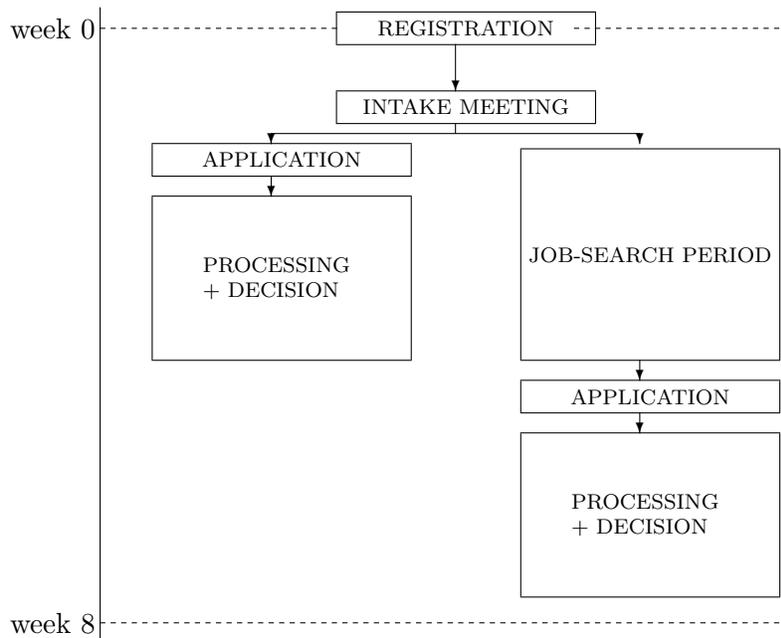


Figure 1: Welfare application procedure

default option, which makes our design similar to an encouragement design (Duflo et al., 2007). We instructed caseworkers to deviate from the default option only in cases where the default option is really not appropriate. The possibility to deviate in special cases helped to make the experiment more acceptable for caseworkers and in getting their commitment to the experiment. Our design exploits that within local offices welfare applicants are randomly allocated to caseworkers. The matching of applicants to caseworkers is an administrative process, in which welfare applicants are matched to the caseworker with the lowest case load.

There are three different default options:

- Never: never impose a job-search period.
- Always: always impose a job-search period if the financial situation of the individual allows for it.
- Normal: the decision to impose a job-search period is left to the discretion of the

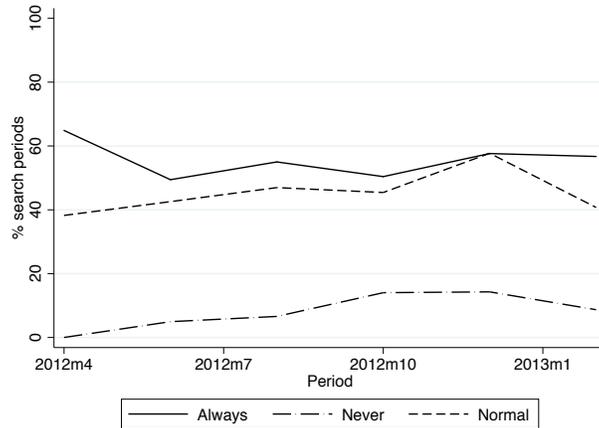


Figure 2: Percentage of applicants that was given a job-search period by default option, over time

caseworker.

In the remainder we refer to these default options as ‘never’, ‘always’ and ‘normal’. The default option ‘normal’ shows what the caseworkers would do with the applicant in absence of the experiment, which allows us to study targeting by caseworkers. The experimental period was divided into four periods of three months. Each period the caseworkers received a new default option which they had to apply to all new applicants assigned to them in this period. This allows us to control for business cycle effects. The randomization of default options over caseworkers took place at the level of the welfare office.

For the success of the experiment it is crucial that compliance to the default options is sufficiently high. Before the start of the experiment we organized meetings with all caseworkers to inform them about the experiment. At the start of every three-month period each caseworker was instructed individually about her new default option. Caseworkers were asked to fill in a form about the applicant at each intake meeting. The forms were personalized for each caseworker and had the period-specific default option printed on the form, such that they were constantly reminded of their current default option. We kept track of the inflow and regularly tried to visit caseworkers if they had not filled in the forms for new applicants or deviated substantially from their given default option. During the experiment, we visited the welfare offices almost weekly to answer questions from

caseworkers, pick up forms and to keep an eye on the implementation of the experiment.⁷

Compliance rates

Figure 2 shows the fraction of job-search periods applied per default option during the experimental period. The distinction between the three default options is most pronounced at the start of the experiment. The percentage of job-search periods given under the default ‘always’ remained relatively stable over time, while it increased for the other two default options. During the experiment period, we communicated with 112 caseworkers. Some caseworkers left, and new caseworkers entered the organization and took over their caseload. Furthermore a few times the caseload of a sick caseworker was taken over by a colleague. Such cases are associated with a larger noncompliance in both the default options ‘always’ and ‘never’. On average, during the experimental period, caseworkers with the default option ‘never’ gave a job-search period to nine percent of the applicants, caseworkers with default option ‘always’ gave a job-search period to 55 percent of the applicants, and caseworkers with default option ‘normal’ gave a job-search period to 46 percent of the applicants.

3 Data

Data sources

Our analysis employs data from three different sources, that are linked using unique identifiers for each individual. The welfare agency of Amsterdam provides administrative information on the date of registration at the welfare office, date of application for welfare, start and end date of collecting welfare benefits, whether a job-search period is applied and the identity of the caseworker that conducted the intake meeting. The individual

⁷The forms are filled in for 72 percent of the observations. Given that all information is also available through the administrative records (for the full sample), we do not use the information from the forms in our empirical analysis. Initially, we introduced the forms because we were not sure if the administrative system of the welfare office would also include applicants that applied for benefits but never returned after a job-search period. This turns out to be the case. However, for the experiment the forms were very useful as the period-specific default option was printed on them and it gave us a reason to regularly check caseworkers.

characteristics of applicants that are registered in these data are date of birth, gender, household composition and highest level of education. Furthermore, we observe the exact benefits payments. Second, we use data from the national social insurance administration to observe daily information for each individual on the amount of income from employment, hours worked and income from other benefit schemes.⁸ We have this information for all participants in the experiment from 2008 until October 2013. The retrospective nature of the data allows us to construct a labor-market history for all individuals, which is a control variable. Third, we link the data to individual records of all Dutch citizens kept by Statistics Netherlands. Using these records we determine whether an individual was suspect of a crime in a given year. The data from both Statistics Netherlands and from the national social insurance administration cover the full population of the Netherlands such that the experiment sample is matched without attrition.

Sample

Based on inflow in previous years we expected 2500 individuals to participate in the experiment. Our initial power calculation was based on this inflow number. Our final sample consists of 2860 welfare applications (2709 unique individuals).⁹ The economic downturn may explain the slightly higher inflow. 38 welfare applicants have an incorrect personal identifier, such that we can not match them to their outcomes. For eight applicants information on their caseworker is missing, so we can not determine under which default option they were treated. Furthermore, for 24 applicants information on the job-search period is incomplete, and for seven applicants information on essential controls is missing (gender, age or household situation). This means that in total we exclude 72 observations from the analysis, which leaves us with 2788 observations (2640 unique individuals). For 64 of the 72 excluded observations we have information on the default option. The excluded observations are evenly distributed among the three default options (joint p-value is 0.70).

⁸The other benefit schemes include unemployment insurance benefits and disability insurance benefits. We also observe whether someone receives welfare benefits in another municipality. Only data on income from self-employment are missing.

⁹An individual can have multiple applications if he/she applies for benefits multiple times within our experimental period. The average number of days elapsed between consecutive applications is 112.

Table 1: Characteristics of applicants under different default options

	Full Sample	Default option			P-value difference		
		Normal	Always	Never	Normal vs Always	Normal vs Never	Always vs Never
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Background characteristics</i>							
Female	39%	39%	41%	36%	0.45	0.44	0.07
Partner	11%	12%	9%	8%	0.79	0.68	0.57
Children	14%	14%	15%	13%	0.47	0.72	0.58
Age under 30	25%	24%	25%	27%	0.56	0.13	0.30
Age 31 - 36	25%	23%	28%	27%	0.07	0.12	0.82
Age 37 - 45	26%	27%	24%	24%	0.17	0.22	0.96
Age above 45	24%	25%	23%	22%	0.31	0.07	0.37
Bachelor/Master	28%	27%	28%	31%	0.66	0.59	0.52
Vocational	23%	24%	24%	21%	1.00	0.53	0.47
High school	13%	11%	16%	13%	0.06	0.85	0.08
Preparatory vocational	20%	21%	17%	20%	0.04	0.68	0.19
Primary education	10%	10%	11%	10%	0.35	0.29	0.94
Education missing	2%	3%	2%	1%	0.37	0.30	0.79
Annual income in previous 2 years (x1000 €)	13.6	13.5	13.4	14.2	0.31	0.95	0.34
<i>Treatment</i>							
Job-search period applied	40%	46%	55%	9%	0.00	0.00	0.00
Number of observations	2788	1657	571	560			

Note: The p-values in the last three columns are weighted by the office of registration, as randomization took place within welfare office.

The first column of Table 1 provides information about background characteristics. The majority of the applicants (over 60 percent) are male. Applicants are relatively young, the average age in the sample is only 38.4 years, with the median at 36 years. Recall that our sample includes only individuals older than 27 years. Young people have had less time to build up work history and, therefore, have in general shorter maximum unemployment insurance entitlement (they also have, on average, less wealth and are more often single). The average annual income in the two years before the welfare application is approximately 13,600 euro. Given that welfare is means tested and the income of the partner is taken

into account in the means test, couples are less likely to qualify for welfare. In our sample only 11 percent of the applicants have a partner, and 14 percent have children. The singles with children are almost exclusively women. Finally, 28 percent of the applicants have at least a bachelor's degree (compared to 35 percent in the total population aged 25-55 in the Netherlands).

Random assignment

Our design hinges on the fact that within local welfare offices applicants are randomly allocated to caseworkers, and are, therefore, also randomly assigned to default options. The second to fourth column of Table 1 show the mean characteristics of applicants under the default options 'normal', 'always' and 'never'. Columns five to seven show the p-value of the difference between two groups, for all different combinations. The characteristics are well balanced over the three treatment groups. There are no systematic differences and for only five out of 42 reported characteristics the difference is significantly different from zero. The lower panel in Table 1 shows the treatment probability for applicants in each of the three treatment groups and the number of observations per treatment group. The treatment group with the default 'normal' is the largest, as this was agreed upon with the welfare agency in the original research design.

Table A2 in the appendix shows the characteristics of only the applicants under the default option 'normal'. In this treatment group the decision to impose a job-search period was left to the caseworker. The first column describes the applicants that did not get a job-search period and the second column those that did get a job-search period. The third column gives the p-value of the difference. This table provides insight on how caseworkers target the job-search period in a non-experimental setting. The results in this table confirm that without experimental manipulation caseworkers target job-search periods., for example, to young applicants and those without children.

Descriptive statistics of the outcome variables

Table 2 displays information on the outcome variables. We mainly look at outcomes up to six months after registration. At that moment 50 percent of the applicants are still collecting welfare benefits. Average cumulative welfare payments are almost 3,000 euro in the six months after registration and average earnings are about 2,500 euro. Income from other benefits (mainly unemployment insurance benefits and disability insurance benefits) are 511 euro. Total income is the sum of these three income sources (welfare benefits, wage earnings and other benefits).

Job-search periods delay the first welfare benefits payment causing that people who have very limited wealth will be without income for some period. This may trigger criminal behavior. Therefore, we also consider crime as a relevant outcome. Our data are informative on whether someone was a crime suspect in a given year. The exact date of the crime is not reported. We consider crime outcomes in 2012 and 2013, which implies that for some people the crime might have taken place before applying for welfare benefits. However, given our randomized design there is no reason to suspect a difference in crime rates before the start of the experiment between groups. Table 2 shows that about nine percent of the applicants were suspect of a crime in 2012 or 2013.

4 Empirical strategy and graphical evidence

Empirical strategy

To estimate the effect of a job-search period on welfare receipt and other income variables we assume a linear relationship. The outcomes of individual i observed t time periods after registering for welfare benefits at time τ at welfare office w are denoted by $Y_{i\tau tw}$ and $SP_{i\tau w}$ is an indicator for a job-search period. Our regression model is specified as

$$Y_{i\tau tw} = \alpha_{\tau t} + \gamma_{wt} + \delta_i SP_{i\tau w} + X_i \beta_t + u_{i\tau tw} \tag{1}$$

Table 2: Descriptive statistics of the outcome variables

	Mean	SD
<i>Labor-market outcomes</i>		
Applicants on welfare after 6 months	51%	50
Cum. welfare benefits received over 6 months (in €)	2953	2413
Cum. earnings over 6 months (in €)	2494	3766
Cum. other benefits over 6 months (in €)	511	1715
Cum. total income over 6 months (in €)	5957	3653
Nr of weeks with benefits >0 over 6 months	16.2	10.8
Cum. hours worked over 6 months	190	266
Mean hourly wage 6 months after (in €, if wage>0)	13.3	6.3
<i>Crime outcomes</i>		
Suspect of a crime in 2012 or 2013	9%	28
Suspect of a property crime in 2012 or 2013	4%	20
Number of applicants	2788	

We estimate this model separately for different elapsed durations t since applying for welfare benefits. The vector X_i contains a set of covariates including age, gender, partner status, an indicator for children, cumulative income in the 24 months before registration and dummies for five education categories. $\alpha_{\tau t}$ and γ_{wt} are fixed effects for the quarter of registration and the local welfare office, which are allowed to differ for any duration. The former takes business cycle effects into account and the latter controls for possible differences between the local labor market in the city districts. The parameters of interest are δ_t which describe the effect of a job-search period t weeks after registration. Standard errors are clustered at the level of the applicant to account for multiple applications per individual.

Before, we showed that caseworkers usually target job-search periods towards younger workers and those without children. If caseworkers are more likely to impose a job-search period on applicants with better labor-market prospects, the OLS estimator of δ_t is biased. We exploit our experimental design to estimate the causal effect of the job-search period using two strategies. First, we replace SP_i by the default option of the caseworker that

conducted the intake meeting of individual i :

$$Y_{i\tau tw} = \alpha_{\tau t} + \gamma_{wt} + \delta_{1,t}Normal_i + \delta_{2,t}Always_i + X_i\beta_t + u_{i\tau tw} \quad (2)$$

Because compliance was not perfect, δ_1 and δ_2 are the intention-to-treat effects (ITT). The advantage of the ITT parameters is that they reflect the change in outcomes if the welfare agency moves from abandoning job-search periods to the current policy (δ_1), or to a more strict ‘always’ policy (δ_2). However, since the parameters average over all applicants (including those that did not receive a job-search period), they do not reflect the effect of actually imposing a job-search period. Therefore, we employ a second strategy, where we instrument SP_i with the default option of the caseworker that conducted the intake. We estimate a first-stage equation of the form:

$$SP_{i\tau w} = \kappa_{\tau} + \phi_w + \lambda_1Normal_i + \lambda_2Always_i + X_i\theta + v_{i\tau w} \quad (3)$$

In equation (3) λ_1 and λ_2 reflect the difference in the probability to receive a job-search period for caseworkers with the default options ‘normal’ and ‘always’, compared to the default option ‘never’. These parameters thus show the differences in compliance rates to the different default options. We saw before that caseworkers with the default option ‘normal’ (‘always’) give 36 percentage points (45 percentage points) more job-search periods than caseworkers with the default option ‘never’.

Three key assumptions underlie our empirical strategy. First, for the default options to be valid instruments, applicants’ assignment to a caseworker must be uncorrelated with unobserved characteristics that can influence labor-market outcomes (conditional on observed characteristics). As discussed before, applicants are assigned to caseworkers in a process that is unrelated to applicant characteristics. Because applicants are assigned to caseworkers within a local welfare office it is crucial to control for welfare office fixed effects in our analysis. Otherwise differences in the number of job-search periods applied under the default options could reflect differences in populations of applicants between districts, for example, arising from differences in local labor-market conditions.

A second assumption that is needed for the interpretation of the instrumental variables estimates is instrument monotonicity. No individual would have received a job-search period from a caseworker with default option ‘never’ and would not have received a job-search period from a caseworker with default option ‘always’. This is very likely to hold for the same caseworker. Formally, we see that within a local welfare office some caseworkers have a higher job-search period rate under ‘never’ than (other caseworkers) under ‘normal’ or ‘always’. However, in our experiment some caseworkers only have a small number of applicants per default option and observed differences in the fraction of applied job-search periods can also reflect differences in the average characteristics of the applicants. Furthermore, recall that only very few job-search periods are applied under the default option ‘never’. Therefore, it is likely that an applicant who would receive a job-search period under ‘never’ would also get this under any of the other options in which case the monotonicity assumption holds.

Finally, the probability that an individual finds employment (with or without a job-search period) should not be related to whether other individuals receive a job-search period (stable unit treatment value assumption (SUTVA)). Our design increased the probability for some applicants to receive a job-search period, but decreased it for others, so on average approximately the same amount of job-search periods were given as before the experiment started. Furthermore, the treated population in the experiment is only a small fraction of the total population of unemployed workers in Amsterdam, which consisted of around 42,000 individuals in 2013. Therefore, it is not likely that the applied job-search periods in our population have substantial spillover or general equilibrium effects.

Using instrumental variables, we estimate a Local Average Treatment Effect (LATE) (Angrist et al., 1996). The effect of a job-search period is only identified for the group applicants for which the caseworkers complied to their default options. In section 7 we elaborate further on the definition of compliers in the setting of the experiment and the interpretation of our estimated effect of a job-search period.

Graphical evidence

We start with a graphical description of how labor-market outcomes are related to the three default options, before turning to a more detailed regression analysis. Figure 3 presents the evolution of the fraction of applicants receiving welfare benefits by default option. Three issues are important before interpreting the figure. First, an individual is counted as receiving benefits in a certain week if payments of benefits were made that are assigned to that week. For example, if an applicant returns after a job-search period and receives benefits retrospectively from the moment of initial registration, this person is counted as benefits recipient from registration onwards (even though the first payment took place after eight weeks). A mechanical effect of the job-search period, the delayed payment of the first benefits, can, therefore, not be a driver of possible effects. Second, from the figure it is clear that take-up of welfare benefits is less than 100 percent for all three default options. This arises because eligibility for welfare benefits is only determined if the application for benefits is activated, so after the intake meeting and a possible job-search period. Our sample describes individuals who registered for benefits and who had an intake meeting, which was the moment of randomization. Conditioning on actual welfare-benefits entitlement leads to possible confounding effects, since for individuals who do not re-apply for welfare benefits after a job-search period, welfare-benefits entitlement is not determined. Third, we also see that during the first five weeks the fraction of benefits recipients increases for all three default options. The increase is due to people who register at the welfare office before the date of exhaustion of UI benefits. This is advised by the unemployment insurance office to prevent financial problems due to the processing time in which no welfare benefits are received.

Figure 3 shows that under the default option ‘never’ the fraction of people receiving welfare benefits is higher than under the default ‘always’. The fraction receiving welfare benefits under the default option ‘normal’ lies in between, but is closer to the default option ‘always’. This suggests that a job-search period has a substantial effect on the probability to receive welfare benefits. Over time the differences between the three default options decrease, but 26 weeks after registration applicants under the default option ‘never’ are

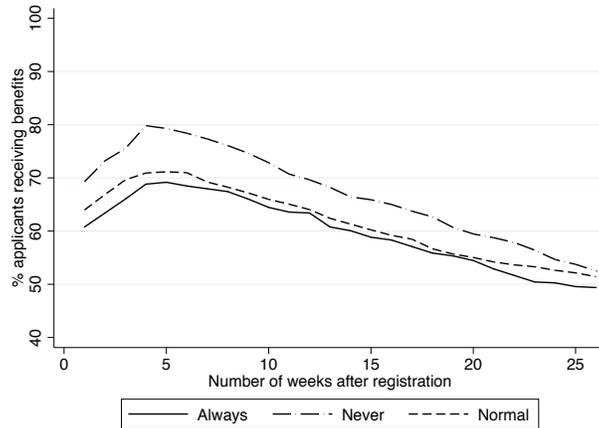


Figure 3: Fraction of applicants receiving welfare benefits by default option

still more likely to receive benefits.

5 Results

The discussion of our empirical results is divided into five parts. We first present estimates of the effect of a job-search period on the likelihood to receive welfare benefits. We continue with estimates of the effects on earnings and other benefits, followed by a discussion of the impact on crime. Finally, we look at the long-term effects of a job-search period.

Welfare benefits

In Figure 4a we plot the point estimates and 90 percent confidence intervals of the intention-to-treat effect on the probability to receive welfare benefits (0/1) estimated separately for each week after registration (following equation (2)). In this figure, the default options ‘always’ and ‘normal’ are compared to the default option ‘never’. Individuals with a caseworker with default option ‘always’ have a ten percentage points lower probability to collect some welfare benefits. Over time the effect becomes slightly smaller. Individuals with a caseworker with the default ‘normal’ have a six percentage points lower probability to receive welfare. More than 20 weeks after registration both effects are still significantly different from zero. The results are summarized in Table 3, which reports the effects of the

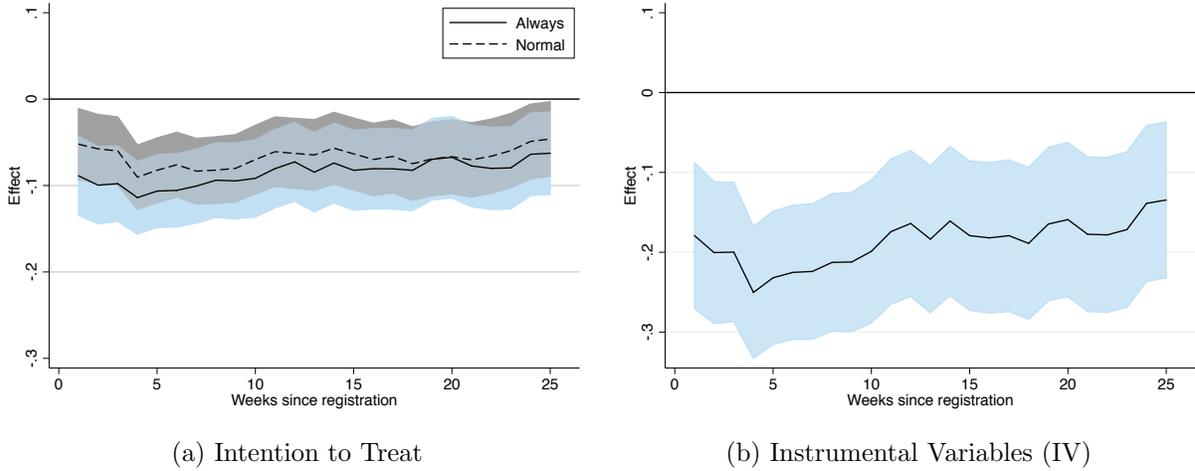


Figure 4: Probability to receive welfare benefits (colored areas are 90 percent confidence intervals)

default options on cumulative outcomes half a year after registration. The first and second column provide estimates that only include controls for calendar time and welfare office, the third and fourth column also control for characteristics of the applicant. The estimates hardly change when including these additional control variables, which is expected since the default options are randomly assigned. Table 3 shows that compared to the default option ‘never’, applicants in the default group ‘normal’ receive benefits, on average, 1.72 weeks shorter. For applicants in the group ‘always’ the effect is larger and amounts to a reduction of 2.18 weeks. Recall that this is not a mechanical effect of the job-search period, since entitlement to benefits starts at the date of initial registration.

To get an idea about the size of the effect of a job-search period, we next estimate the effect using our instrumental variables approach. The first-stage estimates of the default options on the probability to impose a job-search period are respectively 0.34 (s.e. 0.02) for the default ‘normal’ and 0.46 (s.e. 0.02) for the default ‘never’, with an F-statistic for joint significance of the instrumental variables equal to 237. Figure 4b plots the point estimates of the instrumental variables estimates of the probability to receive welfare. A job-search period lowers the probability to receive welfare benefits with around 20 percentage points in the first ten weeks. Given that in the group with the default option ‘never’ total take up of welfare benefits is around 80 percent this implies a reduction of about 25 percent.

Table 3: Effect of job-search period on cumulative outcomes 26 weeks after initial registration

	Intention to treat		Intention to treat		IV
	Always	Normal	Always	Normal	
	(1)	(2)	(3)	(4)	(5)
Weeks on welfare	-2.05*** (0.63)	-1.64*** (0.56)	-2.18*** (0.62)	-1.72*** (0.56)	-4.78*** (1.25)
Benefits received (in €)	-372*** (142)	-260** (127)	-389*** (141)	-261** (126)	-814*** (287)
Earnings (in €)	348 (225)	291 (197)	407* (220)	342* (195)	909** (449)
Other benefits (in €)	-50 (98)	-51 (88)	-42 (95)	-68 (87)	-122 (199)
Total income (in €)	-74 (218)	19 (193)	-25 (213)	13 (189)	-27 (438)
Weeks with earnings	0.71 (0.35)	0.35 (0.54)	0.75 (0.60)	0.40 (0.54)	1.48 (1.23)
Hours worked	26* (16)	18 (14)	30* (16)	21 (14)	64** (31)
Hourly wage (in €) (conditional on work)	0.71 (0.60)	0.30 (0.43)	0.72 (0.58)	0.26 (0.43)	1.78 (1.33)
Observations	2788		2788		2788
Included controls:					
Calendar time fixed effects	Yes		Yes		Yes
Local office fixed effects	Yes		Yes		Yes
Applicant characteristics	No		Yes		Yes

Note: Columns (1) and (2) and columns (3) and (4) in each row represent each one regression. The fifth column represents a separate regression. The applicant characteristics are age at registration, gender, household composition, cumulative income in 24 months before registration and dummies for five education categories. Standard errors are clustered at the level of the applicant. The outcome variable hourly wage is not a cumulative outcome, it refers to the average hourly wage in 26 weeks after registration. *** =significant at 1% level, ** =at 5% level, * =at 10% level

After the tenth week the effect somewhat decreases to 11 percentage points in week 26. The effect of a job-search period is still significantly different from zero 26 weeks after registration. This implies that the job-search period does not only have a short-run effect, but also has a more profound and longer lasting effect on welfare uptake. The first row of Table 3 shows that a job-search period reduces the period of receiving welfare benefits with 4.8 weeks.

Figure 4b uses a binary variable for receiving welfare benefits. Welfare recipients are,

however, obliged to also accept part-time jobs and, therefore, also partial outflow can take place. To look at the total impact of the job-search period on welfare receipt Figure 5a shows the effect on the amount of welfare benefits received by an applicant. The pattern is quite similar to the pattern in Figure 4b with the binary welfare variable. The job-search period has a strong effect on welfare receipt that is long lasting and only becomes insignificant after 24 weeks. On average, about 30 euro per week is saved on welfare benefits if a job-search period is imposed. Table 3 shows that total welfare benefits payments in the first six months decrease with 814 euro, which amounts to a 25 percent reduction of the mean cumulative amount of welfare during the first 26 weeks.

Income from wages and other benefits

Not receiving welfare benefits does not necessarily imply that someone is employed. A job-search period can discourage applicants to apply for welfare because of increased application costs and the complexity of the application process. The remainder of Figure 5 shows the effect of a job-search period on other sources of income. In particular, Figure 5b shows the effect of the job-search period on weekly income from employment. A job-search period has a positive effect on earnings of about 30 euro per week that becomes significant after five weeks. After 14 weeks the effect steadily increases to 50 euro a week. Table 3 shows that during the first half year after registration imposing a job-search period, on average, induces individuals to earn in total 909 euro more. These additional earnings thus completely compensate (112 percent) the forgone welfare benefits of a job-search period.¹⁰

A spillover effect of a job-search period can be that individuals try to get income from other benefit schemes or apply for welfare benefits in a different municipality. We do not expect large effects here, as welfare should be the safety net and people can only apply if they are not entitled to any other benefit scheme. Furthermore, to apply for welfare benefits in a different municipality, people need to move. Figure 5c shows the

¹⁰Recall that our experiment took place during an economic downturn (see Figure B1 in the appendix). The effects can be different when labor-market conditions change. During the second half of the experimental period, there was a sharp increase in the inflow into welfare benefits. However, we find that the estimated effects of a job-search period are the same for applicants in both time periods. Within this limited time frame, we, therefore, do not find strong evidence that the effects vary with labor-market conditions.

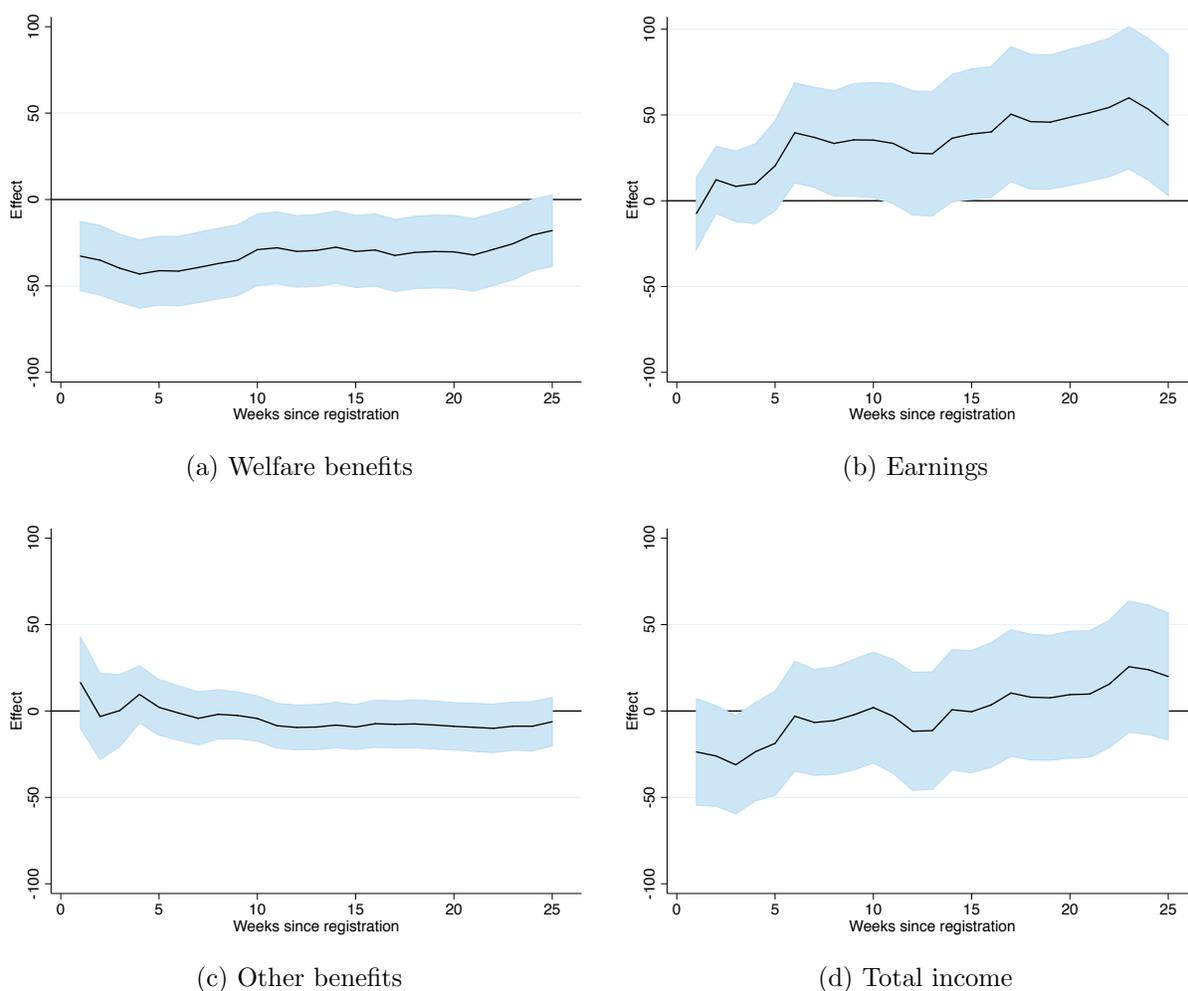


Figure 5: Instrumental variable estimates of the effects of a job-search period on labor-market outcomes $t - \tau$ weeks after registration (colored areas are 90 percent confidence intervals)

estimated effect on income from other benefit schemes and these are indeed small and insignificant. This is confirmed by the insignificant effect on cumulative other benefits in Table 3. Finally, the effect of a job-search period on total income (the sum of income from welfare, wages and other benefits) is shown in Figure 5d. During the first four weeks the effect on total income is negative and (almost) significant. After that, the effect is close to zero and insignificant. The effect on cumulative total income is small and not significantly different from zero (see Table 3). The negative effect during the first four weeks might be caused by individuals that find employment during the job-search period, and, therefore,

never collect any welfare benefits.

The increase in earnings can be driven by three different channels. First, an increased probability to be employed, second, an increase in hours worked, and third, an increase in the hourly wage. In the remainder of Table 3 we look at several outcomes to distinguish these channels. A job-search period does not have a significant effect on the number of weeks that an applicant has a non-zero wage, although the point estimate is positive. The number of hours worked is, however, significantly larger after a job-search period. In particular, imposing a job-search period increases the total working hours with, on average, 64 hours, an increase of 36 percent. This signals that a job-search period mainly increases the likelihood to find a full-time job compared to finding a part-time job in which the applicant does not earn enough to leave benefits. Finally, we estimate the effect on the hourly wage (conditional on having a job).¹¹ A job-search period can reduce the quality of the job that individuals are willing to accept, in case they are liquidity constraint and quickly accept a job to have income. On the other hand, receiving benefits can have a stigma effect on future employers by giving a bad signal about the quality of the employee. In that case, the reduced likelihood to receive benefits can translate into a positive effect on the hourly wage. The positive estimated effect on wages reported in Table 3 points in the direction of the second explanation, although the effect is not significant.¹²

Crime

A job-search period increases the time that an individual has to bridge before receiving the first welfare-benefits payment. Applicants of welfare benefits are unlikely to have access to savings or credit that can cover such temporary cash shortfalls. The policy is not to impose job-search periods on welfare applicants with severe financial problems. But, individuals

¹¹We condition on having a job in order to abstract from the possible employment effect of a job-search period. If a job-search period has a positive effect on the employment probability including the zeros could wrongly lead to the conclusion that a job-search period leads to a better paid job. Conditioning on having a job implies that there is a possible composition effect. If more people with a job-search period work and also those with a lower wage potential work, the results for hourly wage are biased downwards.

¹²A third explanation could be an experience effect. Applicants with a job-search period find a job earlier and the longer experience is reflected in the hourly wage. However, given the short time frame (26 weeks) this explanation is not likely to play a large role.

with low earnings prospects may still turn to crime to supplement their income during the job-search period. Foley (2011) finds that in the US crime rates increase in the amount of time that has passed since welfare payments occurred. Table 2 shows that in our sample nine percent has been suspect of a crime in 2012 and/or 2013 (compared to approximately two percent of the total population aged between 27 and 65) which indicates that for our population crime is not an irrelevant outcome.

Before turning to the results a few things need to be mentioned. First, by the nature of the data, we only consider registered crime, which likely underestimates actual crime. There is, however, not a clear reason to suspect a systematic difference between treatment groups. Second, we only know whether an individual was registered as a suspect of a crime, not whether she was actually convicted. However, in the Netherlands, on average, 90 percent of the registered suspects are declared guilty (Statistics Netherlands et al., 2013), such that this is a very strong indicator for actually having committed a crime. Third, we do not know the exact date of the crime, just the year in which the crime was registered. This means that we cannot consider the exact time elapsed between registration for benefits and the crime. We take as an outcome whether an individual was suspected of a crime in the year 2012 and/or 2013 (remember that the experiment started in April 2012 and ended in March 2013). This means that the crime could have taken place before the job-search period started. Given the randomized design there is no reason to suspect a difference in crime rates before the start of the experiment.

We find no evidence that a job-search period increases crime rates. With a point estimate of -0.02 (s.e. 0.04), the instrumental variables estimate of a job-search period on total crime is both not significantly different from zero and has the wrong sign. Following Foley (2011), we also look only at property crimes, to separate out crimes that have a financial motivation. For these crimes the point estimate is -0.01 (s.e. 0.03). To check whether the effects are sensitive to the definition of our outcome measure we repeat the analysis by period of registration, and only for crimes committed in 2013. Again, we find no evidence of an effect of the job-search period on crime.

Long-run effects

So far, we have mainly looked at outcomes up to half a year after registration. In Figure 4b we saw that 26 weeks after registration the effect of a job-search period becomes insignificant. A job-search period can still have long-run effects. If a job-search period causes individuals to accept temporary low-wage jobs and individuals without a job-search period find more stable employment, then the long-run effect of a job-search period can be negative. Table 4 presents the instrumental variable estimates for 26, 52 and 78 weeks after registration. For 78 weeks the sample is smaller because in the current dataset not everyone is observed in the data for such a long period.

Table 4 shows that, although precision decreases due to larger standard errors, the point estimates are fairly stable over time. The main effect of a job-search period thus takes place during the first 26 weeks, and after that not much changes, neither in a positive nor in a negative way. The reductions in welfare-benefits payments are, therefore, permanent savings, that are not offset by a later increase in benefit dependency.¹³ One thing that stands out is the positive and significant effect on the hourly wage one year after registration, an increase of 21 percent. This result contradicts the idea that a job-search period induces applicants to accept lower quality jobs. A possible explanation is that collecting welfare benefits has a negative stigma. Another possible explanation is that during the job-search period individuals can decide themselves which jobs to apply to, which may result in a better match than when the welfare agency assists in the job-search process.

6 Heterogeneous treatment effects

In this section we explore heterogeneity in the effect of a job-search period. First, we consider three important determinants of labor-market outcomes: age, gender and education. Second, we study variation in the effect by looking at the income distribution. This is motivated by the concern that although a job-search period might, on average, have

¹³To know whether this also holds for the really long run we have to repeat our analysis in the future. Because our data are linked to the microdata database of Statistics Netherlands, this can be realized later.

Table 4: Instrumental variable estimates for the effect of a job-search period on cumulative long-run outcomes

	Weeks since registration		
	26 weeks (1)	52 weeks (2)	78 weeks (3)
Weeks on welfare	-4.78*** (1.25)	-5.87** (2.47)	-8.40* (4.98)
Benefits received (in €)	-814*** (287)	-824 (536)	-971 (1063)
Earnings (in €)	909** (449)	965 (971)	765 (2069)
Other benefits (in €)	-122 (199)	-7 (390)	542 (856)
Total income (in €)	-27 (438)	133 (938)	337 (2020)
Weeks with earnings	1.48 (1.23)	0.42 (2.39)	0.10 (4.82)
Hours worked	64** (31)	62 (68)	40 (144)
Hourly wage (in €) (conditional on work)	1.38 (1.17)	1.76** (0.79)	1.72 (1.26)
Observations	2788	2788	1399
Included controls:			
Calendar time fixed effects	Yes	Yes	Yes
Local office fixed effects	Yes	Yes	Yes
Applicant characteristics	Yes	Yes	Yes

Note: Each cell represents one regression. The applicant characteristics are age at registration, gender, household composition, cumulative income in 24 months before registration and dummies for five education categories. Standard errors are clustered at the level of the applicant. The outcome variable hourly wage is not a cumulative outcome, it refers to the average hourly wage in 26, 52 or 78 weeks after registration
*** =significant at 1% level, ** =at 5% level, * =at 10% level

a positive effect on labor-market outcomes, there can still be a group that is seriously harmed by a job-search period.

Gender, age and education

We distinguish between three education levels, at least a bachelor degree (28 percent of the sample), a basic qualification (36 percent) and less than a basic qualification (36

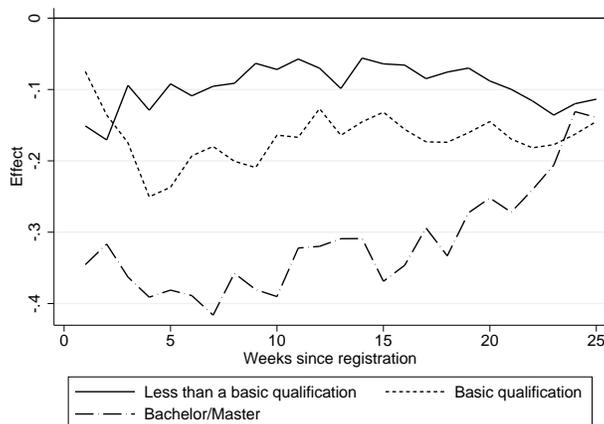


Figure 6: Effect on receiving welfare benefits by level of education

percent).¹⁴ Next, we estimate a model with interactions between these education levels and treatment. Figure 6 shows the effects on receiving welfare benefits for three education levels. For ease of presentation, the confidence intervals are not shown. The effects are significant for almost all weeks for applicants with at least a bachelor degree and those with a basic qualification. For the group with less than a basic qualification the point estimates are also negative but the effect is not significant.

The effect of a job-search period on the probability to receive welfare benefits monotonically declines with the level of education, and for the group with at least a bachelor degree the effect is significantly different from the effect for applicants with less than a basic qualification. For individuals with at least a bachelor degree a job-search period reduces the probability to receive welfare benefits with almost 40 percentage points, translating into a 50 percent reduction in the uptake of welfare benefits.

Table A3 in the appendix reports the effects on the other (cumulative) outcomes for the different education groups. For most outcomes the effect sizes monotonically increase with the level of education, and have the same sign for all three groups. For the number of weeks on welfare benefits and total welfare benefits received the estimates for the group with at least a bachelor degree are significantly different from those for applicants with less

¹⁴A basic qualification is the government definition for the minimum level of education needed to be self-sufficient on the labor market. Such a qualification requires at least senior general secondary education, pre-university education, or level-2 of senior secondary vocational education.

than a basic qualification. Overall the results by education level suggest that a job-search period is most effective for individuals that have better prospects on the labor market, as reflected by their level of education.

We do not find any differential effects when we consider gender and age. Figure B2 in the appendix shows the effect of a job-search period on the probability to receive benefits split by gender (left) and above/below median age in our data (right). In both cases the estimated coefficients are very similar. The lack of differential effects on these dimensions can be a result of the characteristics of the target population. Compared to the general population, welfare applicants are quite young (older workers have longer entitlement to unemployment insurance benefits) and are less likely to have an (income-generating) partner. Furthermore, job-search periods are only applied to individuals for whom there are no direct restrictions to work.

Distributional impacts

In section 5 we found that the loss in benefits is completely (112 percent) compensated by an increase in earnings. However, given that the minimum wage for a full-time job (approximately 1200 euro per month) is substantially higher than the level of welfare benefits the earnings gain can be unequally distributed along the income distribution. If that is the case, a job-search period can still be harmful for part of the applicants. To inquire this further we estimate the marginal distribution of the outcome under different treatments for the subpopulation of compliers, following Imbens and Rubin (1997).¹⁵ For sake of representation we take the default options ‘always’ and ‘normal’ together, thereby reducing the instrument to a binary instrument. The analysis could, however, be extended to consider both instruments separately.

Figure 7 plots the estimated distributions of cumulative income for applicants that did or did not receive a job-search period because their caseworker complied with the default options. Both 26 and 52 weeks after registration we see that the income distribution of the ‘treated’ compliers shifts slightly to the right. The treated compliers are, however,

¹⁵This method is briefly explained in appendix C.

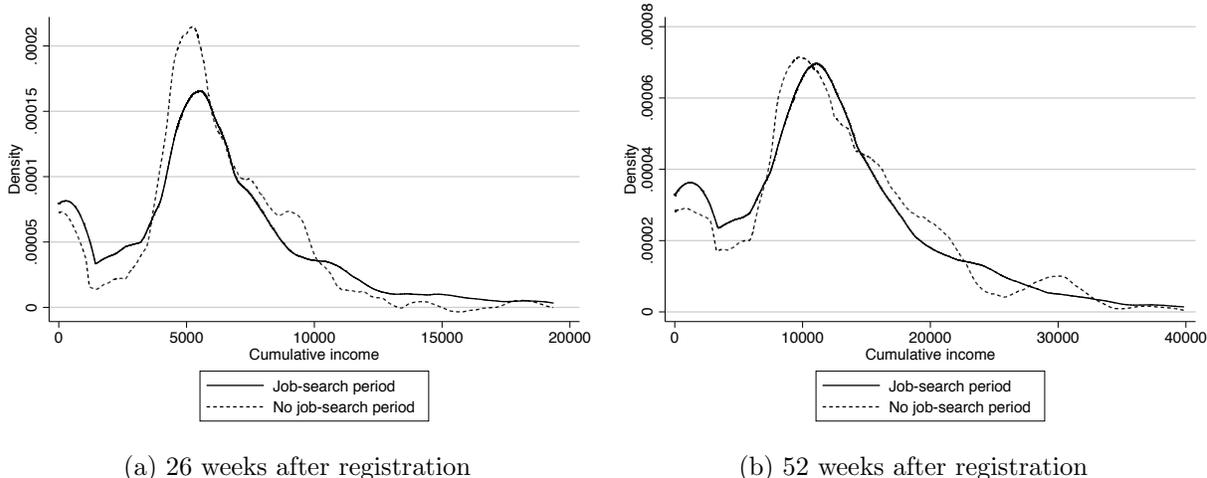


Figure 7: Income distribution of compliers to the default options

more likely to have a very low (close to zero) cumulative income. This suggests that for the majority of the applicants a job-search period has a positive effect on income because they find a job in which they earn more than the benefits level. However, for a small fraction of the applicants a job-search period leads to a higher probability of having a very low income. Regression results (not reported) confirm that during the first ten weeks after registration a job-search period leads to a higher probability to have an income below 150 euro per week. This is below the welfare-benefits level so after a job-search period some individuals neither have earnings nor receive benefits. After ten weeks this effect is no longer significant. This indicates that job-search periods cause that a few additional individuals have a very low income for some weeks after registration.

In the previous subsection we reported that the effect of a job-search period is particularly large for highly educated applicants. Figure 8, therefore, estimates the complier distributions split by level of education. We see that the higher probability of having very low income is mainly driven by the highly educated applicants (with a bachelor and/or master degree). Our preferred explanation is that there are two effective signals of a job-search period. First, individuals are forced to actively look for employment during the job-search period. Second, once starting collecting benefits, all jobs should be accepted, even if the work requirements are far below the educational and experience level of

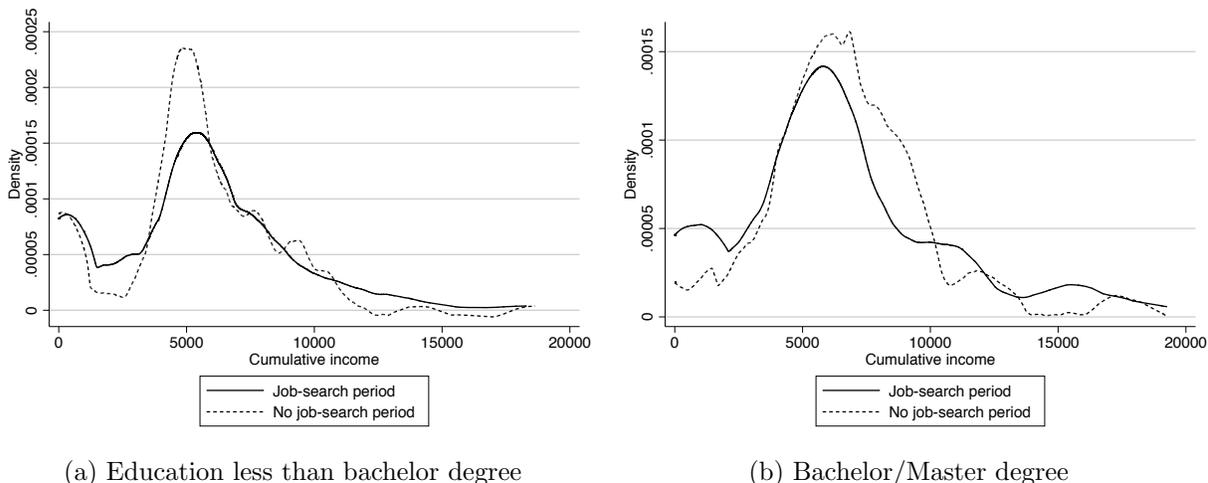


Figure 8: Income distribution (after 26 weeks) of compliers to the default options

the applicant. For highly educated individuals the second signal is more important than for individuals with lower education, which may even cause that some highly educated individuals choose no income over collecting welfare benefits.

7 Interpretation

The estimated effect of a job-search period should be interpreted as a local average treatment effect. Recall that caseworkers decide about imposing a job-search period and that decisions depend on the randomly assigned default options. In this section we provide some interpretation of the compliers for which we estimate the average treatment effect and test if the estimated effect changes when we consider other groups of compliers. The latter is also informative about how well caseworkers can target job-search periods.

Suppose it is possible to rank applicants according to an (unobserved) index which we refer to as the propensity *not* to receive a job-search period. This propensity can be based on characteristics observed by the caseworker, but not by us. Applicants that are always given a job-search period have propensity zero and applicants with a propensity of one never receive a job-search period. For ease of presentation we assume that these propensities are uniformly distributed. And for the moment we assume that there is no heterogeneity among caseworkers when executing the three default options. Then, the

default options can be translated into thresholds. Under the default option ‘never’ job-search periods are only assigned to applicants with a propensity less than 0.09 and these applicants are the always takers. Under the default option ‘always’ all applicants with a propensity less than 0.55 are assigned a job-search period. Therefore, applicants with a propensity above 0.55 are the never takers and those with a propensity between 0.09 and 0.55 are compliers. The default option ‘normal’ splits the compliers in two groups. First, applicants with a propensity between 0.09 and 0.46 who comply to both the default option ‘normal’ and ‘always’. And second, applicants with a propensity between 0.46 and 0.55, who only comply to the default option ‘always’.

If the effect of a job-search period is the same for all applicants, or if it only depends on applicants’ characteristics that do not affect the propensity, the average treatment effects will be the same for the compliers to the default options ‘normal’ and ‘always’ (Heckman and Vytlacil, 2001). However, there may be applicant-level heterogeneity in the response to the job-search period. For example, caseworkers may target job-search periods mainly to applicants for whom they expect the largest effects. In that case the effect is decreasing in the applicant’s propensity and the estimated effect depends on the thresholds chosen by the caseworker. This relates to Heckman and Vytlacil (2001), who develop a framework in which they express treatment parameters as different weighted averages of the marginal treatment effects. In our case this means that the average treatment effect is smaller for the compliers to the default option ‘always’ than for the compliers to the default option ‘normal’.

Nonparametric identification of the full set of marginal treatment effects requires an instrument that generates variation on the full support of the probability of treatment assignment. Our default options do not have this property.¹⁶ However, we consider two alternative approaches to investigate if the effect of a job-search period declines in the applicant’s propensity to receive a job-search period. First, we exploit that we have two default options deviating from ‘never’ (i.e. ‘normal’ and ‘always’), which generate different

¹⁶We could exploit differences in rates at which caseworkers assign job-search periods under the different default options (e.g. Maestas et al. (2013)). This yields much more variation, but in our case the average number of applicants per caseworker is low. This analysis would incorporate too much noise.

Table 5: Comparing the two instruments, outcomes 26 weeks after registration

	Total sample (1)	P-value over-id (2)	Normal vs Never (3)	Always vs Normal (4)
Weeks on welfare	-4.78*** (1.25)	0.88	-5.17*** (1.63)	-4.02 (4.59)
Benefits received (in €)	-814*** (287)	0.75	-723* (371)	-1063 (1024)
Earnings (in €)	909** (449)	0.83	979* (582)	733 (1650)
Other benefits (in €)	-122 (199)	0.61	-246 (258)	125 (689)
Total income (in €)	-27 (438)	0.84	10 (562)	-204 (1526)
Weeks with earnings	1.48 (1.23)	0.72	1.10 (1.60)	3.19 (4.51)
Cum. hours worked	64** (31)	0.91	64 (41)	87 (121)
Mean hourly wage (in €) (conditional on work)	1.38 (1.17)	0.49	0.55 (1.27)	3.30 (3.30)
First stage coefficient default normal	0.34*** (0.02)		0.34*** (0.02)	
First stage coefficient default always	0.46*** (0.02)			0.12*** (0.03)
Observations	2788		2217	2228
Included controls:				
Calendar time fixed effects	Yes		Yes	Yes
Local office fixed effects	Yes		Yes	Yes
Applicant characteristics	Yes		Yes	Yes

Note: Each cell represents one equation. The applicant characteristics are age at registration, gender, household composition, cumulative income in 24 months before registration and dummies for five education categories. Standard errors are clustered at the level of the applicant.

*** =significant at 1% level, ** =at 5% level, * =at 10% level

groups of compliers. And second, we increase the group of compliers by only considering caseworkers with high compliance rates to the default options.

For the first approach we perform overidentification tests. In the LATE-framework with two valid instrumental variables, rejection of the overidentification test indicates that treatment effects are heterogeneous (Angrist and Fernandez-Val, 2013). The intuition

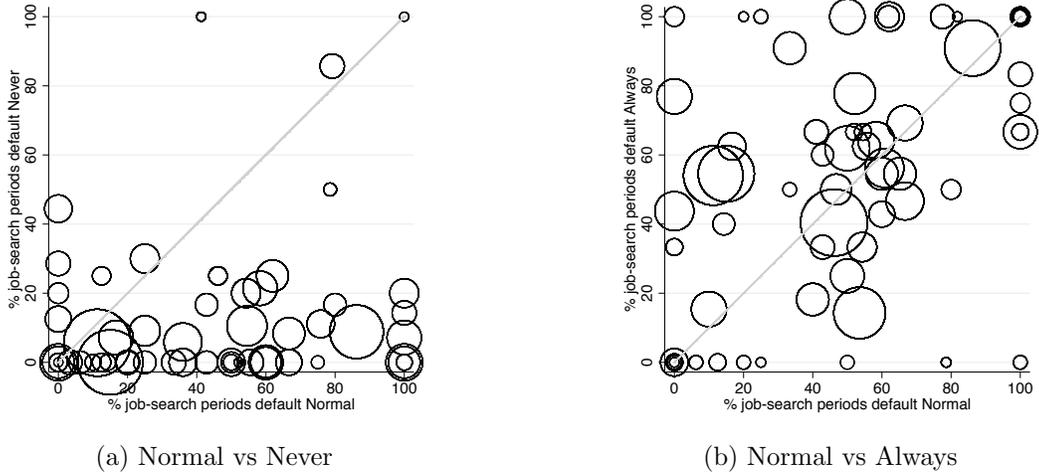


Figure 9: Percentage of job-search periods applied by default option and caseworker

is that both instrumental variables define different groups of compliers which may have different average treatment effects. The second column of Table 5 shows that the p-values for this overidentification test are for none of the outcomes below 0.49. This suggests either that both complier populations are very similar or that there is not much heterogeneity in the marginal treatment effects. Given that the default options ‘always’ and ‘normal’ yield complier populations which largely overlap, the first explanation is not unlikely.

Next, we estimate the effect of a job-search period using two different strategies. First, we instrument the job-search period with the default option ‘normal’ with as a reference the default option ‘never’. Second, we instrument the job-search period with the default option ‘always’ with as a reference the default option ‘normal’. The first strategy captures the average treatment effect for applicants with a propensity between 0.09 and 0.46, while the second strategy captures the average treatment effect for applicants with a propensity between 0.46 and 0.55. For the first strategy, we exclude all applicants under the default option ‘always’, and for the second strategy we exclude all applicants under the default option ‘never’. If caseworkers, indeed, mainly target job-search periods to applicants for which effects are highest, the estimated effects should be smaller when estimated under the second strategy. Table 5 presents the results for the total sample (column (1)) and for the two strategies (columns (3) and (4)). The estimates in the third column are not very

precise, due to less power in the first stage, but the point estimates are very similar to those in the second column. Again, there is no evidence that marginal treatment effects are decreasing in the propensity not to receive a job-search period.

This first approach exploits the difference in the rate at which job-search periods are applied under the default option ‘normal’ and ‘always’. Recall that the rates are not very different, which affects the power of the analysis. Furthermore, it can be that marginal treatment effects are very different for those applicants with much higher propensities. Therefore, we consider the individual compliance of each caseworker. Figure 9a shows for each caseworkers the rate at which they impose job-search periods under the default options ‘normal’ and ‘never’. Each circle represents a caseworker and the size of the circle describes the number of applicants a caseworker has under the default option ‘never’.

Three things come to the attention. First, there is substantial variation in the rate at which caseworkers assign job-search periods under the default option ‘normal’. There are also caseworkers who almost never apply a job-search period. Second, caseworkers who normally assign many job-search periods, substantially reduce this if they are assigned the default option ‘never’. And third, there are a few caseworkers who do not change behavior when being assigned the default option ‘never’, i.e. under this default option they impose job-search periods as often as under the default option ‘normal’.

Figure 9b shows the same figure but now comparing the default option ‘always’ with the default option ‘normal’. The figure is less pronounced than the previous figure. Obviously, many caseworkers find it difficult to impose job-search periods more often than usually. There are also some caseworkers who normally already impose job-search periods so often that this can hardly be increased when they are assigned the default option ‘always’. Finally, there are some caseworkers who more or less refuse to give job-search periods. Even under the default option ‘always’ they (almost) never apply job-search periods. Because within local offices applicants are randomly assigned to caseworkers, we can restrict our sample to caseworkers with substantial compliance rates without harming the randomization. In Table 6 we proceed in three steps and remove applicants of caseworkers who do not comply with the following rules:

Table 6: Outcomes for different groups of compliers

	Total sample (1)	Never<20% Always>40% (2)	Never<10% Always>40% (3)	Never<10% Always>60% (4)
Weeks on welfare	-4.78*** (1.25)	-4.28*** (1.21)	-4.66*** (1.26)	-5.72*** (1.55)
Benefits received (in €)	-814*** (287)	-647** (275)	-749*** (285)	-813** (326)
Earnings (in €)	909** (449)	-832* (429)	912** (440)	1386*** (518)
Other benefits (in €)	-122 (199)	-127 (184)	-86 (199)	-83 (239)
Total income (in €)	-27 (438)	58 (422)	77 (434)	490 (511)
Weeks with earnings	1.48 (1.23)	1.60 (1.17)	1.92 (1.23)	2.99** (1.39)
Cum. hours worked	64** (31)	54* (30)	58* (31)	83** (37)
Mean hourly wage (in €) (conditional on work)	1.38 (1.17)	1.63 (1.19)	0.81 (0.97)	1.05 (1.08)
First stage default normal	0.34*** (0.02)	0.37*** (0.02)	0.37*** (0.02)	0.39*** (0.03)
First stage default always	0.46*** (0.02)	0.55*** (0.03)	0.56*** (0.03)	0.63*** (0.03)
Observations	2788	2373	2207	1603
Included controls:				
Calendar time fixed effects	Yes	Yes	Yes	Yes
Local office fixed effects	Yes	Yes	Yes	Yes
Applicant characteristics	Yes	Yes	Yes	Yes

Note: Each cell represents one equation. The applicant characteristics are age at registration, gender, household composition, cumulative income in 24 months before registration and dummies for five education categories. Standard errors are clustered at the level of the applicant.

*** =significant at 1% level, ** =at 5% level, * =at 10% level

1. Fraction ‘never’ <20% & fraction ‘always’ >40%
2. Fraction ‘never’ <10% & fraction ‘always’ >40%
3. Fraction ‘never’ <10% & fraction ‘always’ >60%

We only remove observations if a caseworker had more than five applicants in the

relevant default options, such that we do not run the risk of removing caseworkers that had a very peculiar draw of applicants. A regression of an indicator for the different samples that remain under these selection criteria on applicant characteristics shows that there are no observable differences between applicants in the different groups (see Table A4 in the appendix). Only the indicator for the local welfare office is significant. This confirms that compliance to the experiment differ between local offices. Table 6 presents the estimated effects for the job-search periods for the different subsamples. At the bottom of the table we see that by removing non-complying caseworkers, the first-stage coefficients increase. In the most strictly defined sample (column 4) the estimates now cover propensities from 0.06 to 0.71 percent. Looking at the results, there is no evidence that the effect of a job-search period decreases with the propensity not to apply a job-search period. The effect on earnings even becomes slightly larger.

Overall both strategies do not provide evidence that the effect of a job-search period decreases in the propensity not to receive a job-search period. It also shows that the local average treatment effect are not sensitive to changes in the sample and the group of compliers. Obviously, caseworkers do not succeed in targeting job-search periods to those applicants for which the effects are largest. Finally, the results suggest that the estimated local average treatment is generalizable to a larger share of the sample.

8 Conclusion

In this paper we use a field experiment, in combination with detailed administrative data, to study mandatory job-search periods for new applicants of welfare benefits. Our empirical results provide evidence for a strong and persistent effect of a job-search period on receipt of welfare benefits. Six months after applying for welfare the total benefits payments are reduced by, on average, 25 percent. The reduced welfare benefits are fully compensated (112 percent) by increased income from employment and there are no spillovers to other benefit schemes. The effect of a job-search period increases with education, for applicants with at least a bachelor degree the likelihood to collect welfare benefits decreases with 50

percent.

The job-search requirement of a job-search period can increase the likelihood to find a job, but job-search periods can also increase the complexity of the welfare-benefits application. Our results suggest that the increased complexity of the application process does not hurt the most vulnerable applicants. In particular, we do not find evidence for negative side-effects, such as engagement in criminal behavior. Therefore, the increased job-finding and reduced welfare dependency indicate that a job-search period is an effective instrument for targeting welfare benefits to those people who need it most.

The job-search requirement and the waiting period, are easily transferable to other situations, which suggests that job-search periods can also be useful policy instruments for unemployment insurance and disability insurance. In addition, the administrative costs of imposing a job-search period are small and it is an early intervention that can prevent more costly interventions later during the period of benefits dependency. However, the population in our experiment has two important features. First, they do not have serious limitations to work. Second, all applicants are liquidity constrained, otherwise they would not be eligible for welfare benefits. These aspects of the applicant population might be essential for a successful implementation of a job-search period in other settings.

Randomized experiments with welfare applicants are still rare. We show that using an encouragement design which allowed for opting out in special cases, it is possible to evaluate (existing) policies for welfare applicants. The opt-out possibility has been important to obtain support of caseworkers, which ensured sufficient compliance to our randomization. Finally, the design of our experiment allowed us to also study whether caseworkers are able to target search periods to those clients for whom they are most effective. We find that this is not the case, a job-search period is effective for a larger share of the applicants than the population on which caseworkers normally impose a job-search period. Our results suggest that the welfare administration should instruct the caseworkers to more frequently apply job-search periods.

References

- Angrist, J. D. and Fernandez-Val, I. (2013). ExtrapolATE-ing: External validity and overidentification in the LATE framework. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics*. Cambridge University Press.
- Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455.
- Bhargava, S. and Manoli, D. (2015). Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11):3489–3529.
- Bitler, M. P., Currie, J., and Scholz, J. K. (2003). WIC eligibility and participation. *Journal of Human Resources*, 38:1139–1179.
- Black, D. A., Smith, J. A., C., B. M., and Noel, B. J. (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.
- Brien, M. J. and Swann, C. A. (1999). Prenatal WIC participation and infant health: Selection and maternal fixed effects. University of Virginia Dept. of Economics, unpublished manuscript.
- Card, D., Kluve, J., and Weber, A. (2010). Active labour market policy evaluations: A meta-analysis. *Economic Journal*, 120(548):F452–F477.
- Card, D. and Robins, P. (1998). Do financial incentives encourage welfare recipients to work? Evidence from a randomized evaluation of the self-sufficiency project. *Research in Labor Economics*, 17:1–56.
- Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. *Quarterly Journal of Economics*, 128(2):531– 580.

- Currie, J. (2006). The take-up of social benefits. In Auerbach, A. J., Card, D., and Quigley, J. M., editors, *Public Policy and the Income Distribution*. New York: Russell Sage Foundation.
- Currie, J. and Grogger, J. (2001). Explaining recent declines in food stamp program participation. In Gale, W. and Rothenberg-Pack, J., editors, *Brookings Papers on Urban Affairs*, pages 203–244.
- Duflo, E., Glennerster, R., and Kremer, M. (2007). Using randomization in development economics research: A toolkit. In Schultz, T. P. and Strauss, J. A., editors, *Handbook of Development Economics*, volume 4, pages 3895–3962. Elsevier.
- Foley, C. F. (2011). Welfare payments and crime. *Review of Economics and Statistics*, 93(1):97–112.
- Heckman, J. J. and Vytlacil, E. (2001). Policy-relevant treatment effects. *American Economic Review*, 91(2):107–111.
- Imbens, G. W. and Rubin, D. B. (1997). Estimating outcome distributions for compliers in instrumental variables models. *Review of Economic Studies*, 64(4):555–574.
- Ketel, N., Leuven, E., Oosterbeek, H., and van der Klaauw, B. (2015). The returns to medical school in a regulated labor market: Evidence from admission lotteries. *American Economic Journal: Applied Economics*. Forthcoming.
- Kleven, H. J. and Kopczuk, W. (2011). Transfer program complexity and the take-up of social benefits. *American Economic Journal: Economic Policy*, 3(1):54–90.
- Krueger, A. B. (1990). Incentive effects of workers’ compensation insurance. *Journal of Public Economics*, 41(1):73–99.
- Maestas, N., Mullen, K. J., and Strand, A. (2013). Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt. *American Economic Review*, 103(5):1797–1829.

O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.

Parsons, D. O. (1991). Self-screening in targeted public transfer programs. *Journal of Political Economy*, 99(4):859–876.

Statistics Netherlands, WODC, and Raad voor de Rechtspraak (2013). Criminaliteit en rechtshandhaving 2013. Report, Justitie in Statistiek.

Van der Klaauw, B. and Van Ours, J. C. (2013). Carrot and stick: How re-employment bonuses and benefit sanctions affect exit rates from welfare. *Journal of Applied Econometrics*, 28(2):275–296.

A Additional tables

Table A1: Characteristics of the target population

	Mean
<i>Financial situation</i>	
Good	26%
Average	48%
Bad	26%
<i>Estimated level of independence (by the caseworker) of the applicant</i>	
Very independent	40%
Somewhat independent	54%
Completely not independent	6%
<i>Estimated time (by the caseworker) until exit to employment</i>	
Less than 1 month	13%
2 to 3 months	39%
4 to 5 months	31%
6 months or more	17%
<i>Reason application for welfare</i>	
Lost job	16%
End self-employment	8%
End unemployment insurance benefits	34%
Other reason	42%
<i>Other characteristics</i>	
Non-western immigrant	56%
Ever received welfare before	47%

Note: The information in the first four panels of this table is taken from the forms that caseworkers filled in for the applicants that were part of the experiment. These forms were completed for 72 percent of the sample. The information in the lowest panel is for the complete sample and was derived by linking the experimental sample to administrative data of Statistics Netherlands. The degree of independence indicates whether the applicant is self-reliant and is able to independently search for work.

Table A2: Targeting job-search periods by caseworkers under default ‘normal’

	Job-search period		p-value diff
	No	Yes	
	(1)	(2)	(3)
Female	40%	37%	0.32
Partner	13%	11%	0.01
Children	16%	11%	0.00
Age under 30	20%	30%	0.00
Age 31 - 36	21%	26%	0.02
Age 37 - 45	30%	23%	0.00
Age above 45	29%	21%	0.00
Bachelor/Master	26%	27%	0.32
Vocational	24%	24%	0.90
High school	13%	9%	0.02
Preparatory vocational	22%	20%	0.36
Primary education	10%	10%	0.73
Education missing	1%	5%	0.00
Annual income in previous 2 years (x1000 €)	13.7	13.3	0.32
Number of observations	899	758	

Note: The p-values in the last three columns are weighted by the office of registration, as randomization took place within welfare office.

Table A3: Effects by level of education 26 weeks after registration

	Education level		
	Low (1)	Middle (2)	High (3)
Weeks on welfare	-2.49 (2.15)	-4.32** (1.98)	-7.94*** (2.16)
Benefits received (in €)	-277 (483)	-838* (470)	-1442*** (488)
Earnings (in €)	562 (633)	1113 (687)	1060 (953)
Other benefits (in €)	-283 (366)	-24 (298)	-66 (358)
Cum. total income (in €)	1 (617)	251 (673)	-447 (943)
Weeks with earnings	0.57 (2.12)	2.19 (1.79)	1.61 (2.10)
Cum. hours worked	36 (53)	101** (46)	46 (58)
Mean hourly wage (in €) (conditional on work)	-0.69 (1.80)	1.79 (1.50)	3.03 (2.14)
Observations	1011	1007	770
Included controls:			
Calendar time fixed effects	Yes	Yes	Yes
Local office fixed effects	Yes	Yes	Yes
Applicant characteristics	Yes	Yes	Yes

Note: Each row in this table represents one regression including interactions for the different subgroups. The applicant characteristics are age at registration, gender, household composition, cumulative income 24 months before registration and dummies for five education categories. Standard errors are clustered at the level of the applicant.

*** =significant at 1% level, ** =at 5% level, * =at 10% level

Table A4: Characteristics of different groups of compliers

	Nev.<20% Alw.>40% (1)	Nev.<10% Alw.>40% (2)	Nev.<10% Alw.>60% (3)
<i>Applicant characteristics</i>			
Female	-0.01 (0.01)	-0.01 (0.02)	-0.00 (0.02)
Partner	0.02 (0.02)	-0.01 (0.03)	-0.02 (0.03)
Children	0.02 (0.02)	0.02 (0.02)	0.01 (0.03)
Age 31 - 36	-0.02 (0.02)	-0.03 (0.02)	-0.01 (0.03)
Age 37 - 45	-0.01 (0.02)	-0.00 (0.02)	0.02 (0.03)
Age above 45	0.03 (0.02)	0.04** (0.02)	0.05* (0.03)
Bachelor/Master	-0.04* (0.02)	-0.01 (0.03)	-0.01 (0.03)
Vocational	-0.03 (0.02)	-0.00 (0.03)	-0.01 (0.03)
High school	0.02 (0.02)	0.04 (0.03)	-0.01 (0.04)
Prep. vocational	0.00 (0.02)	0.03 (0.03)	0.00 (0.03)
Education missing	-0.01 (0.01)	-0.01 (0.02)	-0.00 (0.02)
Income 2 yrs before (x1000 €)	0.00* (0.00)	0.00 (0.00)	0.00 (0.00)
<i>Location of registration</i>			
Southeast	0.08*** (0.02)	0.25*** (0.02)	-0.32*** (0.03)
North	0.03 (0.02)	0.12*** (0.03)	0.19*** (0.03)
Centrum/East	0.02 (0.02)	0.19*** (0.02)	0.11*** (0.03)
New West	-0.01 (0.02)	0.17*** (0.03)	0.02 (0.03)
<i>Quarter of registration</i>			
Quarter 2	0.02 (0.02)	-0.00 (0.02)	-0.03 (0.03)
Quarter 3	-0.01 (0.02)	-0.02 (0.02)	-0.05** (0.03)
Quarter 4	-0.00 (0.02)	0.01 (0.02)	0.02 (0.03)

Note: *** =significant at 1% level, ** =at 5% level, * =at 10% level

B Additional figures

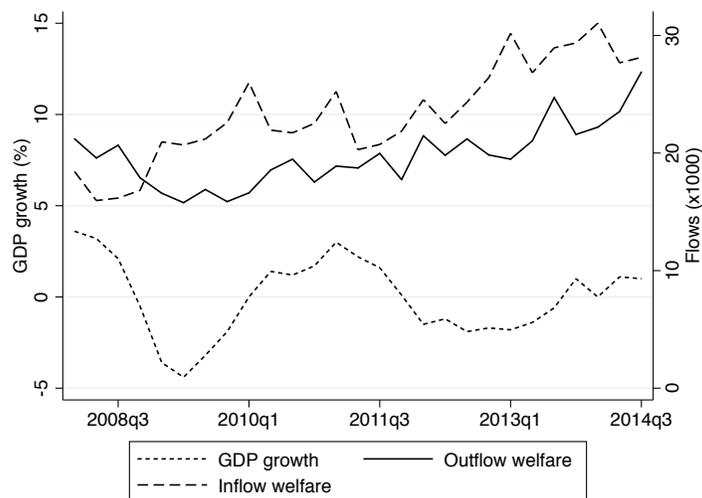


Figure B1: Inflow and outflow of welfare between 2008 and 2014 and macroeconomic indicators (source: Statistics Netherlands)

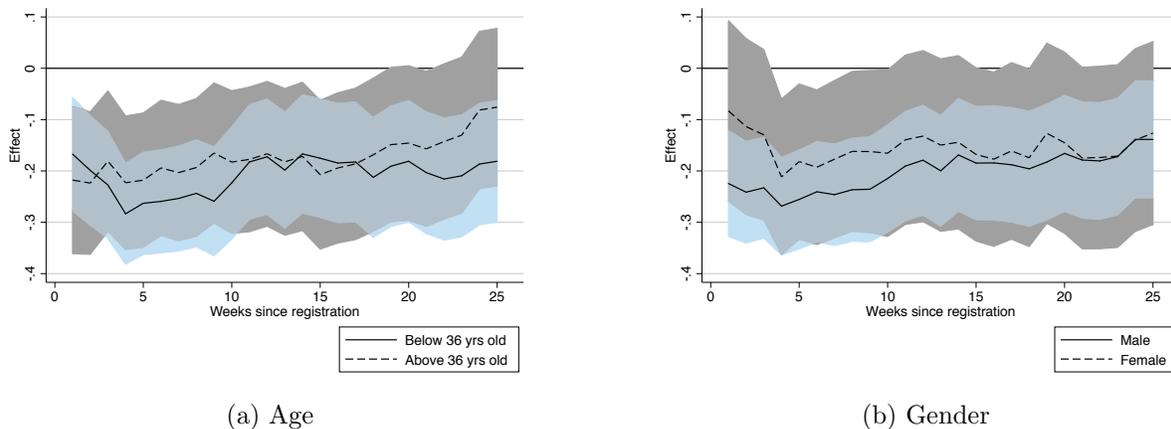


Figure B2: Instrumental variables estimates of the effects of a job-search period by age and gender $t - \tau$ weeks after registration (colored areas are 90 percent confidence intervals)

C Estimation of outcome distributions for compliers

Imbens and Rubin (1997) show how to estimate the marginal distribution of outcomes under different treatments for the subpopulation of compliers. In this section we briefly review their approach.¹⁷ $Y_i(1)$ and $Y_i(0)$ will denote the applicant’s total income with and without a job-search period, respectively. D_i denotes the default option of the caseworker to whom the applicant is matched, where $D_i = 1$ indicates the defaults ‘always’ or ‘normal’ and $D_i = 0$ the default ‘never’. SP_i indicates whether the applicant had a job-search period (0/1).

For each observation we observe the triple (D_i, SP_i, Y_i) , where Y_i is the observed outcome. From the data we cannot directly identify the compliers, but we can identify some always takers ($D_i = 0$ and $SP_i = 1$) and some never takers ($D_i = 1$ and $SP_i = 0$). Because of the randomization, the value of the default option will be independent of the applicant’s type, we can infer the distribution of $Y_i(0)$ for never takers and $Y_i(1)$ for always takers. These distributions are denoted by $g_a(y)$ and $g_n(y)$. The population proportions of compliers (ϕ_c), always takers (ϕ_a) and never takers (ϕ_n) can be identified from the data.

The distributions of interest are the distributions of $Y_i(0)$ and $Y_i(1)$ for compliers, described as $g_{c0}(y)$ and $g_{c1}(y)$. These cannot be directly observed from the data because the outcome distribution of applicants for whom $D_i = 0$ and $SP_i = 0$ consists both of never takers and compliers. Correspondingly, in the outcome distribution of applicants with $D_i = 1$ and $SP_i = 1$ there will be both always takers and compliers.

We denote the directly estimable distributions of Y_i for the subsample defined by $D_i = d$ and $SP_i = sp$ as $f_{d,sp}(y)$. This implies that $g_a(y) = f_{01}(y)$ and $g_n(y) = f_{10}(y)$. Imbens and Rubin (1997) show that the distributions for applicants that did or did not receive a job-search period because their caseworker complied to the default options can be expressed in terms of the directly estimable distributions in the following way:

$$g_{c0}(y) = \frac{\phi_c + \phi_n}{\phi_c} f_{00}(y) - \frac{\phi_n}{\phi_c} f_{10}(y), \quad (4)$$

¹⁷The notation in this appendix is partly taken from Ketel et al. (2015).

and,

$$g_{c0}(y) = \frac{\phi_c + \phi_a}{\phi_c} f_{11}(y) - \frac{\phi_a}{\phi_c} f_{01}(y). \quad (5)$$