**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series
ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 91

**The Effectiveness of Public Sponsored Training
Revisited: The Importance of Data and
Methodological Choices**

Martin Biewen, Bernd Fitzenberger, Aderonke Osikominu and Marie Paul

September 2012

The Effectiveness of Public Sponsored Training Revisited: The Importance of Data and Methodological Choices*

MARTIN BIEWEN,[‡] BERND FITZENBERGER,[‡] ADERONKE OSIKOMINU,[#] and MARIE PAUL[§]

[‡]University of Tübingen, IZA, DIW, martin.biewen@uni-tuebingen.de

[‡]Albert-Ludwigs-University Freiburg, ZEW, IZA, IFS, bernd.fitzenberger@vwl.uni-freiburg.de

[#]University of Zurich, Albert-Ludwigs-University Freiburg, CESifo, IZA, aderonke.osikominu@econ.uzh.ch

[§]née Waller, University of Duisburg-Essen, marie.paul@uni-due.de

This version: August 29, 2012

Abstract: As the first, substantive contribution, this paper revisits the effectiveness of two widely used public sponsored training programs, the first one focusing on intensive occupational training and the second one on short-term activation and job entry. We use an exceptionally rich administrative data set for Germany to estimate their employment and earnings effects in the early 2000s. We employ a stratified propensity score matching approach to address dynamic selection into heterogeneous programs. As a second, methodological contribution, we carefully assess to what extent various aspects of our empirical strategy such as conditioning flexibly on employment and benefit histories, the availability of rich personal information, handling of later program participations, and further methodological and specification choices affect estimation results. Our results imply pronounced negative lock-in effects in the short run in general and positive medium-run effects on employment and earnings when job-seekers enroll after having been unemployed for some time. We find that data and specification issues can have a large effect on impact estimates.

Keywords: public sponsored training, dynamic treatment effects, multiple treatments, kernel matching, administrative data

JEL: C 14, J 68, H 43

---

*We are indebted to Peter Mueser, Jeffrey Smith and participants in numerous seminars for many useful comments and suggestions. We thank Stefan Bender for his efforts to make available the data on which this study is based. The study is part of the project “Employment effects of further training programs 2000-2002 – An evaluation based on register data provided by the Institute for Employment Research, IAB (Die Beschäftigungswirkung der FbW-Maßnahmen 2000-2002 auf individueller Ebene – Eine Evaluation auf Basis der prozessproduzierten Daten des IAB)” (IAB project no. 6-531.1A, joint project with University of St. Gallen and the IAB). We gratefully acknowledge financial and material support by the IAB. Aderonke Osikominu gratefully acknowledges financial support from the German Excellence Initiative and the Austrian Science Fund (FWF grant no. S 10304-G16). The usual caveat applies.

Contents

1	Introduction	1
2	Review of the Training Literature	4
3	Background and Data	7
3.1	Training as Part of Active Labor Market Policy	7
3.2	Administrative Database Used	9
3.3	Evaluation Sample and Training Programs Analyzed	11
4	Econometric Approach	13
4.1	Evaluating Multiple Treatments in a Dynamic Context	13
4.2	Combining Exact Matching and Kernel Matching	15
4.3	Empirical Content of the DCIA	18
4.4	Specification of the Propensity Scores	20
4.5	Testing for Covariate Balance	21
4.6	Regression Adjustment and Test for Effect Heterogeneity	22
5	Which Program Works and for Whom?	24
5.1	Effectiveness of Training Compared with Searching	24
5.2	Heterogeneity of Training Impacts	27
5.3	Assessing the Quality of Our Benchmark Estimates	29
6	How Sensitive Are Impact Estimates to Data Features and Methodological Choices?	31
6.1	Sensitivity Analysis 1: Employment History	31
6.2	Sensitivity Analysis 2: Personal Characteristics	33
6.3	Sensitivity Analysis 3: Information on Other Programs	36
6.4	Sensitivity Analysis 4: 12-Months Evaluation Window	37
6.5	Sensitivity Analysis 5: Hypothetical Program Starts	39
6.6	Sensitivity Analysis 6: Future Program Participation	41
6.7	Statistical Significance of the Sensitivity Analyses	43
7	Summary and Conclusions	44

1 Introduction

There has been an enormous interest in the evaluation of active labor market policies in general, and of public sponsored training in particular, both in the US and Europe.¹ While earlier studies typically focused on the evaluation of a single program, recent developments in methodology and data access allow researchers to study the heterogeneity of treatment effects across different subgroups of people and the comparative effects of narrowly defined subprograms.² Detailed information on employment and earnings histories prior to program participation seems important to justify matching estimators for treatment effects that rely on a selection on observables assumption.³ Accurate longitudinal information on labor market transitions is also useful to account for the dynamics of program assignment and to carefully align treated and comparison units in their elapsed unemployment experience.⁴

While methodological progress is likely to improve the policy relevance of scientific evaluations, it dramatically increases the range of possible choices researchers can make. A better understanding of how training programs work and to what extent evaluation results hinge on particular data features and methodological choices is crucial to assess research findings based on different empirical strategies and data sets. This paper makes two contributions. On the substantive side, we revisit the effectiveness of two widely used public sponsored training programs – the first one involving comprehensive classroom further training and the second one focusing on short-term activation and reemployment – using a dynamic propensity score matching approach. On the methodological side, we undertake a detailed sensitivity analysis to investigate the importance of data features and specification choices as well as econometric modeling strategies in shaping estimation results.

Our first methodological contribution relates to the availability of rich data. We have access to unique data for Germany that merges information from different administrative registers. They contain precise and extensive information on individual

---

¹For comprehensive overviews see Heckman, LaLonde, and Smith (1999), Martin (2000), Martin and Grubb (2001), Kluve and Schmidt (2002), Carcillo and Grubb (2006), and Card, Kluve and Weber (2010).

²Examples of evaluations involving multiple comparisons of different programs are Lechner (2002), Gerfin and Lechner (2002), Larsson (2003), Hardoy (2005), Dyke et al. (2006), Dorsett (2006), Frölich (2008), and Sianesi (2008).

³See, e.g., Dolton and Smith (2011), Heckman et al. (1998), Heckman, LaLonde, and Smith (1999), Heckman and Smith (1999)).

⁴See, e.g., Crépon et al. (2009), Fredriksson and Johansson (2008), Lechner (1999), Osikominu (2012), Sianesi (2004).

employment and benefit histories covering more than ten years before and about 2.5 years after treatment start, detailed information on participation in all active labor market programs as well as detailed profiles of job-seekers (i.e. personal characteristics and goals of job search) that are reported by the caseworkers at the local employment agencies. Using different balancing tests, we first show that the information on personal characteristics and individual employment and benefit histories in our data balances individual characteristics and pre-treatment outcomes in our treatment and comparison groups well. We then carry out a careful sensitivity analysis that compares our benchmark evaluation results to different variations in which we omit important aspects of our data or our specifications. We believe that such an analysis provides important information for the evaluation of labor market programs in other countries. For example, many data sets used for evaluation purposes lack the information about what other programs comparison group members might participate in, e.g. that of Mueser, Troske, and Gorislavsky (2007). We investigate the sensitivity of the results if we proceed as if we did not have this information. Similarly, in the spirit of Card and Sullivan (1988), Dolton and Smith (2011), and Heckman and Smith (1999), we examine what difference conditioning on various aspects of the employment and benefit history makes. In order to assess to what extent evaluation results hinge on the availability of rich personal information, we compare our results to several specifications in which we omit information e.g. about the motivation of participants or on health and other characteristics which are often not available in data sets used for evaluation purposes.⁵

Our second methodological contribution relates to the implementation of a matching approach under dynamic treatment assignment. People who become unemployed may at some later point take up a job, participate in one of the training programs under investigation or in some other labor market program. According to many labor market policy regulations, as is the case also in Germany, these three options are to be viewed as competing risks. In particular, somebody who has found employment again is not eligible for program participation anymore. In this setting, a simple static evaluation approach in which unemployed receiving treatment are compared to unemployed who never receive treatment is not feasible. Defining non-

---

⁵Complementary to our first methodological contribution, Lechner and Wunsch (2011) examine the importance of different sets of conditioning variables using the same administrative data and propensity score matching in combination with matching on actual and hypothetical program starts. Their study involves an Empirical Monte Carlo analysis that presumes a known data generating process using real data and the validity of the conditional independence assumption. Our analysis, in contrast, investigates the sensitivity of the estimated treatment effects. In another Empirical Monte Carlo analysis, Huber, Lechner, and Wunsch (2010) investigate the performance of different propensity score matching estimators.

participants as never participants implies conditioning on future outcomes because for these persons the exit to employment has occurred before the potential start of treatment (Fredriksson and Johansson (2008)). Put differently, excluding comparison individuals who receive treatment in the future excludes possible sequences of events which may have occurred under the counterfactual situation that a treated individual had not started the program at the particular point of time observed. Similarly, we may only observe a treatment starting at some point of time during an unemployment spell because the unemployed did not happen to find a job before. Following Sianesi (2004), we define treatment parameters conditional on the unemployment experience at the time of treatment start and include all those in the comparison group for a particular treatment start who do not enroll at the elapsed unemployment experience under consideration. As a consequence, the future outcomes of comparison units may be affected by the effect of the treatment. In our sensitivity analysis we investigate the sensitivity of the estimated treatment effects to not using the information on future program participation and on participation in other labor market programs as well as to not aligning treated and comparison persons by elapsed unemployment duration. We further compare the results obtained in our benchmark approach to an alternative approach suggested by Lechner (1999). The latter assigns to each nontreated individual who does not receive the treatment of interest within a fixed time window of e.g. twelve months a hypothetical program start date that is simulated based on covariates observed at the start of the unemployment spell. A nontreated individual is then only used as comparison if he/she is still unemployed before the hypothetical starting date. In an attempt to align treated and comparison units in their prior unemployment experience, similarity of the actual and hypothetical program start dates is introduced as a further matching requirement. Our second methodological sensitivity analysis thus provides new insights on the importance of econometric modeling choices to control for dynamic selection effects.

Finally, as our substantive contribution, we assess the effectiveness of the two training programs during a time when resources were shifted away from longer-term training programs to short-term programs and to more job oriented programs (see Bundesagentur für Arbeit (2005c), and figure 1 below). Given the wide range of training programs implemented in Germany, our comparison of different forms of training may also hold lessons for other countries that employ similar programs. Short-term training is particularly interesting because, in addition to the provision of skills, it includes elements of job search assistance as well as profiling and monitoring of the unemployed (see the literature on such programs reviewed in the next

section). Our results show sizeable lock-in effects for all programs in the short run. In the medium-run, we find positive effects on employment and earnings when job-seekers enroll after having been unemployed for some time. Although both types of training programs generate significant earnings gains in the cases in which there are also positive employment effects, earnings gains induced by classroom further training attain €200 per month and are generally larger than those of short-term training (€70 per month for men and €175 for women). We find little evidence for additional impact heterogeneity after stratifying by gender and elapsed unemployment duration.

The remainder of the paper proceeds as follows. Section 2 reviews the training literature. Section 3 describes the institutional background and data. We introduce our econometric approach in Section 4. The evaluation results and our sensitivity analysis are presented in Sections 5 and 6. Section 7 concludes. The Online-Appendix provides complementary empirical results and background information.

2 Review of the Training Literature

The earlier literature for the US (see the reviews by Barnow (1987) Bloom et al. (1997) or Heckman et al. (1999)) distinguishes classroom training, on-the-job training, and work experience. Programs tend to be more effective for women than for men and on-the-job training tends to be more effective than classroom training and work experience. More recent evaluations for the US make use of advances in econometric methods, often investigate impacts over a longer time period, and put a stronger emphasis on the heterogeneous effects of different training programs (Dyke et al. (2006), Hotz, Imbens, and Klerman (2006), Heinrich et al. (2010)). These studies investigate programs involving job search assistance, improvement in job readiness, or training and they compare longer to shorter programs. Dyke et al. (2006) and Hotz et al. (2006) analyze training programs for a specific group of welfare recipients, whereas our study concerns training among a sample of prime-age unemployed who were employed in the recent past. Heinrich et al. (2010) find positive medium-run to long-run effects for the broader group of participants in programs under the Workforce Investment Act. Their findings imply higher gains for participants in general programs for adults than for participants from the group of displaced workers.

In a study for Canada, Park et al. (1996) distinguish training programs for the long-term unemployed, re-entry programs for women, and programs addressing specific

skill shortages. The study finds that only the latter two involve positive earnings effects. Hui and Smith (2003a, b) distinguish training programs by the source of financing (private, public, employer). The study concludes that the quality of their data is not sufficient to provide credible estimates. For the UK, Dorsett (2006) analyzes the relative effectiveness of the five options offered under the New Deal for Young People. The five options are remaining on the gateway of intensive job search, subsidized full time employment, full-time education or training, work placement in a voluntary sector, and placement in the so-called Environmental Task Force. In most cases, entering one of the other options was not found to be more effective than intensive job search.

As one of the first studies providing a pairwise evaluation of multiple, mutually exclusive treatments (Imbens (2000) and Lechner (2001)), Gerfin and Lechner (2002) analyze the comparative effectiveness of five types of public sponsored training programs in Switzerland with durations ranging between 5 and 13 weeks using high quality administrative data. The study finds that, one year after program start, the employment rate of participants is lower than that of comparable non-participants. Longer, more intensive training programs show less negative effects than shorter ones.

For the Nordic countries, a number of studies investigate the effects of public sponsored training programs, often based on high quality administrative data. Larsson (2003) finds zero or even negative effects of two youth labor market programs in Sweden, a subsidized work program and a training program, compared to no treatment. Sianesi (2008) analyzes the comparative effectiveness of different labor market programs in Sweden, among them vocational classroom training, work practice (subsidized employment including elements of training), and trainee replacement (training on-the-job mostly in the public sector). She finds moderate positive effects for trainee placement but no positive effects for other forms of training. For Norway, Hardoy (2005) and Jespersen, Munch, and Skipper (2008) find negative effects of classroom training. The evidence on more practically oriented training is mixed.

A number of studies use high quality administrative data for Germany during the 1990s with an evaluation period after program start of up to eight years. Lechner, Miquel, and Wunsch (2007, 2011) distinguish between medium-term programs (mean duration 4 months), longer programs (mean duration 9 to 12 months) and long programs with specific contents. Most of the programs show positive effects in the long run, even in East Germany. An important finding is that medium-term programs outperform longer programs as they exhibit a much shorter lock-in period

with otherwise similar employment effects after the end of the program. Similar findings are obtained by Fitzenberger and Speckesser (2007), Fitzenberger and Völter (2007), and Fitzenberger, Osikominu, and Völter (2008) based on the same data but using a different methodological approach building on Sianesi (2004).

The administrative data for Germany during the 1990s are yet still much less informative than the administrative data used in this study for the 2000s. For instance, the data for the 1990s only contain information on periods with unemployment benefits but not about being registered as unemployed. Furthermore, the older data involve less rich information on personal characteristics and no information about other active labor market programs apart from training. As to be expected from the evidence reported in Sianesi (2008) for Sweden, a country with a comprehensive set of active labor market policies like Germany, information about other programs is important, as our analysis will show. Our data set is not unmatched in terms of informational content compared to data sets used for the evaluation of public sponsored training programs in other countries. Our impression is that there are administrative data sets available for the Scandinavian countries, Switzerland, and Austria, which are more or less just as rich in informational content as the recent administrative data for Germany. In contrast, the data sets used for the US, Canada, UK, and most other countries seem typically much less informative in comparison.

A further drawback of the studies for Germany for the 1990s is that they do not analyze short-term training programs which were abolished in 1993.⁶ Since their reintroduction in 1998, short-term training programs ('Trainingsmaßnahmen') have by now become the most important active labor market program regarding the number of participants (see Section 3.1). Evaluations for other countries (Blundell, Costa Dias, and van Reenen (2004), Weber and Hofer (2004), Fougère, Pradel, and Roger (2005), Crépon, Dejemeppe, and Gurgand (2005), van den Berg and van der Klaauw (2006)) and policy discussions (Martin and Grubb (2001), OECD (2005)) suggest that short-term training programs are superior to longer training programs because the former activate the unemployed without a long lock-in effect. Job search assistance may be an inexpensive way to increase the employment chances of job-seekers. Using data for the 2000s, Hujer, Thomsen, and Zeiss (2006) find that participation in short-term training reduces the unemployment duration of West German job-seekers.

---

⁶An exception is Fitzenberger et al. (2012) who compare the effects of different types of short-term training for the 1980s, the early 1990s, and the early 2000s and which finds modest positive employment effects in some but not all cases.

The studies mentioned so far for Germany do not compare short-term training to other programs. Using the data for the 2000s, Schneider et al. (2006) find that shorter training programs, which are longer than short-term training programs, tend to be more effective than longer ones. Osikominu (2012) uses the same administrative data as this paper but a duration framework in continuous time. She finds that short-term training reduces the remaining time in unemployment and has moderate positive effects on subsequent job stability. Longer training programs initially prolong the remaining time in unemployment but former participants have substantially more stable employment spells and earn more. Also using the same administrative data as we do, Wunsch and Lechner (2008) and Lechner and Wunsch (2009) estimate differential effects of various programs including short-term training.⁷ For West Germany, partly in contrast to our results in this study, Wunsch and Lechner (2008) find negative lock-in effects for all programs in the short run and no – or even negative – medium-run effects on employment and earnings. We take the differences to our study as a motivation to analyze the sensitivity of results with regard to methodological choices. We note here three important methodological aspects of the study Wunsch and Lechner (2008).⁸ First, the comparison group is defined based on nonparticipation during an observation window of 18 months (in an additional sensitivity analysis, a time window of 12 months is used). Second, Wunsch and Lechner (2008) simulate hypothetical program start dates for controls and use these dates for the alignment of treated and controls in time. This presumes matching on the duration until treatment, which is unknown for the controls. Third, in addition to the start dates of the program, Wunsch and Lechner (2008) match on an estimated propensity score that does not change over the course of the unemployment spell.

3 Background and Data

3.1 Training as Part of Active Labor Market Policy

The main goal of German active labor market policy is to permanently reintegrate unemployed individuals (and individuals who are at risk of becoming unemployed)

---

⁷Similar to Lechner and Wunsch (2009), Biewen et al. (2007) do not find significantly positive employment effects for East Germany.

⁸Another difference between the evaluation sample in Wunsch and Lechner (2008) and ours is that they do not include programs financed under the heading ‘Freie Förderung’ and by the ESF (European Social Fund). Further, they restrict the analysis to workers who receive unemployment assistance or unemployment benefits.

back into employment. There exists a wide range of different programs such as wage subsidies, job creation schemes, youth programs, programs to promote self-employment, and training programs, see figure 1 for an overview. Training programs have traditionally been the most important part of active labor market policy. This paper focuses on training programs during the time period 2000 and 2002.⁹

— Figure 1 about here —

There are three main types of training: *short-term training* (Trainingsmaßnahmen), *further training* (Berufliche Weiterbildung), and *retraining* (Umschulung).¹⁰ Apart from the fact that all three types of training require full-time participation, they differ considerably in length and content.

Short-term training programs last only two to twelve weeks (the mean duration is slightly over four weeks, table 1) and typically have one or several of the following three goals. A first goal is to assess job seekers' labor market opportunities and their suitability for different jobs. This may also entail profiling and developing a strategy to find a job. A second goal is to test the job seeker's willingness to work and to improve job search skills. This may involve job-application training. The third goal is the provision of specific skills that are necessary to improve the job seeker's labor market prospects, e.g. through computer courses or commercial training. In 2001, about a fifth of short-term training programs focus on the first and second goal, respectively, and 28% on the third. 31% serve more than one of these goals (Kurtz (2003)).

— Table 1 here —

The more substantial *further training* programs typically last between several months and one year, thus representing medium-term programs. Their goal is to maintain, update, adjust, and extend professional skills. The programs cover a wide range of fields and may also comprise practical elements such as on-the-job training or working in practice firms. Typical examples include training on marketing and sales strategies, computer assisted bookkeeping, operating construction machines, and specialist courses in specific legal fields. Depending on their practical content

---

⁹We focus on the period before the Hartz-Reforms were implemented that changed some of the rules on training programs (Schneider et al. (2006)).

¹⁰In addition, there are specific training programs for disadvantaged youths and disabled persons as well as German language courses.

we distinguish between *classroom further training* and *practical further training*. Finally, *retraining* programs involve training on a new vocational degree according to the German system of vocational education. They last two to three years.

Eligibility for one of the training programs requires the registration as a job-seeker at the local labor office. This involves a counseling interview with the caseworker. For participation in the longer further training and retraining programs, individuals also have to fulfil a minimum work requirement of one year and they must be entitled to unemployment benefits. However, there are a number of exceptions. The binding criterion is that the training program has to be considered necessary in order for the job-seeker to take up a job, e.g. because the training is required to meet the hiring standards of the job. Training programs are usually assigned by the caseworker, depending on the regional supply of upcoming training slots. A participation in training may take place at any point in time during the unemployment spell. Job-seekers have no entitlements regarding participation. A program assignment is compulsory for the job-seeker and noncompliance may entail benefit sanctions and the exclusion from further services. The employment agency covers all direct training costs. In addition, participants in short-term training may continue to receive unemployment benefits or means-tested unemployment assistance, if eligible. Participants in further training and retraining usually receive a subsistence allowance of the same amount as unemployment benefits or unemployment assistance, provided they fulfill the minimum work requirement.

Table 1 shows that the average monthly training costs per participant are lower for short-term training courses (€570 in 2001) than for the medium- and long-term programs (€664). Given that the average length of short-term training is only 1.1 months while that of the longer programs is 9.3 months, this results in ten times higher costs for the medium- and long-term programs as opposed to short-term training (€6,175 vs 627). In light of their substantially lower costs, participation in short-term training programs has been expanded since 2002 compared to longer-term programs (figure 1). Short-term training has become the largest training program regarding the number of participants. These stark differences in costs and durations motivate our comparison of the two types of training programs.

3.2 Administrative Database Used

Our study uses a new and exceptionally rich administrative database, the Integrated Employment Biographies Sample (IEBS) provided by the Research Data Center of

the Federal Employment Agency.¹¹ The IEBS is a 2.2% random sample from a data file combining individual records out of four different administrative registers, i.e. (i) the *Employment History* (Beschäftigten-Historik), (ii) the *Benefit Recipient History* (Leistungsempfänger-Historik), (iii) the *Database of Registered Job-Seekers* (Bewerberangebot), and (iv) the *Database of Program Participants* (Maßnahme-Teilnehmer-Gesamtdatenbank). Start and end dates of the different labor market episodes are measured with daily precision. In addition to individual labor market histories, the IEBS contains extensive information on personal characteristics, occupation and job characteristics as well as regional identifiers, which allows us to merge e.g. the unemployment rate at the county level.¹² For evaluation purposes, a rich set of covariates is essential to reconstruct the circumstances affecting program participation and labor market outcomes.

The *Employment History* involves register data for the time period January 1990 to December 2004 on employment subject to social security contributions.¹³ Earnings before taxes are collected on an annual basis in a given employment relationship if no changes in relevant information, such as a change of the health insurance, occur within the year. Specifically, for somebody continuously working with a given employer we observe a spell that lasts from January 1st to December 31 of each year together with total earnings during each spell. This allows us to calculate earnings per calendar day as well as over more aggregated time intervals. We use these data to construct variables that reflect the employment status and wages during the pre- and post-treatment periods, spanning about ten years in total. We only consider unsubsidized employment spells for our outcome measures. In addition, we use information on year of birth, gender, occupation and industry.

The *Benefit Recipient History* includes spells of all unemployment benefit, unemployment assistance, and subsistence allowance payments between January 1990 and June 2005 together with some personal characteristics. The Benefit Recipient History provides information on periods during which people are not employed and therefore not covered by the Employment History. We use the information on benefit payments to construct individual benefit histories dating back for seven years. Moreover, we use information on reasons for suspensions of payments to proxy a

---

¹¹See Jacobebbinghaus and Seth (2007) and <http://fdz.iab.de/en.aspx>. Being among the first to use the IEBS, we performed extensive consistency checks, see Bender et al. (2004, 2005), Fitzenberger, Osikominu and Völter (2006), and Waller (2008).

¹²Counties do not necessarily coincide with local labor markets. Nevertheless, the county level unemployment rate is a good proxy for the local labor market conditions.

¹³This is equivalent to about 80% of the employed persons (Bender, Haas, and Klose (2000)). The main groups not covered are the self-employed and civil servants.

lack of motivation to find a job.

The *Database of Registered Job-Seekers* contains information on job search during the time period January 1997 to June 2005.¹⁴ First, the data source provides additional information about the labor market status of a person, i.e. whether somebody is looking for a job while employed or nonemployed or whether somebody is temporarily sick while being registered as unemployed. Second, it includes rich information about personal characteristics, e.g. educational attainment, nationality, family situation, health status, characteristics of the last job as well as goals of job search.

The *Database of Program Participants* contains detailed information on participation in all public sponsored labor market programs (5-digit program codes) as well as some additional variables such as the planned end date. This register covers the time period January 2000 to July 2005. We use the information on participation in job creation schemes and wage subsidy programs to identify subsidized employment spells in the Employment History.

3.3 Evaluation Sample and Training Programs Analyzed

Our evaluation approach involves comparisons between a group of people receiving a program of interest and a group of people who do not receive the program but could do so in principle. Correspondingly, one has to determine first who is a potential program participant. We focus on individuals who become unemployed after having been continuously employed for a while, instead of a stock sample of *observed* unemployed at a given point of time. This way, we exclude those individuals who have been out of the labor force and who register as unemployed just because they want to participate in a training program. In interviews, caseworkers told us that especially women returning from maternity leave, divorcees, or university graduates who have difficulties finding a job contact the local employment office inquiring about the possibility to participate in programs. Thus, an evaluation sample based on the observed unemployment status at a particular point in time (instead of an inflow sample from employment into unemployment) has important disadvantages. Since participants returning from out of the labor force are not registered as unemployed before the start of training, the corresponding comparison persons, who are still out of the labor force, are not recorded in the data. Even if it

---

¹⁴According to the data manual for the IEBS, the Database of Registered Job-Seekers is complete only from 2000 onwards. However, this does not seem to be an issue for our analysis as the share of missings in our retrospective variables does not vary systematically over time.

were possible to observe these people in the data, this would entail the danger that these individuals are not considering to take up a job, which is the main purpose of active labor market programs. Therefore, we focus on individuals who are attached to the labor market because they were employed in the recent past. This enables us to construct a suitable comparison group based on the information in the data. Furthermore, the beginning of the unemployment spell defines a natural time scale to align treated and nontreated individuals.

Our sample of inflows into unemployment comprises West Germans who become unemployed between the beginning of February 2000 and the end of January 2002, after having been continuously employed for at least three months. Entering unemployment is defined as ending regular, non-subsidized full-time or part-time employment, and subsequently being in contact with the employment office (not necessarily immediately) as reflected by a job search spell, benefit receipt, or program participation.¹⁵ In order to exclude individuals eligible for specific youth programs or for early retirement schemes, we only consider persons aged between 25 and 53 years at the beginning of their unemployment spell. Our evaluation focuses on the first training program that is attended in the course of a given unemployment spell.

— Table 2 here —

We report results for two different types of training that largely follow the legal grouping of program types: *short-term training (STT)* and *classroom further training (CFT)*, see Section 3.1.¹⁶ In a few cases, we also included programs with similar contents and planned durations but somewhat different legal basis into our program types.¹⁷ While we display treatment effects only for STT and CFT, we consider participation in other programs, e.g. practical further training and subsidized employment, when defining the comparison groups (see table 2). That is, participants in other programs are not counted as in open unemployment. We estimate program effects separately for men and women and for three strata representing different durations of elapsed unemployment. This results in a total of six evaluation samples (stratum 1/2/3, male/female). Table 2 gives an overview of the sample sizes and the

---

¹⁵About 10% of the individuals in our inflow sample appear with more than one unemployment spell according to our inflow definition. We take account of this when calculating standard errors, see Section 4.

¹⁶In Biewen et al. (2007), we also report results for *practical further training* and *retraining* as well as for East Germany.

¹⁷We assigned some programs in the legal category of *Discretionary Support* (Freie Förderung) and programs financed by the *European Social Fund* (Europäischer Sozialfonds) into one of our program categories.

transitions from open unemployment. Figure 2 provides descriptive information on the duration of different program types. STT programs last on average about one month. CFT programs last on average 7.5 months, with spikes in the distribution at 6 and 12 months.

— Figure 2 here —

4 Econometric Approach

4.1 Evaluating Multiple Treatments in a Dynamic Context

Our empirical analysis is based upon the potential-outcome approach to causality (see Neyman (1923), Roy (1951), Rubin (1974) and the survey of Heckman et al. (1999)). Imbens (2000) and Lechner (2001) extend this framework to allow for multiple, exclusive treatments. Our analysis distinguishes two different treatments, STT and CFT, as well as a nontreatment state. Let Y^k represent the potential outcome associated with training program $k = 1, 2$ and Y^0 the potential outcome when not participating in any program. For each individual, only one of the three potential outcomes is observed. We focus on the first treatment in a given unemployment spell. Our goal is to estimate the average treatment effect on the treated (ATT) of receiving treatment $k = 1, 2$ as the first treatment both against nonparticipation $l = 0$ (treatment versus ‘searching’) and against treatment l , with $k \neq l$ and $k, l \neq 0$ (differential effect of two programs).

In our context, participation in training is possible at any point in time during the unemployment spell. A job-seeker who has not yet enrolled at a given point in time may join at some later point in time provided he remains unemployed up to that time. In such a setting, persons with longer completed unemployment durations are more likely to end up receiving treatment than persons with shorter completed unemployment durations. Fredriksson and Johansson (2008) show that in this case a static evaluation analysis that assumes that the treatment is administered only once yields biased treatment effects because the definition of the comparison group conditions on future outcomes.¹⁸ We follow the approach suggested in Sianesi (2004, 2008) to extend the static multiple treatment approach to a dynamic setting. Sianesi focuses on the effect of participating after a given unemployment experience (i.e. pe-

---

¹⁸While Fredriksson and Johansson (2008) then estimate the effect on the hazard of leaving unemployment, our outcome variables of interest are the monthly employment status and earnings.

riod of eligibility) versus not participating up to that point in time and continuing to search for a job. Thus, treatment effects are defined conditional on a given starting date during the unemployment spell and the treated and comparison individuals are aligned by the elapsed unemployment duration. Note that individuals in the nontreatment group may enroll in training at some later point during the unemployment spell. This treatment parameter (treatment versus searching) mirrors the decision problem of the caseworker and the unemployed who recurrently during the unemployment spell decide whether to start any of the training programs now or to postpone possible participation to the future, in case job search will not be successful. When discussing the policy implications of our estimates, one has to be careful because those nontreated individuals in one stratum who receive treatment later experience the impact of the treatment.

We distinguish between treatment starting during months 1 to 3 of the unemployment spell (stratum 1), treatment starting during months 4 to 6 (stratum 2), and treatment starting during months 7 to 12 (stratum 3). We experimented with shorter strata but found that the number of treated and comparison individuals became too small (thus resulting in higher variance of the estimates) to implement our matching approach that relies on both exact matching and kernel matching (see below). Also due to concerns regarding small sample sizes, we do not evaluate training programs starting more than 12 months after the beginning of unemployment.

We evaluate treatments conditional on the unemployment spell lasting at least until the start of treatment k . Let $u = 1, 2, \dots$ denote the month in which treatment k starts in the unemployment spell. For instance $u = 2$ means that treatment starts in month 2 of unemployment. $\tau = 0, 1, 2, \dots$, counts the months since the beginning of treatment k . $Y^k(u, \tau)$ is then the potential outcome in period $u + \tau$ for treatment k starting in month u . Similarly, $Y^l(\tilde{u}, \tau - (\tilde{u} - u))$ is the alternative potential outcome in period $\tilde{u} + (\tau - (\tilde{u} - u))$ associated with the alternative treatment l starting in period $\tilde{u}$. The ATT parameter for treatment k against the alternative l ($k \in \{1, 2\}$, $l \in \{0, 1, 2\}$, $k \neq l$) is given by

$$(1) \quad \theta(k, l; u, \tau) = E(Y^k(u, \tau) | U^o = u - 1, T_u = k) - E_{u \leq \tilde{u} \leq \bar{u}_s}(Y^l(\tilde{u}, \tau - (\tilde{u} - u)) | U^o = \tilde{u} \geq u - 1, T_u = k),$$

where U^o denotes the random time spent in open (i.e. untreated) unemployment and $T_u \in \{0, 1, 2\}$ the treatment status in period u . $\underline{u}_s$ ($\bar{u}_s$) denotes the start (end) of stratum s , $s \in \{1, 2, 3\}$, and $E_{u \leq \tilde{u} \leq \bar{u}_s}(\cdot)$ the expectation with respect to the distribution of $\tilde{u}$ in the remainder of the stratum. Thus, for treatment k starting in period u we require that possible comparison individuals receiving treatment l have spent the same amount of time in open unemployment as of period $u - 1$ and receive

treatment l no earlier during the stratum considered than the treated individual. For $l = 0$, we have $\tilde{u} = u$ and the comparison group comprises all individuals whose unemployment spell lasts at least $u - 1$ periods and who do not participate in any program during the stratum considered. This entails the possibility of exit to employment before the end of the stratum. For $l \geq 1$, $\tilde{u}$ can take any value in the interval $[u, \bar{u}_s]$ and $\tau - (\tilde{u} - u)$ counts the months since start of treatment l , thus aligning elapsed unemployment duration. This way, we compare participation in one program to possibly delayed participation in another program later during the stratum.

We display treatment effects by stratum only, averaging the period specific treatment effects, $\theta(k, l; u, \tau)$, with respect to the distribution of starting dates in the stratum:

$$\theta(k, l; s, \tau) = \sum_u f_s(u) \theta(k, l; u, \tau) ,$$

where $f_s(u)$, $u \in \{\underline{u}_s, \dots, \bar{u}_s\}$, is the distribution of starting dates in stratum s .

We assume the following dynamic version of the conditional mean independence assumption (DCIA):

$$(2) \quad \begin{aligned} & \mathbb{E}_{u \leq \tilde{u} \leq \bar{u}_s} (Y^l(\tilde{u}, \tau - (\tilde{u} - u)) | U^o = \tilde{u} \geq u - 1, T_u = k, \mathbf{X}_s) = \\ & \mathbb{E}_{u \leq \tilde{u} \leq \bar{u}_s} (Y^l(\tilde{u}, \tau - (\tilde{u} - u)) | U^o = \tilde{u} \geq u - 1, T_{\tilde{u}} = l, \mathbf{X}_s) , \end{aligned}$$

where $\mathbf{X}_s$ denotes a vector of covariates that are time-invariant within a stratum. We effectively assume that, conditional on $\mathbf{X}_s$ and on remaining in open unemployment at least until period $u - 1$, individuals are comparable in their outcome for treatment l occurring in month $\tilde{u}$ between u and $\bar{u}_s$. As usual, implicit is the assumption that potential outcomes are independent across individuals, ruling out general equilibrium effects.

4.2 Combining Exact Matching and Kernel Matching

Building on Rosenbaum and Rubin's (1983) result on the balancing property of the propensity score in the case of a binary treatment, Lechner (2001) shows that the conditional probability of treatment k , given that the individual receives treatment k or treatment l , $P^{k|kl}(\mathbf{X}_s)$, exhibits an analogous balancing property for the comparison of program k versus l . This allows us to apply standard binary propensity score matching based on the sample of individuals participating in either program k or l . For this subsample, we simply estimate the probability of treatment k and then apply a bivariate extension of standard propensity score matching techniques.

Implicitly, we assume that matching on the elapsed unemployment duration, prior employment history and calendar time at the start of the spell accounts for any remaining selection within the stratum.

To account for the dynamic nature of treatment assignment, we estimate the probability of receiving treatment k versus l in a given stratum using all the individuals in group k and l who are still unemployed at the beginning of the stratum considered. For treatment during months 1 to 3, we take the total sample of unemployed, who participate in k or l during months 1 to 3 (stratum 1), and estimate a Probit model for participation in k . For $l = 0$, the comparison group includes those unemployed who either never participate in any program or who start some treatment after month 3 (see table 2). Similarly, for treatment during strata 2 and 3, the basic sample consists of those individuals in group k and l who are still unemployed at the beginning of the stratum.

We then combine kernel matching involving the propensity score with cell matching, by imposing additional criteria along which treated and comparison observations are aligned exactly. Specifically, we impose the following four matching requirements: (i) similarity in the pairwise propensity score, (ii) similarity in the calendar date of the beginning of unemployment, (iii) equality in the elapsed time in open unemployment, and (iv) equality in the employment history before the start of unemployment. Given exact matches for (iii) and (iv), we estimate the counterfactual employment and earnings outcomes by means of a Nadaraya-Watson kernel regression to achieve (i) and (ii). We use a product kernel in the estimated propensity score and the calendar month of entry into unemployment

$$(3) \quad KK_j(p_d, c_d) = K\left(\frac{p_d - p_j}{h_p}\right) \cdot h_c^{|c_d - c_j|},$$

where $K(z)$ is the Gaussian kernel function, p_d and c_d are the propensity score and the calendar month of entry into unemployment of a particular treated individual d , p_j and c_j are the estimated propensity score and the calendar month of entry into unemployment of an individual j belonging to the comparison group of individuals treated with l . The parameters h_p and h_c are the bandwidth (type) parameters. We use the bivariate crossvalidation procedure suggested in Bergemann, Fitzenberger, and Speckesser (2009) and Fitzenberger et al. (2008) to obtain the bandwidths h_p and h_c by minimizing the squared prediction error of the average l -outcome for the individuals in the l -group who are most similar to the participants in program k .¹⁹ We obtain standard errors and pointwise confidence bands for our estimated

---

¹⁹For the comparisons of training versus searching, the crossvalidation yields an optimum for h_p of zero, which corresponds to an exact alignment in calendar time.

treatment parameters through bootstrapping based on 250 resamples. We resample individuals with their entire observation vector to obtain standard errors clustered at the individual level, thus also accounting for the multiple appearance of individuals. The propensity score is reestimated in each resample.

Recently, there have been a number of Monte Carlo studies analyzing the performance of various estimators for treatment effects relying on propensity score matching, see Busso, DiNardo, and McCrary (2009, 2011), Galdo, Smith, and Black (2008), and Huber, et al. (2010). These studies reassess the conclusions drawn by the earlier Monte Carlo study of Frölich (2004). The key insights from the different studies are the following. First, there is no estimator which performs best in all settings and according to all criteria. Second, the popular nearest neighbor matching approach is inferior to more involved estimators appropriately adjusted to the estimation problem investigated. Third, optimized versions of both kernel or radius matching improve upon simple implementations, unless sample sizes are really small, and typically do not perform badly even when inverse probability weighting (IPW, the preferred method in Busso et al. (2009, 2011)) performs better. Fourth, regression adjustment in addition to the preferred estimator by the authors improves the performance when the propensity score is incorrectly specified.²⁰ The aforementioned Monte Carlo studies determine the bandwidth for kernel matching approaches based on a standard leave-one-out cross-validation. The Monte Carlo results reported in Galdo et al. (2008) show that it can be useful to base the bandwidth choice on predictions for nontreated individuals who are close to the treated individuals.

Our estimation approach is not covered explicitly by the recent Monte Carlo studies reported in Busso et al. (2009, 2011) and Huber et al. (2010). Both local linear ridge matching (the best performing estimator in Frölich (2004) and Huber et al. (2010)) and IPW are not feasible in our application because we combine cell matching and kernel matching and match on closeness in two dimensions at the same time. In the case of treatment versus searching we have sizeable treatment groups (between 200 and 900 observations) and 17 to 75 times larger nontreatment groups (see table 2) – and we solely estimate the average treatment effect on the treated. Furthermore, the overlap in the covariate distributions and the balancing of the covariates are both very good in these cases (see Section 5.3). We expect our kernel regressions to perform very well in the large nontreatment samples, in particular because we take great care in details such as bandwidth choice and regression adjustments after matching (see Section 4.6). We follow Huber et al. (2010) in that we match exactly

---

²⁰Note that the studies do not analyze the effect of regression adjustment for all estimation approaches.

on employment history variables – and in that we investigate the importance of trimming highly influential comparison observations (see Section 5.3).

4.3 Empirical Content of the DCIA

The DCIA in equation (2) states that conditional on a given unemployment experience and a vector of observed covariates, the sequence of potential outcomes associated with alternative treatment l (including nonparticipation $l = 0$) in the same stratum is mean independent of the treatment status in this stratum. In a dynamic context, nonparticipation in the current stratum entails the possibility of participation in later strata. Our matching approach will produce valid estimates if we consider all the determinants that jointly influence treatment status and potential outcomes in the current stratum as well as in future strata. Conditional on these determinants individuals are randomly allocated to one of the treatments in a given stratum and there is balanced anticipation of future treatment or employment chances.²¹ We argue in the following that these assumptions are plausible in light of program assignment in Germany and the rich information in our data.²²

Although the assignment to a training program typically occurs with the consent of the job-seeker, considering his or her willingness to receive training and to work in a specific field, the assignment decision is up to the discretion of the caseworker during the time period we consider. According to Blaschke and Plath (2000), indicators available to the caseworker like the composition of a group of participants in a particular course or his/her assessment of the motivation of the unemployed played an important role. Job placement has priority over program participation. The unemployed are encouraged to continue job search at any time, even while participating in a training program. The assignment to a particular training program is driven strongly by the supply of courses (Schneider et al. (2006)). Evidence from qualitative surveys in Schneider et al. (2006) suggests that belated assignments and referrals on very short notice are commonly used in order to assure a high capacity utilization of courses booked in advance and to keep up job search incentives. This suggests that program assignment is not targeted but driven by region-specific variation in the supply of courses throughout the year. From the perspective of the unemployed, the assignment to a particular program cannot be perfectly anticipated. Moreover,

---

²¹‘Balanced’ means that treated and nontreated have the same predictions conditional on the observables considered.

²²See Section A in the Online-Appendix for a list of the covariates considered and descriptive statistics. See Osikominu (2012) and her Supplementary Appendix for further institutional information.

the Database of Registered Job-Seekers involves a wide range of information that is collected by the caseworker as a basis for his counseling activities and assignment decisions.

To be specific, we consider the following variables that reflect the caseworker’s information on the motivation, plans and labor market prospects of a particular unemployed: the caseworker’s assessment of the job-seeker’s current health status, information on his/her previous health status (during the last two years before the start of the current unemployment spell), a dummy variable indicating whether the unemployed person appeared to lack motivation (e.g. failed to show up at regular meetings), dummies indicating whether the job-seeker dropped out of a program, whether benefits were withdrawn, and whether the person participated in a program providing psychosocial support, where all variables refer to the last three years unless said otherwise. In addition, we include variables indicating whether the job-seeker would like to work in a different occupation, whether he/she is looking for a part-time job, and the number of job proposals he/she received from the employment office.²³

The literature (e.g. Heckman, Ichimura, and Todd (1997), Heckman et al. (1998), Heckman et al. (1999), Heckman and Smith (1999)) stresses the importance of conditioning flexibly on lagged employment and wages, benefit receipt history, and local labor market conditions. Our sample of analysis involves individuals who, between 2000 and 2002, exit regular employment lasting three months or longer. Our data allow us to include in the propensity score a rich set of variables depicting past employment and nonemployment histories as well as local labor market conditions. To account for potential anticipatory effects regarding re-employment, we use the complete longitudinal information in our administrative data sources covering employment and benefit receipt and impose several restrictions to ensure that our sample as a whole as well as the different treatment and comparison groups are sufficiently homogeneous. We use several different variables recording wages prior to unemployment as well as occupation and industry of the last job, indicators of whether the last job was part-time, whether it was a white-collar or blue-collar position, the reason why the job ended, and the quarter of the beginning of the unemployment spell. We also include indicators of whether the person was employed

---

²³Note that these variables reflect information that is collected in the course of the counseling process using standardized, closed-ended answer formats. For the three variables ‘occchange’, ‘endlastjob’, and ‘family’ (see Section A in the Online-Appendix), the share of missings is relatively high. Nevertheless, they often prove significant in the propensity scores. This likely reflects the fact that the missings are informative as well. For instance, job-seekers who are undetermined about their future occupation might choose not to answer the question.

in month 6, 12, or 24 before the beginning of the unemployment spell, the number of days employed during the preceding seven years as well as the average annual earnings and the annual employment status in each of the seven years before the beginning of the unemployment spell.

As in Card and Sullivan (1988), we stratify treatment and comparison individuals based on their employment history sequences as well as elapsed unemployment duration (see also Dolton and Smith (2011)). We construct ten different employment history groups that indicate different combinations of the annual employment status of an individual during the four years preceding the unemployment spell.²⁴ We use the nine sequences (1000), (1001), (1010), (1011), (1100), (1101), (1110), (1111), and (0000). E.g. the sequence (1010) represents employment during the first and the third year before entry into unemployment and nonemployment during the second and the fourth year. The nine sequences considered so far cover over 95% of the cases. The remaining sequences are subsumed in a residual category. We include dummy variables for each sequence in the propensity score.

Turning to nonemployment, we include the entitlement period for unemployment benefits at the beginning of unemployment as well as variables measuring the duration of unemployment transfer payments (i.e. unemployment benefits, unemployment assistance, or subsistence allowance) in days during the last three years, the number of days without any information in the data set (e.g. periods of self-employment or out of the labor force), indicators for whether somebody was disabled or unable to work at some point in time during the preceding three years, and a variable recording previous contacts with the employment office during the last three years. Finally, to model local labor market conditions we use several different variables on the unemployment rate in the county of residence, indicators for the federal state and indicators for the economic situation of the local labor market.

4.4 Specification of the Propensity Scores

Using all of the variables described in the previous section as possible regressors, we fit the propensity scores separately for each of the six evaluation subsamples (men/women, stratum 1/2/3), and each treatment comparison pair. In each case, we run an extensive specification search. The final specification was chosen based on which variables (according to institutional and economic knowledge) may drive the

---

²⁴A person is considered as employed in a given year if he/she is employed for at least 50% of the days. The results are not sensitive to the exact choice of the cutoff point.

selection into programs, based on statistical significance of the variables included, and based on the balancing tests described below. We usually start with a fairly general specification and drop variables that are grossly insignificant. This leads to satisfactory specifications in most cases. In the few cases in which the balancing condition fails, we modify the specification and reinclude covariates or add interactions in order to eliminate the covariate imbalance. We take particular care in ensuring that pre-treatment employment and earnings outcomes are well balanced between treatment and comparison group members. The final specifications typically include between 20 and 35 covariates (see Section B.1 in the Online-Appendix for the estimation results).

4.5 Testing for Covariate Balance

We employ two different balancing tests to check whether our treatment and comparison groups are sufficiently balanced.²⁵ First, we test whether the means of important covariates differ significantly between treatment and matched comparison groups using bootstrap t -tests and Wald tests (for groups of variables that belong together). Our bootstrap approach accounts for clustering at the individual level which is of particular importance as members of the comparison group can be used as comparisons more than once. In each subsample, the comparison of means is based on the same set of 90 regressors that we consider particularly important (cf. the discussion in Section 4.3). We calculate the matched mean of a given regressor by applying the matching procedure described above to this regressor in exactly the same way as we do to predict the counterfactual employment and earnings outcomes. Second, we investigate whether treated and matched nontreated individuals differ significantly in their employment and earnings outcomes over a period of seven years before the beginning of unemployment. We estimate these differences in the same way as the treatment effects after the beginning of the program. By construction, treated individuals and their matched counterparts exhibit the same unemployment duration until the beginning of treatment as well as the same type of four-year employment history before the beginning of the unemployment spell. We report the results of both balancing tests in Section 5.3 and in Section B.2 in the Online-Appendix

---

²⁵See Lee (2012) for a recent critical assessment of balancing tests.

4.6 Regression Adjustment and Test for Effect Heterogeneity

The estimated ATT may be biased if treated and matched comparison individuals are not matched well by observed characteristics affecting the treatment effect.²⁶ Following Abadie and Imbens (2011) and Mueser et al. (2007), we apply an regression based adjustment to investigate the existence of any remaining mismatch between treated and comparison individuals after kernel matching. While Abadie and Imbens (2011) propose the ex post regression adjustment as a device to remove the leading bias term of the nearest neighbor matching estimator (Abadie and Imbens (2006)), in our case of kernel matching, the bias of the nonparametric kernel regression vanishes asymptotically. Nevertheless, as some mismatch in the levels of the matching variables between each treated unit and the comparison units may remain in any finite sample (Mueser et al. (2007, p. 765)) we use the ex post regression adjustment to investigate the existence of such mismatch bias.²⁷ At the same time, we use it to test for treatment effect heterogeneity.²⁸ In fact, estimation of the ATT provides a semi-parametric, aggregate impact measure of potentially heterogeneous individual treatment effects depending upon observed characteristics of the participating individuals.

In our outcome regressions, the dependent variable, $\bar{Y}_i(u)$, is the average of the individual period specific outcomes over 24 months from month u onwards. We focus on this global outcome in order to abstract from the volatility over time of the monthly outcomes. Pooling over all possible starting dates u in stratum s , i.e. $u \in [\underline{u}_s, \bar{u}_s]$, $s = 1, 2, 3$, we run the following weighted regression:

$$(4) \quad \min_{\hat{\alpha}, \hat{\beta}, \hat{\gamma}, \hat{\delta}} \sum_i \sum_{u=\underline{u}_s}^{\bar{u}_s} g_i(u) [\bar{Y}_i(u) - \hat{\alpha} - \mathbf{X}_i \hat{\beta} - \hat{\gamma} D_{s,i} - D_{s,i} (\mathbf{X}_i - \bar{\mathbf{X}}) \hat{\delta}]^2,$$

where the summation is over all individuals i who either receive treatment k or l

---

²⁶A possible mismatch may be caused by a misspecification of the propensity score or by the fact, that despite the validity of the DCIA, the particular sample draws for the observed characteristics may differ between treated and matched comparison individuals. The former involves a bias, that remains in a large sample if a parametric propensity score is used, while the latter involves a finite sample bias.

²⁷Mueser et al. (2007) adjust each set of comparison cases for a particular treated case using a linear regression in the weighted data set of comparison cases (see also Heinrich et al. (2010) for this approach). Unlike us they do not use the ex post regression to study effect heterogeneity.

²⁸To our knowledge, recent evaluation studies rarely estimate ex post outcome regressions to study effect heterogeneity. In the literature on statistical treatment choice, Frölich (2008) uses a similar approach to predict potential outcomes as a function of individual characteristics. Huber et al. (2010) report Monte Carlo evidence for treatment effect estimators based on such ex post outcome regressions, but do not investigate effect heterogeneity.

in the stratum considered. $\mathbf{X}_i$ denotes a vector of observed characteristics and $\bar{\mathbf{X}}$ the average of the characteristics in the sample of individuals receiving treatment k .²⁹ $D_{s,i}$ is a dummy variable for receiving treatment k in stratum s , i.e. $D_{s,i} = 1$ ($D_{s,i} = 0$) if individual i receives treatment k (l).

$g_i(u)$ is the weight associated with treatment start u . For a treated individual ($D_{s,i} = 1$), the weight picks the observed start month, i.e. $g_i(u) = 1$ for persons starting treatment k in month u and zero otherwise. For a nontreated individual j ($D_{s,j} = 0$) serving as a comparison for treatment k starting in month u the weight $g_j(u)$ equals

$$\sum_{d \in \{T_u = k \wedge U^o = u-1; u \in [\underline{u}_s, \bar{u}_s]\}} \frac{K\left(\frac{p_d - p_j}{h_p}\right) \cdot h_c^{|c_d - c_j|}}{\sum_{n \in \{T_{\bar{u}} = l \wedge U^o \geq u-1; u, \bar{u} \in [\underline{u}_s, \bar{u}_s]\}} K\left(\frac{p_d - p_n}{h_p}\right) \cdot h_c^{|c_d - c_n|}},$$

where d indexes those individuals starting treatment k in month u ($T_u = k \wedge U^o = u-1$) and n those comparison units who are still in open unemployment in the month before treatment k starts ($U^o \geq u-1$). The function inside the sum corresponds to the normalized kernel weight as given in equation (3), where the normalization guarantees that the total sum of weights for d across nontreated individuals j sums to one. Note, that j is used repeatedly for each possible treatment starting date u in the respective stratum for which j serves as a comparison unit, i.e. observation j is used up to $(\bar{u}_s - \underline{u}_s + 1)$ times in the pooled regression (4).

The term $\mathbf{X}_i \hat{\boldsymbol{\beta}}$ in equation (4) controls for the effect of the characteristics on the average outcome variable. The coefficient on the treatment dummy, $\hat{\gamma}$, gives the mismatch corrected estimate of the aggregate ATT. This mismatch corrected estimate can be compared to the estimate obtained without the additional regression adjustment. If the two estimates are similar, then our matching procedure works well in aligning the distributions of the conditioning variables in the comparison group to that of the treatment group. The coefficients on the vector of interaction terms, $\hat{\boldsymbol{\delta}}$, give the ceteris paribus changes in the treatment effect when changing the corresponding characteristics. If $\boldsymbol{\delta} = 0$, then there is no effect heterogeneity by the level of covariates. The standard errors of the estimated coefficients in (4) are obtained through a full bootstrap for the matching estimator by rerunning the regression for all resamples.

---

²⁹All the explanatory variables in $\mathbf{X}_i$ are dummies.

5 Which Program Works and for Whom?

5.1 Effectiveness of Training Compared with Searching

Figures 3 to 6 show the evaluation results for the comparison of training vs not participating in any active labor market program (i.e. ‘searching’) in a given time window. We distinguish three different time windows (strata) of elapsed unemployment experience: 1 to 3 months (stratum 1), 4 to 6 months (stratum 2), and 7 to 12 months (stratum 3). Each graph displays, for different points in time given on the horizontal axis, the average treatment effect on the treated (ATT), i.e. the difference between the observed outcome with training and the estimated counterfactual outcome without training averaged over those who participate in the program in a given time window. On the time axis in our graphs, positive values denote months since program start, while negative values represent pre-unemployment months and years. We omit the period between the start of unemployment and the start of the program in which both comparison and treatment persons are unemployed. The dashed lines around the estimated ATT are bootstrapped 95% pointwise confidence bands. Treatment effects for a particular time period are statistically significant if zero is not contained in the confidence band. In each figure, the results for men (women) are in the left (right) column. While figures 3 and 5 refer to the difference in employment rates, figures 4 and 6 represent differences in pre-tax earnings. For both outcome measures, subsidized employment is counted as nonemployment.³⁰

— Figure 3 and Figure 4 about here —

Turning to the estimated employment effects of short-term training (STT) in figure 3 we find negative effects shortly after the program start in the order of minus five percentage points (ppoints). This suggests that, while enrolled in the program, participants have a five ppoints lower monthly employment rate than they would have if they did not participate in the program. These lock-in effects do not last for more than two to three months, which is in line with the short duration of STT programs of about one month. After the lock-in period, the treatment effects turn positive in general, but they are not always statistically significant. The results differ strongly

---

³⁰Our earnings information comes from social security records which is collected only once in a given calendar year if a person is employed with the same employer throughout and no changes in relevant information, such as a change of the health insurance, occur within the year. Observed changes in monthly earnings are therefore due to wage changes with the same employer, job-to-job changes, or job-to-nonemployment transitions.

across strata. While there is no evidence for statistically significant treatment effects for individuals participating in the first three months of their unemployment spell (stratum 1), treatment effects for men starting an STT program in months 7 to 12 of their unemployment spell (stratum 3), and women starting one in month 4 or later are significantly positive. Accordingly, the monthly employment rate of men (women) participating in STT is increased by about 5 ppoints (about 11 ppoints) after the end of the lock-in period (see also table 3). The impact estimates for earnings in figure 4 show a corresponding pattern. There are statistically significant earnings gains of up to €100-200 for female participants in STT in strata 2 and 3, and marginally significant effects of up to €100 for male participants in stratum 3 and female participants in stratum 1. Averaging over the second year after program start, the monthly effect is €87 for men in stratum 3 and €169 for women in strata 2 and 3 (see table 4).

— Figure 5 and Figure 6 about here —

Estimates of the employment effects for the more substantive classroom further training programs (CFT) are given in figure 5. The most conspicuous difference between these results and those for STT programs is the long and pronounced lock-in effect. During the first six months after program start, participants have an employment rate that is up to 25 ppoints lower than it would have been if they had not taken part in the program. From month 7 onwards, the treatment effects recover, eventually turning positive 10 to 12 months after program start. The lock-in period is longer and entails larger employment losses for people who take up their treatment during the first 6 months of unemployment (strata 1 and 2). In contrast, people who have been unemployed for more than 6 months (stratum 3) experience less deep and shorter lock-in effects. While there is little evidence for statistically significant employment effects for men starting a CFT program in months 1 to 6 of their unemployment spell (strata 1 and 2) or women enrolling in the first 3 months of unemployment (stratum 1), treatment effects for longer-term unemployed men (stratum 3) and medium to longer-term unemployed women (strata 2 and 3) are sizeable and statistically significant. After the initial lock-in phase, they eventually exceed 10 ppoints, with an average of 7 ppoints throughout the second year since program start (see also table 3).

Figure 6 reveals a similar pattern for the earnings impacts of CFT compared with searching. After an initial lock-in effect, there are significant earnings gains for male participants in CFT in stratum 3 (and marginally in stratum 2), and in all

three strata for women. Especially for men, the earnings gains after lock-in tend to be larger than those for STT, often attaining €200 per month. This means that although, after lock-in, CFT and STT yield comparable employment gains, CFT participants benefit on average more from their program than participants in STT, which is consistent with the fact that CFT involves a more sizeable human capital investment compared to STT. Hence, CFT programs, unlike STT programs, contribute to the development of human capital that is rewarded on the labor market.

Tables 3 and 4 summarize the estimated treatment effects during the first two years after program start. We report the difference of the estimates regarding the average monthly employment and earnings gains or losses over the course of the first and the second year. For STT vs Searching (rows 1 and 5 of tables 3 and 4), in cases with slightly negative effects during the first year these are generally offset by positive effects in the second year. Thus, the average effect tends to be zero or positive for STT after two years. In contrast, it remains negative in the majority of cases for CFT vs searching after two years (rows 3 and 7 of tables 3 and 4). Consistent with the graphical results, when we compare the estimates across strata the estimated effect increases (becomes less negative) the later a STT or CFT program starts during unemployment.

What explains the differing treatment effects across strata? First, a large number of those having become unemployed recently find new jobs easily without participating in a training program. If these people are assigned to a CFT program anyway, they will be ‘locked-in’, while many of their counterparts in the comparison group quickly exit to employment. The evidence in Osikominu (2012) suggests that the opportunity cost of participating in a long program is in fact higher for people with a high probability of exiting unemployment without treatment. Our findings in Section 6 are also consistent with this view because early program participants tend to be a positive selection of all unemployed and are therefore matched with nonparticipants who exit unemployment very quickly. Second, treatment effects tend to be higher for the late program starts if many of the long-term unemployed in the comparison group abandon their job search and move out of the labor force. Our additional evaluation results for the outcome variable ‘leaving the labor force’ (see Section C.1 in the Online-Appendix) confirm that participants enrolling after a longer unemployment experience have a reduced probability to drop out of the labor force compared to the situation if they did not receive training. Hence, an important channel through which training programs work consists in keeping the long-term unemployed in the labor force.

5.2 Heterogeneity of Training Impacts

Given that, in many cases, training programs show positive treatment effects in the medium-run when compared to attending no program in the same stratum, the questions arise (i) who benefits most from a given program and (ii) which program is the best for a given subpopulation. The results in the previous section suggest that STT may have similar or even more positive effects compared to CFT when each of the programs is compared to attending no program. However, this does not necessarily mean that participants in STT could not have improved their employment chances by attending CFT instead, or that participants in CFT would not have lost from taking part in STT instead. This is because the selection of the treated may differ by treatment.

Tables 3 and 4 also contain estimates for the direct comparisons between STT and CFT, which allows us to address question (2) (see also Section C.2 in the Online-Appendix for further evidence). Row 2 for men and row 6 for women in tables 3 and 4 show the average effect of participating in STT instead of CFT for the STT participants. The estimates are generally positive suggesting that those participating in STT could not have fared better with the more intensive CFT program. For instance, the numbers in row 2, column 1-2 of table 3 suggest that male participants in STT in stratum 1 have on average a 16 ppoint higher employment rate in year 1 and a 2 ppoint higher employment rate in year 2 compared to the situation in which they would participate in CFT. The advantage of STT over CFT for those who participate in STT is generally higher in year 1 than in year 2, which reflects the differing enrollment lengths of the two programs. When we look at the opposite comparison of CFT vs STT for those who participate in CFT (rows 4 and 8 in tables 3 and 4) similar patterns emerge during the first year after program start. During the second year after program start, when most CFT participants have completed their training, the estimated effect changes in favor of CFT. Overall however, CFT participants could in most cases have done similarly well or even better with a STT program. This is remarkable since it means that these people could have been assigned to the much less expensive STT programs without deteriorating their labor market chances.

Now, we turn to the question of who benefits most from a given program. The results presented so far reveal some heterogeneity in the effects for different subpopulations.

A consistent finding is that programs are only effective for those individuals who start a program after having been unemployed for some time. However, note that this is subject to the caveat that the selection of eligible individuals changes over the course of the unemployment spell. Another finding is that training effects are generally larger for women than for men (figures 3 and 5). In the following, we further investigate whether the treatment effect depends upon further observable personal characteristics. For this purpose, we regress the individual-level outcomes averaged over the first 24 months since program start on the treatment dummy, a set of observed covariates and the interactions with the treatment dummy. The comparison observations are weighted according to their matching weights, see Section 4.6. We estimate a first regression including all the personal characteristics and interactions with the treatment dummy and three further regressions where we only include interaction effects for one out of three subsets of characteristics (i.e. for demographic variables, human capital variables, and characteristics of the previous job). By adding the interaction effects separately, the power of the tests is likely to be higher.

— Table 5 here —

The results of the Wald-tests for effect heterogeneity for all interaction variables are reported in table 5 and for the separate groups in Sections D.2-D.3 in the Online-Appendix. The full coefficient estimates are reported in Section D.1 in the Online-Appendix. In almost all cases, the Wald-test statistics for the null hypothesis that the interactions between the treatment dummy and the characteristics have zero coefficients are nowhere close to significance, neither when the specification includes interactions with all individual characteristics nor when we undertook separate analyses by type of variables. The lowest p -values are around 10% in two cases and, they are much higher in the remaining 34 cases. A small number of single coefficients are significant but no clear patterns emerge. The overall insignificance suggests that we should not interpret the single significant coefficients. We conclude for the comparisons of training vs searching that there is little heterogeneity along observed characteristics after conditioning on gender and stratum.³¹

---

³¹We are not aware of state-of-the-art evaluation studies which investigate effect heterogeneity within subgroups defined by gender and elapsed unemployment. Thus, it is difficult to relate our findings to the literature. Most closely related to our analysis is the paper by Lechner et al. (2011). They stratify the treatment samples for men and women in a number of dimensions and find some differences beyond gender differences. According to our reading of the paper, Lechner et al. (2011) do not seem to provide evidence for significant effect heterogeneity within subgroups defined by gender and the elapsed duration of unemployment.

What explains the generally higher treatment effects for women compared with men? While our finding is in line with the literature (see Bergemann and van den Berg (2008) for an overview), this view has been challenged in a recent study by Lechner and Wiehler (2011) who, using Austrian data, argue that positive effects might be overstated if fertility responds to prolonged unemployment. We cannot rule out such an effect because we do not observe pregnancies in our data. However, as indirect evidence against the hypothesis of an effect due to pregnancies, we find that treatment effects do not differ significantly between women below age 40, who are in their fertility phase, and women above age 40, who are mostly past their fertility phase. Other studies (e.g. Lechner et al. (2007), Osikominu (2012)) have attributed part of the gender differences in training impacts to differences in the occupational composition. We could not detect heterogeneous effects across occupations and industries after conditioning on gender and stratum. Our additional results for the outcome indicating a permanent exit from the labor force (see Section C.1 in the Online-Appendix) suggest that the labor market attachment effect of training is slightly stronger for women than for men. This difference may partly explain the gender gap in employment and earnings impacts. Unfortunately, we cannot explore this issue further because our data do not allow to determine exactly whether (and for what reason) an exit from the labor market occurred.

5.3 Assessing the Quality of Our Benchmark Estimates

We provide extensive checks of the quality of our benchmark matching approach, for which the estimates are reported above. First, we investigate the balancing of covariates by comparing the means of 90 covariates between treated and matched comparison groups. The detailed results are contained in Section B.2 in the Online-Appendix, rows labeled ‘Benchmark’ refer to our benchmark matching procedure, rows labeled ‘Raw’ to the raw, unmatched samples. We conduct t -tests Wald-tests of equality of means. The figures in Section B.2 in the Online-Appendix show that, in nearly all cases, the p -values of the Wald-test statistics clearly exceed 5% after conducting our benchmark matching procedure. The Monte Carlo evidence reported by Lee (2012), which is, however, not based on clustered standard errors, suggests that balancing tests of the type used here tend to reject covariate balance too often. Thus, we are confident that covariates are well balanced in our benchmark estimates.

Second, we find that the distributions of the propensity scores and covariates on which we match exactly overlap sufficiently between treated and comparison groups (see Section B.3 in the Online-Appendix). For the comparisons of training versus

searching, the smallest (largest) propensity score values in the treatment group are larger (smaller) than the minimum (maximum) in the comparison group in most cases (see Section B.3.1 in the Online-Appendix). In the few cases where this does not hold, the extrema in the treated sample do not exceed the corresponding extrema in the comparison sample by more than five percentage points. For the comparisons of one training program against another, the overlap is slightly worse but still sufficient. Thus, we do not need to trim any of the propensity score distributions of the treated. Further, Section B.3.2 in the Online-Appendix shows how many treated and comparison observations had to be dropped because of empty cells for any of the covariates on which we match exactly. For the comparisons of training versus searching no treated observations are dropped. For the pairwise comparisons of training programs, between 1 and 3% of the treated observations are dropped, except for CFT versus STT where 17% cannot be used.

Third, as a balancing test on the employment history, we compute placebo treatment effects also for the seven years before the start of the unemployment spell considered. This is also similar to the preprogram test of Heckman and Hotz (1989), but note that we partly match on employment history. If our matching procedure works well, treatment effects should be zero in that period. In 23 out of the 24 cases depicted in figures 3 to 6, pretreatment outcomes of treatment and comparison groups do not differ significantly during any of the seven years before the start of the current unemployment spell. A slight outlier is the earnings impacts of STT for women in stratum 1 (panel A of figure 4) where treated individuals have significantly *lower* earnings in year six before the start of unemployment. Consequently, the earnings gain induced by the program may be even larger than the estimates suggest. Section C.2 in the Online-Appendix contains the graphical results for the comparisons of one training program against another. Here again, employment outcomes are not significantly different in any of the preceding seven years in any of the graphs.

Fourth, as in Huber et al. (2010), we examine the existence and role of overly influential observations on our benchmark estimates in Section B.4 of the Online-Appendix. Only in one out of the 24 cases a comparison person has a weight that exceeds 4%, i.e. for the case ‘STT vs CFT, males, stratum 3’. In ten cases – all of them concerning pairwise comparisons of training programs –, the maximal weight exceeds 2%. The difference between the benchmark estimates and those obtained without the comparison observations with weights above 4% and 2%, respectively, is in general small. In terms of kernel matching weights the influence of the most important comparison individuals is thus much smaller than reported in Huber et al. (2010). This is due to the fact that we align individuals by elapsed unemployment

duration and not by actual and hypothetical program start dates, which implies a smaller pool of potential comparison units than in our case. Correspondingly, excluding (trimming) the most influential comparison individuals (even with a lower threshold than used in Huber et al. (2010)) has little effect on the estimates. Thus, in our benchmark estimates we do not need to trim the most influential observations.

Finally, we investigate whether after matching there remains mismatch between treated and comparison units. We use the outcome regressions on the treatment dummy, personal characteristics and their interactions with the treatment dummy to compute ‘mismatch corrected’ treatment effects (see Section 4.6). The coefficient on the treatment dummy gives the treatment effect net of any mismatch in personal characteristics that may have remained after matching. If our matching approach works well, the mismatch corrected average treatment effects should coincide with uncorrected average treatment effects. The results reported in table 5 show that the uncorrected ATT lies in no case outside of the 95%-confidence interval of the mismatch-corrected ATT. Thus, mismatch does not appear to be a problem.

6 How Sensitive Are Impact Estimates to Data Features and Methodological Choices?

Given the richness of our data and the wide range of possible specification choices, it is a highly relevant question to what extent data features and methodological choices influence the empirical results. The following analysis is therefore meant to contribute to the understanding of what matters for the design and outcome of an evaluation analysis. In the following, we start from our most comprehensive specification (whose results we reported in the previous sections) as a benchmark. We then drop or modify certain specification features, e.g. particular information in the data or the matching approach. We analyze the sensitivity of the results for employment in a given time period. The complete sensitivity analyses are available in Section E in the Online-Appendix. Here we only report results for selected cases in order to illustrate the range of variation that may result as a consequence of data or methodological changes.

6.1 Sensitivity Analysis 1: Employment History

In our first sensitivity analysis, following Card and Sullivan (1988), Heckman and Smith (1999), and Dolton and Smith (2011), we investigate how much difference

it makes to condition on past labor market outcomes in a flexible way. Starting from our benchmark specification, we first drop the requirement of exact matching on four-year history sequences. In a second step, we remove in addition the history sequence dummies from the propensity score. In the third step, the remaining seven-year history variables are dropped. The fourth step further removes all information on the benefit history. In a fifth step, we return to the benchmark specification but omit the requirement to exactly match on the elapsed unemployment duration within strata. Note that this fifth step is conceptually different from the first four steps as these omit information about the pre-unemployment period, while the fifth step omits information about the period after the beginning of unemployment.

— Figures 7 and 8 here —

The results of this exercise are shown in figures 7 and 8. The left column always shows the case ‘CFT vs searching for men in stratum 2’, while the right column presents one further case illustrating what else may happen. The results show that omitting the requirement of exact matching on employment histories in general changes impact estimates only a little. In the case of ‘CFT vs searching for men in stratum 2’ treatment effects are pulled downwards by a moderately small amount suggesting that not controlling for employment sequences leaves too many relatively successful individuals in the comparison group. However, this result does not generalize to the other cases (see Section E.1 in the Online Appendix). More typical is what is shown for the case ‘SST vs searching for men in stratum 2’ (panel B of figure 7), where not conditioning on employment sequences makes hardly a difference. If anything, the estimates are slightly pulled upwards compared with the benchmark. The same is true if in addition the employment history dummies are dropped from the propensity score and the window of pre-unemployment information is reduced from seven to three years (steps 2 and 3). If, in addition, benefit-related information is dropped, the estimates also change very little (step 4, figure 8).

By contrast, dropping the requirement of exact matching on elapsed unemployment duration in the benchmark specification (step 5, figure 8) considerably reduces estimated treatment effects in most cases, especially for individuals with long unemployment durations. The case ‘SST vs searching for men in stratum 3’ involves the strongest change: positive treatment effects are completely eliminated. The explanation is that treated individuals tend to be incorrectly matched with comparison group members who have already found employment if one does not control for the elapsed unemployment duration at program start. This pulls down estimated treatment effects.

Apart from the result that matching on elapsed unemployment duration is important, our results suggest that conditioning flexibly on details of the employment history makes little difference (steps 1 to 4). This is in contrast to some of the results reported in Card and Sullivan (1988), Heckman and Smith (1999), and Dolton and Smith (2011). Unlike in these studies, our evaluation sample is already constructed in a way that strongly conditions on past labor market experience. If this basically eliminated the heterogeneity, it would not be surprising that the way in which we control for past employment information does not matter that much. Therefore, we further investigate the role played by the employment history variables in driving the selection into treatment in our application. First, employment and benefit related variables are in many cases significant determinants of the estimated propensity scores (see Section B.1 in the Online-Appendix).³² Second, our detailed balancing tests reported in Section B.2 in the Online-Appendix show that our treatment and the comparison samples significantly differ with respect to most employment and benefit variables before matching. Moreover, when we carry out those balancing tests for selected variants of our sensitivity analysis, we find that matching based on more parsimonious specifications (e.g. steps 2 and 4) fails to balance important employment and benefit variables. Therefore, it is even more remarkable that omitting the detailed information contained in steps 1 to 4 causes such little changes in the estimated treatment effects.

6.2 Sensitivity Analysis 2: Personal Characteristics

In our second sensitivity analysis, we vary the amount of personal information used for matching and the way in which it enters the propensity score. First, we investigate what happens if we omit the personality variables (i.e. information on whether the person took part in a program providing psychosocial support in the past, whether there were signs of a lack of motivation during previous unemployment spells, information on past penalties or past dropout from a program, the number of job proposals received, and the desire to change the occupation). In a second step, we omit in addition rich personal information which was included in the propensity score in the benchmark specification (i.e. information on household type, number and age of dependent children, marital status, information on health, disability, previous part-time employment and reasons why the last job was ended). In a related

---

³²An exception from this are the four-year history dummies. Apart from the case of STT vs searching for women, the four-year history dummies are insignificant in the estimated propensity score. Other employment variables such as the number of days employed or receiving benefits in any year before unemployment start are often significant.

but separate step 3, we compare the benchmark specification with a specification in which all variables are only used in a fixed, basic way without undertaking a separate specification search for each group. In step 4, we compare our benchmark specification to one in which we only include variables used in the evaluation study by Mueser et al. (2007) which represents a typical application in this field. In step 5, we drop the matching on the propensity score. In step 6, we omit in addition the exact matching on employment histories and align treated and comparison units only with respect to their prior unemployment experience and calendar time.

— Figure 9 here —

The results in panels A and B of figure 9 show that omitting the personality variables changes impact estimates surprisingly little as long as all other variables of the benchmark specification are controlled for. By contrast, omitting other personal information (panels C, D) may have a considerable effect, not so much in the case of CFT vs searching, but in practically all other cases (especially for STT after short unemployment). The direction of the changes is not entirely consistent, but omitting rich personal information tends to overstate treatment effects, suggesting that treated individuals generally have favorable personal characteristics. The fact that the inclusion of personality variables does not change much suggests that the kind of information they represent are already largely included in the employment history variables and in the other personal characteristics. This is in contrast to the findings in Dolton and Smith (2011), where similar attitudinal variables make a big difference.³³ The results for step 3 (panels E, F) suggest that varying the amount of specification search leads to relatively minor changes in most cases. The changes also do not display any clear direction.

Mueser et al. (2007) (henceforth MTG) provide a recent state-of-the-art evaluation of job training programs based on administrative data from Missouri. Their setup is representative of what is possible using administrative data taken from different sources in the US and other countries. Our data are also based on administrative registers but seem to be richer in terms of individual characteristics and details of employment histories. Step 4 compares our benchmark specification to a basic specification which is as close as possible to the one used in MTG (p. 768, basic

---

³³Note, however, that the attitudinal variables in Dolton and Smith (2011) were directly elicited in a survey, while our variables are derived from information on program participation, receipt of unemployment benefits and other information available to the caseworker. The population examined in Dolton and Smith (2011) also differs a lot from that considered here.

demographic and educational variables – i.e. age, education, occupation, nationality and region – as well as four quarters of earnings history and four quarters of employment dummies based on whether earnings in the given quarter were positive). We stick to our requirement to exactly match on the unemployment duration prior to treatment even though MTG do not measure unemployment durations at a monthly frequency. This implies that the treatment and comparison groups in MTG are more heterogeneous compared to our sample. Dropping this requirement has a strong impact on our results as step 5 of sensitivity analysis 1 has shown, but it may be less important in the institutional environment of MTG in which program participation does not require unemployment.

— Figure 10 here —

The results for step 4 are shown in panels A and B of figure 10. It turns out that in most cases treatment effects are overstated if matching is based on this more parsimonious information set (panel A, ‘CFT vs searching for men in stratum 2’, is relatively atypical in that respect, most cases look more like panel B, ‘CFT vs searching for men in stratum 3’). This specification differs from our benchmark specification mainly in dropping detailed history information and rich personal information. Our detailed sensitivity analysis on history information showed that dropping this kind of information led only to small changes in the results, as long as the requirement of the same elapsed unemployment duration before program start is not dropped (we do not drop this requirement here). But dropping rich personal information such as information on family type and health status, leads to an overstatement of the treatment effects. In fact, part of the figures of the step 2 of sensitivity analysis 2 (panels C and D of figure 9 and the respective graphs in Section E.2.2 in the Online-Appendix) look very similar to the figures of the sensitivity analysis involving the MTG specification for the propensity score. The remaining difference must be due to dropping even more information from the propensity score (mainly industry variables, a white collar dummy and exact measurement of days in employment during the last three years) and to dropping detailed history information and rich personal information at the same time.³⁴

Panels C-F of figure 10 show the effects of omitting the propensity score altogether (step 5 and 6). In fact, the results of these steps are remarkably similar to the results

---

³⁴In line with our results, Lechner and Wunsch (2011) also find that the availability of a rich set of covariates, including e.g. information on health, job search behavior and the timing of unemployment and program start, matters to account for the selection into treatment.

of the parsimonious MTG specification (panels A, B, step 4). Again, the case ‘CFT vs searching for men in stratum 2’ is relatively untypical for the range of results. In general and particularly for women, estimated treatment effects are pulled upwards considerably when omitting propensity score matching. This demonstrates that treated individuals are in general positively selected (because treatment effects are lower when personal characteristics are accounted for). The similarity between the estimates that omit the propensity scores altogether and the relatively parsimonious MTG specification reinforces the point that such a parsimonious specification may fail to provide unbiased treatment effect estimates. Furthermore, the similarity of the estimates from step 5 and step 6 underscores that this bias is not related to the omission of exact matching on employment histories.

As before, it is interesting to have a closer look at which of the variables representing rich personal information are significant determinants of the propensity score. In particular, it turns out that for both men and women variables relating to family type, presence of children, and health or disability status significantly influence the probability of being treated (see Section B.1 in the Online-Appendix). In many cases, variables describing personality information also significantly enter the propensity score, although ignoring this information did not make much difference for our estimated treatment effects (see step 1 of the current sensitivity analysis). For example, information on whether there were any penalties or signs of lack of motivation mattered for both men and women in many of the cases considered. We also investigated to what extent treatment and comparison groups were unbalanced with respect to the personal information defined above (see Section B.2 in the Online-Appendix). It turns out that there were significant imbalances with respect to the second set of variables (children, family type, health etc.) for men and women in the case of STT vs searching both in the raw data and after the imperfect matching of step 2. This is not, or to a much lesser extent, the case for the the variables dropped in step 1 (past penalties, lack of motivation etc.). For CFT vs searching, the number of significant imbalances before and after matching was low for both groups of variables.

6.3 Sensitivity Analysis 3: Information on Other Programs

In comparison to the data used in MTG and in many other evaluation studies, our data are also more comprehensive in another important respect: they contain information on all possible programs a person might participate in. Hence, our third sensitivity analysis addresses the point to what extent the information on other

programs matters. Specifically, we compare our benchmark specification to one that ignores information on programs other than the one in question. This mimics the situation in MTG who acknowledge that an unknown number of comparison group members may have participated in other programs. In a second step, we combine this with the MTG specification of the propensity score.

— Figure 11 here —

The results for the comparison with a situation where no information on other programs is available are shown in panels A and B of figure 11. The estimated effects are often pulled downwards as in ‘CFT vs searching for men in stratum 2’ and ‘STT vs searching for women in stratum 1’. The reason for this seems to be that, if information on other programs is available, participants enrolling in the same stratum in these other programs are not in the comparison group. Including them in the comparison group pulls the estimated treatment effects upwards to the extent that the participants in the other programs are locked into these programs and have low employment rates. When judging the value of this information one should consider that the lack of information on other programs is likely less severe for countries with a relatively low program density (such as the US) than for those continental European countries with a high program density.

The results for step 2 (in addition, parsimonious propensity score) are shown in panels C and D of figure 11. Here, the tendency to pull down estimated treatment effects by ignoring information about other programs is more than offset by the tendency to overstate treatment effects if the positive selection of training participants is not accounted for. Even if not completely generalizable to every institutional environment, our results demonstrate that limitations in data availability may have sizeable effects on evaluation results.

6.4 Sensitivity Analysis 4: 12-Months Evaluation Window

The purpose of this section is to investigate how estimated treatment effects change when the width of the evaluation window is changed. In Section 4, we argued that a dynamic rather than a static evaluation setup was appropriate for our application. In the following, we define a 12-months evaluation stratum and compare its results to those obtained in our three-strata benchmark setup.³⁵ We expect two ef-

---

³⁵For example, Wunsch and Lechner (2008) use a 18-months stratum in their main analysis and provide results for a 12-months stratum in a sensitivity analysis. In addition, they use the method

fects when going from our finer stratification scheme to a single 12-months window. First, as pointed out by Fredriksson and Johansson (2008), excluding individuals from the comparison group who receive treatment within a time window of 12-months conditions on their future outcomes. Such future participants represent a negative selection because they would not receive treatment in the future if they had found employment before. Excluding them is therefore likely to pull down estimated treatment effects. Second, using a 12-months evaluation window does not account for potential interactions between elapsed unemployment duration and covariates as only one propensity score is estimated for the whole sample, rather than three different ones referring to the beginning of each stratum.

— Figure 12 here —

Figure 12 presents the comparison between the results of a 12-months evaluation window and the results from our benchmark aggregated over the three strata used there.³⁶ Surprisingly, the two sets of results are rather close together. It turns out however that this is a coincidence and the result of two countervailing effects: excluding future participants pulls down estimated treatment effects, while using a pooled instead of a stratum-specific propensity score does the opposite. Panels A, C, and E of figure 13 show the treatment effects obtained when we use our benchmark propensity scores but exclude comparison persons who receive training at a later point during the first 12 months of unemployment (step 2).³⁷ Panels B, D, and E show the treatment effects obtained when we use the pooled propensity score in addition to excluding comparison persons who are treated within 12 months (step 3).³⁸ As expected, excluding future participants pulls down the estimated treatment effects (step 2). The effect of using pooled rather than stratum-specific propensity scores is given by the difference between step 2 and step 3 (the two steps only differ in their use of the pooled vs the stratum-specific propensity scores). It can be seen that using a pooled propensity score pulls up estimated treatment effects considerably.³⁹

---

of hypothetical program starts which is investigated in the next section.

³⁶For the aggregated benchmark, we use the individual treatment effects computed with the stratum-specific matching procedure and aggregate them across all participants in the three strata.

³⁷We only exclude future participants when calculating the matching weights. The propensity scores are exactly as in the benchmark, i.e. they are estimated on samples that include participants in later strata.

³⁸The graphs show the case ‘STT vs searching for men’. The results for women and for CFT are very similar, see Section E.4 in the Online-Appendix.

³⁹We found that the correlation between the pooled propensity score and the stratum-specific propensity scores decreased when going from earlier to later strata.

Combining the effects of excluding future participants, using a pooled propensity score, and aggregating over strata yields the result of the 12-months evaluation window shown in figure 12.

— Figure 13 here —

As a further sensitivity analysis related to a 12-months evaluation window, panels A and C of figure 14 present the results of a 12-months stratum but in which we omit matching on elapsed unemployment duration, history sequences and calendar month of unemployment start. This drastically pulls down the estimated treatment effects because the treated individuals are to a large extent matched to comparison persons who have already found employment. As expected, this effect is much larger than the one observed in step 5 of sensitivity analysis 1 in which we omit matching on elapsed unemployment duration within the smaller strata considered there. Finally, panels B and D of figure 14 show that this is exclusively due to the omission of elapsed unemployment duration and not due to the omission of history sequences and calendar month (the analysis in panels B and D of figure 14 still does not match on the latter but it is already very close to the results of the original 12-months window given in step 1 in which we match on everything).

— Figure 14 here —

6.5 Sensitivity Analysis 5: Hypothetical Program Starts

In this sensitivity analysis, we compare our benchmark estimates (aggregated over the three strata) with the results from a 12-months evaluation window where we align comparison units to treated units according to their hypothetical program start (HPS). The method of hypothetical program starts is due to Lechner (1999) and is also used, e.g., in Lechner and Wunsch (2011) and Wunsch and Lechner (2008). To construct the hypothetical program start dates for the comparison units, we regress the log program start dates of the treated persons on a set of covariates measured at the beginning of unemployment. Then we use the estimated coefficients plus a draw from the empirical distribution of the residuals to predict a start date for the nonparticipants. Nonparticipants whose predicted program start date lies after the end of their unemployment spell are dropped. The comparison group members are then matched based on the similarity of their hypothetical program start date with the start dates of the participants. In addition, the program start dates enter the

propensity score model. Similar to Wunsch and Lechner (2008), we perform kernel matching with respect to the Mahalanobis distance of the estimated propensity score and the program start.

Our sensitivity analysis includes three steps. In step 1, we use a 12-months stratum in combination with the hypothetical program starts methodology where the program starts are included as dummies in the propensity score. In addition, we exactly match on the calendar month of unemployment start and on four-year employment history sequences as in our benchmark. In step 2, we omit the exact matching on calendar time and history sequences. Step 2 is closest to the scenario used in Wunsch and Lechner (2008). In a final step 3, we proceed as in step 2 but include the hypothetical program starts only linearly in the propensity score.

— Figure 15 here —

A summary of the results using the HPS methodology is shown in figure 15 (see Section E.5 in the Online-Appendix for further results). The most conspicuous differences to the benchmark relate to the more pronounced negative treatment effects. Qualitatively, they are much closer to those in Wunsch and Lechner (2008), who use the same data and study similar programs, than to our benchmark.⁴⁰ The comparison of steps 1 through 3 suggests that neither matching on four-year history sequences and calendar time nor the way how the program start dates enter the propensity score make much difference. This is confirmed by further sensitivity analyses (not reported here) varying a number of further aspects of the HPS methodology. For example, we vary the random element in the construction of the start dates (by drawing from a different distribution or by resetting the seed). This made surprisingly little difference. We also impose exact matching on elapsed unemployment duration (represented by the hypothetical start dates in the case of comparison persons) or omit predicted start dates from the propensity score. Finally, we even change the time axis for the treatment effects from counting months since program start to counting months in unemployment. All this did not change the estimated treatment effects in a substantial way.

What explains yet the differences between the HPS results and our benchmark? The use of a 12-months stratum implies excluding future participants, who have disproportionately long unemployment spells, from the comparison group. This

---

⁴⁰There are further differences that may lead to differences in estimated treatment effects. For example, Wunsch and Lechner (2008) pool men and women and consider only persons who receive unemployment benefits or unemployment assistance in the month of program start.

tends to pull down treatment effects (see step 2 of sensitivity analysis 4, panels A, C, E of figure 13). As seen in step 3 of sensitivity analysis 4 (panels B, D, F of figure 13), this effect happens to be offset by the use of a static propensity score, which tends to pull up estimated treatment effects. Our conjecture is that matching with respect to hypothetical programs starts allows one to account for dynamic selection effects in the otherwise static setup of the 12-months stratum (which is a good thing). In the HPS framework, treated individuals are matched with comparison group members who have a similar hypothetical program start date. As the program start date is predicted on the basis of observed characteristics, comparison group members will be similar in terms of their observed characteristics to individuals who receive treatment at a particular point in time. Thus, the downward bias from using a 12-months window is not alleviated by the use of a static propensity score. Actually, the HPS methodology yields even more negative effects than the results obtained for the 12-months window combined with matching on strata specific propensity scores (sensitivity analysis 4, step 2). This suggests that there is an additional negative shift that arises through matching on hypothetical start dates. Comparison units that are matched to early participants appear to be highly positively selected. We see this confirmed in our data as the majority of the controls with early hypothetical program starts exit unemployment within three months after their hypothetical program start. Thus, the HPS methodology seems to introduce a strong interaction between personal characteristics and unemployment experience.

6.6 Sensitivity Analysis 6: Future Program Participation

When estimating the effect of treatment versus searching, the question arises to what extent our results are determined by the fact that the comparison group includes individuals who may participate in some program in the future (i.e. after the end of the stratum considered). Tables 6 and 7 provide some descriptive evidence on the incidence of later program participation in the comparison groups. Future participants make up between 11% and 20% of the raw comparison groups (see column ‘Overall’ in table 6). Their share increases further to 15-27% when we weight the comparison observations with their matching weights (see table 7). For example, in the evaluation of ‘STT vs searching for men in stratum 1’, future participants in the same or in another program represent 22.42% of the matched sample. Table 6 further demonstrates that the incidence of future participation declines over time because the shares in columns 3-5 remain approximately stable while the time intervals of elapsed unemployment lengthen from three to twelve months.

In the following decomposition analysis, we investigate the effects of sequentially excluding the outcomes of future participants from the pool of comparison units. This will shed light on how including or excluding future participation influences the results. In the first step, we proceed as in the benchmark but exclude the information on the outcomes of future participants from the end of their future program onwards. In step 2 of our analysis we exclude the information on their outcomes already from the start of their future program onwards.⁴¹ Finally, in step 3 of our analysis, we exclude future participants from the beginning of the stratum.⁴² Selected results from this exercise are shown in figure 16.⁴³

First, consider the differences between step 1 and the benchmark (in which future participants are fully included). In step 1, potentially positive post-treatment outcomes of future participants are excluded as they are deleted from the comparison group from the end of their future program onwards. We would expect treatment effects in step 1 to be pulled upwards compared with the benchmark estimates because future participants remain only in the comparison group so long as they have negative outcomes. However, in panels A and B of figure 16, treatment effects for step 1 lie clearly below those obtained in the benchmark case. Thus, reducing the weight of future participants in the control group results in a deterioration of treatment effects. This shows that potentially positive outcomes of future participants are unlikely to play any important role in our analysis. It rather implies that once we remove future participants, treated units are matched to a larger extent with unemployed who quickly take up employment. In comparison, future participants represent a negative selection.

Now consider the difference between step 2 (censoring at the beginning of the future program) and step 1 (censoring at the end of the future program). Going from step 1 to step 2 means reducing further the weight of future participants in the comparison group. Consistent with the downwards change from the benchmark to step 1, we now expect treatment effects to be pulled down even further. Indeed, treatment effects are more negative in panels C and D of figure 16 compared with panels A and B. Finally, step 3 completely excludes future participants from the

---

⁴¹A similar approach has been suggested in Fredriksson and Johansson (2008) and de Luna and Johansson (2010).

⁴²In this last step, we also estimate the propensity scores on samples that exclude the future participants.

⁴³The complete range of results can be found in Section E.6 in the Online-Appendix.

control group. Panels E and F of figure 16 show that treatment effects deteriorate again.

To sum up, participants enrolling later during unemployment constitute a disproportionately high share of the comparison group in terms of their matching weights. Removing them pulls estimated treatment effects down for two reasons. First, one systematically deletes persons known to have relatively bad outcomes because these people would not have received treatment later if they had already found employment again. Second, the next best matches to treated units at a given point in time are unemployed with relatively high chances to exit to employment in the near future.

6.7 Statistical Significance of the Sensitivity Analyses

To assess whether the differences between the benchmark and the sensitivity variants considered by us are statistically significant, we computed bootstrap standard errors for selected cases, jointly resampling the benchmark and the sensitivity estimates.⁴⁴ Results for the average differences between the benchmark and different sensitivity variants in year 1 and 2 after program start can be found in Section E.7 in the Online Appendix. It turns out that the differences are in most cases statistically insignificant.⁴⁵ Notable exceptions are step 5 of sensitivity analysis 1 (no matching on elapsed unemployment duration) and step 3 of sensitivity analysis 6 (completely excluding future participants from the comparison group).

More important than whether or not differences between the benchmark and particular sensitivity variants are significant is whether the use of a particular sensitivity variant leads to conclusions that differ substantively from those of the benchmark. This may happen if treatment effects appear statistically significant (insignificant) in the sensitivity variant but insignificant (significant) in the benchmark. As we report in Section E.7 in the Online-Appendix, in a number of sensitivity cases the substantive conclusions would differ from the benchmark, e.g. in step 5 of sensitivity analysis 1 (no matching on elapsed unemployment duration), in steps 2 and 4 of sensitivity analysis 2 (rich personal information), and in steps 2 and 3 of sensitivity

---

⁴⁴We implemented a semiparametric bootstrap where we resampled persons from the combined samples of the benchmark and the sensitivity variants and repeated the matching analysis in every resample with a draw from the asymptotic distribution of the coefficients in the propensity score. We compared this procedure to that of the purely nonparametric bootstrap that we used for our benchmark estimates for all cases and found virtually no differences.

⁴⁵This is also the general finding in Lechner and Wunsch (2011) in the part of the analysis that is comparable to ours.

analysis 6 (future program participation). Note that even if the difference between the benchmark and a given sensitivity variant is not significant in the vast majority of cases, it is still informative about the direction of possible biases and changes in the conclusions based on standard tests of significance of treatment effects.

7 Summary and Conclusions

This paper analyzes and compares the employment and earnings effects of two important types of public sponsored training in West Germany in the early 2000s: short-term training and classroom further training. In light of recent policy reforms fostering shorter training programs, we are particularly interested in the question of how short-term training programs compare in terms of effectiveness to traditional longer-term further training schemes. Another goal of our analysis is to investigate to what extent data features and methodological choices influence the evaluation results.

Our impact estimates suggest that the effectiveness of the different programs strongly depends on the gender of the participants and the timing of program participation during unemployment. We find statistically significant positive employment effects for male and female participants in short-term training and classroom further training who started their training not too early during their unemployment spell. According to our results, both short-term and classroom further training may increase the employment rate of the participants by some 5 to 10 percentage points for men and by an even larger amount for women. Although both types of training programs generate significant earnings gains in the cases in which there are also positive employment effects, earnings gains induced by classroom further training attain €200 per month and are generally larger than those of short-term training (€70 per month for men and €175 for women). The somewhat surprising effectiveness of short-term training regarding employment is also confirmed in pairwise comparisons to classroom further training. While our results suggest some heterogeneity in the effects for different subpopulations, we could not detect further effect heterogeneity regarding the personal characteristics of participants. Part of the explanation why positive treatment effects are more likely for women and people with longer unemployment durations may be that training programs have an ‘attachment effect’, i.e. nonparticipants are more likely to permanently leave the labor market than participants.

Given that short-term and classroom further training show somewhat comparable

employment gains but classroom further training is more effective in terms of earnings, it is not entirely clear whether the policy shift that took place in 2002 and that largely replaced longer-term training courses with inexpensive short-term courses was justified. One has to keep in mind however, that short-term training programs are (due to their much shorter durations) associated with much shorter lock-in periods and considerably lower total costs. Although we lack detailed information on training costs, the results in table 3 combined with the information on average costs in table 1 suggest that an average short-term training course costing only €627 may lead to a long-term employment gain comparable to that of an average further training course costing €6,175. In the short run, longer-term training is dominated by short-term training because of the longer lock-in period associated with longer-term programs. Although a general cost-benefit analysis seems difficult, simple calculations raise some doubts whether the more expensive longer-term programs can be cost-effective. Considering a period of five years and even ignoring the lock-in effect during the first year, providing longer-term further training to 100 unemployed individuals would cost some €617,500 (compared to €62,700 in the case of short-term training). Assuming a positive employment effect of plus 10 percentage points, this would bring on average 10 out of the 100 unemployed back into employment, saving some $10 \times \text{€}1200 \times 12 \text{ (months)} \times 4 \text{ (years)} = \text{€}576,000$ in unemployment compensation the Federal Employment Agency would have to pay otherwise (based on a monthly unemployment benefit level of €1200, see table 1). Note that in many cases, employment effects are smaller or even zero. Even if earnings gains are taken into account, it will take a rather long time before longer-term further training programs amortize their costs. By comparison, the picture looks more positive in the case of short-term training.⁴⁶

Looking at the particular purpose of short-term training analyzed here – profiling, job search assistance, and monitoring – it seems plausible that these programs do indeed help to activate participants who otherwise do not look as intensively for new jobs. Our results are therefore in line with a number of recent studies that document positive effects of increased job-search assistance and activation, see e.g. Blundell et al. (2004), Weber and Hofer (2004), Crépon et al. (2005), Fougère et al. (2005), Hujer et al. (2005, 2006), and van den Berg and van der Klaauw (2006), and Fitzenberger et al. (2012). As a qualification, the post-treatment observation period in our study is rather limited, covering about two and a half years after program start. Studies considering longer time horizons often find that, over the longer run, the balance

---

⁴⁶Note that the framework used by us stresses the possibility that individuals who do not receive treatment now may receive it in the future. A proper cost-benefit analysis would have to take this point into consideration.

shifts in favor of the longer-term training programs, see e.g. Hotz et al. (2006), Lechner et al. (2011), and Osikominu (2012). Further, our study does not account for potential general equilibrium effects (Calmfors (1994)). The limited available evidence (e.g. Boone and van Ours (2009), Crépon et al. (2011), and Speckesser (2004)), does not suggest, however, that there are adverse macroeconomic effects of job-search assistance and training programs.

The second contribution of our paper is a careful sensitivity analysis examining to what extent data features and methodological choices influence our evaluation results. Each time starting from our benchmark specification, we conduct six sets of sensitivity analyses, three relating to data features and three relating to choices of evaluation methodology. Our sensitivity analysis conveys useful information about the adequate design of evaluation studies for similar settings as ours.

In our first sensitivity analysis, we show that flexibly conditioning on possibly remote details of the employment history before the current unemployment spell makes little difference to our estimated treatment effects. On the one hand, this may reflect the fact that we already condition on some aspects of recent prior employment when constructing our evaluation samples. On the other hand, this may be due to the richness of our data along other dimensions, substituting the detailed employment history. More importantly, we find that not conditioning on the exact unemployment experience in the current spell typically changes estimated treatment effects to a large extent because not taking account of this information results in incorrectly matching treated individuals with comparison persons who have already found employment again.

In our second sensitivity analysis, we find that excluding information on the personality and conduct during past unemployment spells, such as whether there were signs of a lack of motivation, past benefit sanctions or past program dropout, does not affect estimated treatment effects. Rather, rich information on household type, number and age of dependent children, marital status, health, disability, and reasons why the last job ended is highly important in order to balance treatment and comparison groups. We also show that using all the available covariate information but without investing in a detailed specification search does not change the estimated treatment effects by much. Further, we compare our benchmark estimates to results we would get if we had the same kind of information as standard state-of-the-art evaluations such as Mueser et al. (2007). Reducing the information to a set of basic demographic covariates may make a considerable difference. In our application, the resulting differences are often as large as if we did not match on the propensity score

altogether.

In our third sensitivity analysis, we address the important point that in many evaluation studies, information on participation in programs other than the one under consideration is not available. In our case, this often leads to strong biases in the estimated effects of training programs. In general, this issue may be less relevant for countries such as the US in which the density of programs is low.

Our fourth, fifth and sixth sensitivity analyses investigate the consequences of different methodological choices. In our benchmark model, we compare persons receiving treatment after a given unemployment experience to persons not yet treated, similar to Sianesi (2004). We stratify program participations into three strata: program start between months 1-3, 4-6, and 7-12 of unemployment. We impose that potential comparison units have the same elapsed unemployment duration as the treated unit under consideration in the month before the program starts and match on strata-specific propensity scores.

In our fourth sensitivity analysis, we show that using a wide 12-months evaluation window rather than three narrower strata implies two different effects. First, such an evaluation window misses interactions between personal characteristics and elapsed unemployment in shaping the probability to receive treatment over time. Second, widening the evaluation window entails that participants enrolling several months later are excluded from the comparison group. This means that the importance of ‘conditioning on future outcomes’ increases because persons are systematically removed from the comparison group who remain unemployed long enough to receive treatment several months in the future. We find that excluding future participants pulls down estimated treatment effects and may thus bias the assessment of how effective a program is. Not conditioning on elapsed unemployment duration has even more drastic consequences in a 12-months evaluation window than with the three narrow ones.

In our fifth sensitivity analysis, we examine the consequences of using different methods to align treated and comparison units in their elapsed unemployment experience. The alternative method of hypothetical program starts (HPS) uses a regression of start dates on personal characteristics in the treated sample to predict hypothetical program starts for the nonparticipants. It then matches treated and comparison units based on the similarity of their start dates and a static propensity score obtained using a 12-months evaluation window. Compared with our benchmark estimates the HPS methodology yields substantially more negative treatment effects. To some extent this is due to the use of a 12-months evaluation window that

excludes future participants from the comparison group. In addition, the matching on hypothetical program starts entails that participants tend to be matched to nonparticipants who take up employment shortly after their predicted start date. These findings are important because they show that evaluation results may depend strongly on the kind of evaluation method used.

Our sixth sensitivity analysis addresses in more detail the consequences of excluding future participants from the comparison group. Doing so may bias estimated treatment effects downwards because members of the comparison group who are known to have bad outcomes (otherwise these individuals would not receive future treatment) are systematically excluded. In our application, future participants make up roughly 20% of the matching weights in the comparison group. Reducing their share pulls estimated treatment effects down as treated units now tend to be matched with positively selected nonparticipants who take up a job shortly. Our results are important as they highlight the differences between ‘dynamic’ approaches in the spirit of Sianesi (2004) and Fredriksson and Johansson (2008) and standard approaches modeled on the static setup of a random experiment. On the one hand, the treatment effects identified in the Sianesi approach differ in a substantive way from those of an ‘experimental’ estimator where participants in future programs do not exist. Our dynamic impact estimates measure the effect of the caseworker’s decision to assign a job-seeker to a program now rather than to continue job search and reconsider her for treatment at a later point in time. On the other hand, excluding future program participants from the comparison group in order to get close to the ideal of measuring a pure ‘experimental’ effect implies conditioning on future outcomes. Our sensitivity analyses suggest that this effect is empirically relevant when future program participations are quantitatively important as in our setting. The differences between the two approaches should be less critical if the program density is low.

References

- Abadie, Alberto and Guido Imbens. 2006. Large sample properties of matching estimators for average treatment effects. *Econometrica* 74: 235-267.
- . 2011. Bias-corrected matching estimators for average treatment effects. *Journal of Business and Economic Statistics* 29: 1-11.
- Barnow, Burt S. 1987. The impact of CETA programs on earnings. A review of the literature. *Journal of Human Resources* 22: 157-193.

- Bender, Stefan, Martin Biewen, Bernd Fitzenberger, Michael Lechner, Sonja Lischke, Ruth Miquel, Aderonke Osikominu, Tobias Wenzel, Conny Wunsch. 2004. Die Beschäftigungswirkung der FbW-Maßnahmen 2000-2002 auf individueller Ebene: Eine Evaluation auf Basis der prozessproduzierten Daten des IAB – Vorläufiger Zwischenbericht Oktober 2004. Department of Economics, Goethe University Frankfurt and SIAW, University of St. Gallen.
- Bender, Stefan, Martin Biewen, Bernd Fitzenberger, Michael Lechner, Ruth Miquel, Aderonke Osikominu, Marie Waller, and Conny Wunsch. 2005. Die Beschäftigungswirkung der FbW-Maßnahmen 2000-2002 auf individueller Ebene: Eine Evaluation auf Basis der prozessproduzierten Daten des IAB – Zwischenbericht Oktober 2005. Department of Economics, Goethe University Frankfurt and SIAW, University of St. Gallen.
- Bender, Stefan, Annette Haas, Christoph Klose. 2000. The IAB Employment Subsample 1975-1995. *Schmollers Jahrbuch* 120: 649-662.
- Bergemann, Annette, Bernd Fitzenberger, and Stefan Speckesser. 2009. Evaluating the dynamic employment effects of training programs in East Germany using conditional differences-in-differences. *Journal of Applied Econometrics* 24: 797-823.
- Bergemann, Annette and Gerard J. van den Berg. 2008. Active labor market policy effects for women in Europe: A survey. *Annales d'Economie et de Statistique* 91-92: 385-408.
- Biewen, Martin and Bernd Fitzenberger. 2004. Neuausrichtung der Förderung beruflicher Weiterbildung. In *Hartz-Gesetze – Methodische Ansätze zu einer Evaluierung*, ZEW Wirtschaftsanalysen, Bd. 74, ed. Tobias Hagen and Alexander Spermann. Baden-Baden: Nomos: 190-202.
- Biewen, Martin, Bernd Fitzenberger, Aderonke Osikominu, and Marie Waller. 2007. Which program for whom: Evidence on the comparative effectiveness of public sponsored training programs in Germany. IZA Discussion Paper No. 2885, Institute for the Study of Labor, Bonn.
- Blaschke, Dieter, and Hans-Eberhard Plath. 2000. Möglichkeiten und Grenzen des Erkenntnisgewinns durch Evaluation aktiver Arbeitsmarktpolitik. *Mitteilungen aus der Arbeitsmarkt- und Berufsforschung (MittAB)* 33: 462-482.

- Bloom, Howard S., Larry L. Orr, Stephen H. Bell, George Cave, Fred Doolittle, Winston Lin, and Johannes M. Bos. 1997. The benefits and costs of JTPA Title II-A programs. *Journal of Human Resources* 32: 549-576.
- Blundell, Richard, Monica Costa Dias, Costas Meghir, and John van Reenen. 2004. Evaluating the employment impact of a mandatory job search program. *Journal of the European Economic Association* 2: 569-606.
- Boone, Jan and Jan van Ours. 2009. Bringing unemployed back to work: Effective active labor market policies. *De Economist* 157: 293-313.
- Bundesagentur für Arbeit. 2001, 2002a, 2003a, 2004a, 2005a. *Daten zu den Eingliederungsbilanzen 2000, 2001, 2002, 2003, 2004*, Nürnberg (various issues).
- . 2002b, 2003b, 2004b, 2005b. *Arbeitsmarkt 2001, 2002, 2003, 2004*. Nürnberg (various issues).
- . 2005c. *Geschäftsbericht 2004*. Nürnberg.
- Busso, Matias, John DiNardo, and Justin McCrary. 2009. Finite sample properties of semiparametric estimators of average treatment effects. Unpublished Manuscript, Department of Economics, University of California, Berkeley.
- . 2011. New evidence on the finite sample properties of propensity score matching and reweighting estimators. Unpublished Manuscript, Department of Economics, University of California, Berkeley.
- Calmfors, Lars. 1994. Active labour market policy and unemployment – a framework for the analysis of crucial design features. *OECD Economic Studies* 22: 7-47.
- Card, David, Jochen Kluve, and Andrea Weber. 2010. Active labour market policy evaluations: a meta-analysis. *Economic Journal* 120: F452-F477.
- Carcillo, Stéphane and David Grubb. 2006. From inactivity to work: The role of active labour market policies. OECD Social Employment and Migration Working Papers No. 36, OECD, Paris.
- Card, David and Daniel Sullivan. 1988. Measuring the effect of subsidized training programs on movements in and out of employment. *Econometrica* 56: 497-530.

- Crépon, Bruno, Muriel Dejemeppe, and Marc Gurgand. 2005. Counseling the unemployed: Does it lower unemployment duration and recurrence? IZA Discussion Paper No. 1796, Institute for the Study of Labor, Bonn.
- Crépon, Bruno, Esther Dufflo, Marc Gurgand, Roland Rathelot and Philippe Zamora. 2011. Do labor market policies have displacement effect? Evidence from a clustered random experiment. Unpublished Manuscript, Centre de Recherche en Économie et Statistique (CREST), Paris.
- Crépon, Bruno, Marc Ferracci, Grégory Jolivet, and Gerard J. van den Berg. 2009. Active labor market policy effects in a dynamic setting. *Journal of the European Economic Association* 7: 595-605.
- de Luna, Xavier and Per Johansson. 2010. Non-parametric inference for the effect of a treatment on survival times with application in the health and social sciences. *Journal of Statistical Planning and Inference* 140: 2122-2137.
- Dolton, Peter and Donal O'Neill. 2002. The long-run effects of unemployment monitoring and work-search programs: Experimental evidence from the United Kingdom. *Journal of Labor Economics* 20: 381-403.
- Dolton, Peter and Jeffrey Smith. 2011. The impact of the UK New Deal for Lone Parents on benefit receipt. IZA Discussion Paper No. 5491, Institute for the Study of Labor, Bonn.
- Dorsett, Richard. 2006. The new deal for young people: effect on the labour market status of young men. *Labour Economics* 13: 405-422.
- Dyke, Andrew, Carolyn J. Heinrich, Peter R. Mueser, Kenneth R. Troske, and Kyung-Seong Jeon. 2006. The effects of welfare-to-work program activities on labor market outcomes. *Journal of Labor Economics* 24: 567-607.
- Fitzenberger, Bernd, Olga Orlanski, Aderonke Osikominu, and Marie Paul. 2012. Déjà vu? Short-term training in Germany 1980-1992 and 2000-2003. *Empirical Economics* (forthcoming), doi: 10.1007/s00181-012-0590-4.
- Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter. 2008. Get training or wait? Long-run employment effects of training programs for the unemployed in West Germany. *Annales d'Économie et de Statistique* 91-92: 321-355.
- Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter. 2006. Imputation rules to improve the education variable in the IAB Employment Subsample. *Journal of Applied Social Science Studies* 126: 405-436.

- Fitzenberger, Bernd and Stefan Speckesser. 2007. Employment effects of the provision of specific professional skills and techniques in Germany. *Empirical Economics* 32: 529-573.
- Fitzenberger, Bernd and Robert Völter. 2007. Long run effects of training programs for the unemployed in East Germany. *Labour Economics* 14: 730-755.
- Fougère, Denis, Jacqueline Pradel, and Muriel Roger. 2005. Does job search assistance affect search effort and outcomes? IZA Discussion Paper No. 1825, Institute for the Study of Labor, Bonn.
- Fredriksson, Peter and Per Johansson. 2008. Dynamic treatment assignment – The consequences for evaluations using observational data. *Journal of Business and Economic Statistics* 26: 435-445.
- Frölich, Markus. 2004. Finite-sample properties of propensity-score matching and weighting estimators. *Review of Economics and Statistics* 86: 77-90.
- . 2008. Statistical treatment choice: An application to active labour market programmes. *Journal of American Statistical Association* 103: 547-558.
- Galdo, Jose, Jeffrey Smith, and Dan A. Black. 2008. Bandwidth selection and the estimation of treatment effects with unbalanced data. *Annales d'Economie et Statistique* 91-92: 189-216.
- Gerfin, Michael and Michael Lechner. 2002. Microeconomic evaluation of the active labor market policy in Switzerland. *Economic Journal* 112: 854-893.
- Hardoy, Inés. 2005. Impact of multiple labour market programmes on multiple outcomes: The case of Norwegian youth programmes. *LABOUR* 19: 425-467.
- Heckman, James J. and V. Joseph Hotz. 1989. Choosing among alternative non-experimental methods for estimating the impact of social programs: The case of manpower training. *Journal of the American Statistical Association* 84: 862-874.
- Heckman, James J., Hidehiko Ichimura, and Petra Todd. 1997. Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Review of Economic Studies* 64: 605-654.
- Heckman, James J., Hidehiko Ichimura, Jeffrey A. Smith, and Petra Todd. 1998. Characterizing selection bias using experimental data. *Econometrica* 65: 1017-1098.

- Heckman, James J., Robert J. LaLonde, and Jeffrey A. Smith. 1999. The economics and econometrics of active labor market programs. In *Handbook of Labor Economics*. Vol. 3 A, ed. Orley Ashenfelter and David Card. Amsterdam: Elsevier Science: 1865-2097.
- Heckman, James J., and Jeffrey Smith. 1999. The pre-programme earnings dip and the determinants of participation in a social programme. Implications for simple programme evaluation strategies. *Economic Journal* 109: 313-348.
- Heinrich, Carolyn J., Peter R. Mueser, Kenneth Troske, Kyung-Seong Jeon, and Daver C. Kahvecioglu. 2010. New estimates of public employment and training program net impacts: A nonexperimental evaluation of the Workforce Investment Act program. Working Paper No. 1003, Department of Economics, University of Missouri.
- Hotz, V. Joseph, Guido Imbens and Jacob A. Klerman. 2006. Evaluating the differential effects of alternative welfare-to-work training components: A re-analysis of the California GAIN program. *Journal of Labor Economics* 24: 521-566.
- Huber, Martin, Michael Lechner, and Conny Wunsch. 2010. How to control for many covariates? Reliable estimators based on the propensity score. IZA Discussion Paper No. 5268, Institute for the Study of Labor, Bonn.
- Hujer, Reinhard, Marco Caliendo and Stephan Thomsen. 2005. Mikroökonometrische Evaluation des Integrationserfolges. *Beiträge zur Arbeitsmarkt- und Berufsforschung (BeitrAB)* 291, Institute for Employment Research, Nürnberg.
- Hujer, Reinhard, Stephan L. Thomsen, and Christopher Zeiss. 2006. The effects of short-term training measures on the individual unemployment duration in West Germany. ZEW Discussion Paper No. 06-065, Centre for European Economic Research, Mannheim.
- Hui, Shek-wai and Jeffrey Smith. 2003a. The labour market impacts of adult education and training in Canada. Research Report, Statistics Canada.
- . 2003b. Issues in the design of Canada's Adult Education and Training Survey. Research Report, Statistics Canada.
- Imbens, Guido. 2000. The role of the propensity score in estimating dose-response functions. *Biometrika* 87: 706-710.

- Jacobebbinghaus, Peter and Stefan Seth. 2007. The German Integrated Employment Biographies Sample IEBS. *Journal of Applied Social Science Studies* 127: 335-342.
- Jespersen, Svend T., Jakob R. Munch, and Lars Skipper. 2008. Costs and benefits of Danish active labour market programmes. *Labour Economics* 15: 859-884.
- Kluve, Jochen and Claus M. Schmidt. 2002. Can training and employment subsidies combat European unemployment? *Economic Policy* 35: 409-448.
- Kurtz, Beate. 2003. Trainingsmaßnahmen – Was verbirgt sich dahinter? IAB Werkstattbericht Nr. 8, Institute for Employment Research, Nürnberg.
- Larsson, Laura. 2003. Evaluation of Swedish youth labor market programs. *Journal of Human Resources* 38: 891-927.
- Lechner, Michael. 1999. Earnings and employment effects of continuous off-the-job training in East Germany after unification. *Journal of Business and Economic Statistics* 17: 74-90.
- . 2001. Identification and estimation of causal effects of multiple treatments under the conditional independence Assumption. In *Econometric Evaluation of Active Labor Market Politics in Europe*, ed. Michael Lechner and Friedhelm Pfeiffer. Heidelberg: Physica-Verlag: 43-58.
- . 2002. Program heterogeneity and propensity score matching: An application to the evaluation of active labor market policies. *Review of Economics and Statistics* 84: 205-220.
- Lechner, Michael, Ruth Miquel, and Conny Wunsch. 2007. The curse and blessing of training the unemployed in a changing economy: The case of East Germany after unification. *German Economic Review* 8: 468-507.
- . 2011. Long-run effects of public sector sponsored training in West Germany. *Journal of the European Economic Association* 9: 742-784.
- Lechner, Michael and Stephan Wiehler. 2011. Kids or courses? Gender differences in the effects of active labor market policies. *Journal of Population Economics* 24: 783-812.
- Lechner, Michael and Conny Wunsch. 2009. Active labour market policy in East Germany: Waiting for the economy to take off. *The Economics of Transition* 4: 661-702.

- . 2011. Sensitivity of matching-based program evaluations to the availability of control variables. IZA Discussion Paper No. 5553, Institute for the Study of Labor, Bonn.
- Lee, Wang-Sheng. 2012. Propensity score matching and variations on the balancing test. *Empirical Economics* (forthcoming), doi: 10.1007/s00181-011-0481-0.
- Martin, John P. 2000. What works among active labour market policies: Evidence from OECD countries' experiences. *OECD Economic Studies* No. 30: 79-113.
- Martin, John P. and David Grubb. 2001. What works for whom: A review of OECD countries' experiences with active labour market policies. *Swedish Economic Policy Review* 8: 9-56.
- Mueser, Peter R., Kenneth R. Troske, and Alexey Gorislavsky. 2007. Using state administrative data to measure program performance. *Review of Economics and Statistics* 89: 761-783.
- Neyman, Jerzy. 1923, edited and translated by Dorota M. Dabrowska and Terrence P. Speed, 1990. On the application of probability theory to agricultural experiments, Essay on principles, Section 9. *Statistical Science* 5: 465-472.
- OECD. 2005. Labour market programmes and activation strategies: Evaluating the impacts. *OECD Employment Outlook*. OECD, Paris: ch. 4.
- Osikominu, Aderonke. 2012. Quick job entry or long-term human capital development? The dynamic effects of alternative training schemes. *Review of Economic Studies* (forthcoming), doi: 10.1093/restud/rds022.
- Park, Norm, Bob Power, W. Craig Riddell, and Ging Wong. 1996. An assessment of the impact of government-sponsored training. *Canadian Journal of Economics* 29: S93-S98.
- Rosenbaum, Paul R. and Donald B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70: 41-55.
- Roy, Andrew D. 1951. Some thoughts on the distribution of earnings. *Oxford Economic Papers* 3: 135-146.
- Rubin, Donald B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* 66: 688-701.

- Schneider, Hilmar, Karl Brenke, Lutz Kaiser, Jakob Steinwedel, Birgit Jesske, and Arne Uhlendorff. 2006. Evaluation der Maßnahmen zur Umsetzung der Vorschläge der Hartz-Kommission – Modul 1b: Förderung beruflicher Weiterbildung und Transferleistungen. IZA Research Report No. 7, Institute for the Study of Labor, Bonn.
- Sianesi, Barbara. 2004. An evaluation of the Swedish system of active labor market programs in the 1990s. *Review of Economics and Statistics* 86: 133-155.
- . 2008. Differential effects of Swedish active labour market programmes for unemployed adults in the 1990s. *Labour Economics* 15: 370-399.
- Speckesser, Stefan. 2004. The aggregate impact of active labour market policy in Germany and the UK, evidence from administrative data. In *Arbeitsmarktpolitik und Strukturwandel. Empirische Analysen. Beiträge zur Arbeitsmarkt- und Berufsforschung (BeitrAB)* 286, ed. Günter Schmid, Markus Gangl, Peter Kupka. Institute for Employment Research, Nürnberg.
- van den Berg, Gerard J. and Bas van der Klaauw. 2006. Counseling and monitoring of unemployed workers: theory and evidence from a controlled social experiment. *International Economic Review* 47: 895-936.
- Waller, Marie. 2008. On the importance of correcting reported end dates of labor market programs. *Journal of Applied Social Science Studies* 128: 213-236.
- Weber, Andrea and Helmut Hofer. 2004. Employment effects of early interventions on job search programs. IZA Discussion Paper No. 1076, Institute for the Study of Labor, Bonn.
- Wunsch, Conny and Michael Lechner. 2008. What did all the money do? on the general ineffectiveness of recent West German labour market programmes. *Kyklos* 61: 134-174.

Figures

Figure 1: Active Labor Market Programs in Germany

Source.—Bundesagentur für Arbeit (2002b-2005b), own calculations. Note.—The category ‘Medium/Long-term Training’ represents the programs we also call ‘Further training’ and ‘Retraining’ in the text.

Figure 2: Distribution of Program Durations in the Evaluation Sample

Figure 3: Employment Effects for STT vs Searching

Note.— Difference in employment rates is measured on the ordinate, pre-unemployment (< 0) and post-treatment (≥ 0) months and years on the abscissa. Dotted lines are pointwise 95% bootstrap confidence intervals. Stratum 1 denotes entry into program during months 1 to 3 of unemployment, stratum 2 during months 4 to 6, and stratum 3 during months 7 to 12.

Figure 4: Earnings Effects for STT vs Searching

Note.—The abscissae measures pre-unemployment (< 0) and post-treatment (≥ 0) months and years. The difference in average real earnings (in Euro, reference year 2000) during the corresponding time interval is given on the ordinate. Dotted lines are pointwise 95% bootstrap confidence intervals. Stratum 1 denotes entry into program during months 1 to 3 of unemployment, stratum 2 during months 4 to 6, and stratum 3 during months 7 to 12.

Figure 5: Employment Effects for CFT vs Searching

Note.—Difference in employment rates is measured on the ordinate, pre-unemployment (< 0) and post-treatment (≥ 0) months and years on the abscissa. Dotted lines are pointwise 95% bootstrap confidence intervals. Stratum 1 denotes entry into program during months 1 to 3 of unemployment, stratum 2 during months 4 to 6, and stratum 3 during months 7 to 12.

Figure 6: Earnings Effects for CFT vs Searching

Note.—The abscissae measures pre-unemployment (< 0) and post-treatment (≥ 0) months and years. The difference in average real earnings (in Euro, reference year 2000) during the corresponding time interval is given on the ordinate. Dotted lines are pointwise 95% bootstrap confidence intervals. Stratum 1 denotes entry into program during months 1 to 3 of unemployment, stratum 2 during months 4 to 6, and stratum 3 during months 7 to 12.

Figure 7: Sensitivity Analysis 1: Employment History

Step 1: No exact matching on employment histories

Step 2: In addition, no history sequence dummies in propensity score

Step 3: In addition, the seven-year history information is dropped from in propensity score

Note.—The bold line refers to the alternative specification. The thin line and the dotted confidence bands refer to the benchmark specification (see figures 3 and 5).

Figure 8: Sensitivity Analysis 1: Employment History, continued

Step 4: In addition, everything related to benefit information dropped

Step 5: Benchmark but without exact matching on unemployment duration within strata

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 9: Sensitivity Analysis 2: Rich Personal Information

Step 1: No personality variables

Step 2: In addition, no rich personal information

Step 3: No detailed specification search

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 10: Sensitivity Analysis 2: Rich Personal Information, continued

Step 4: Comparison with specification similar to MTG

Step 5: No propensity score matching

Step 6: No propensity score matching and no exact history matching

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 11: Sensitivity Analysis 3: Information on Other Programs

Step 1: Omitting information on other programs

Step 2: Omitting information on other programs, MTG specification

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 12: Sensitivity Analysis 4: 12-Months Evaluation Window

Step 1: 12-months evaluation window compared to benchmark aggregated over its three strata

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 13: Sensitivity Analysis 4: 12-Months Evaluation Window, continued (1)

Step 2: Strata specific propensity scores and exclusion of participants within 12 months (panels A, C, E)

Step 3: Pooled propensity score and exclusion of participants within 12 months (panels B, D, F)

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 14: Sensitivity Analysis 4: 12-Months Evaluation Window, continued (2)

Step 4: 12-months evaluation window without matching on unemployment duration (panels A, C)

Step 5: 12-months evaluation window with matching on unemployment duration (panels B, D)

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 15: Sensitivity Analysis 5: Hypothetical Program Starts (HPS)

Step 1: HPS included as dummies

Step 2: HPS as dummies, no extra matching on history/calendar month

Step 3: HPS linearly, no extra matching on history/calendar month

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Figure 16: Sensitivity Analysis 6: Future Program Participation

Step 1: Exclusion from end of future program

Step 2: Exclusion from start of future program

Step 3: Complete exclusion

Note.—The bold line refers to the alternative specification. The thin line and the confidence interval refers to the benchmark specification (see figures 3 and 5).

Tables

Table 1: Average Expenditures on Training and Unemployment Compensation and Average Training Durations in Germany

	2000	2001	2002	2003	2004
Short-Term Training					
Enrollment length	1.2	1.1	0.9	1.0	0.9
Training costs	580	570	658	538	421
Further Training and Retraining					
Enrollment length	8.2	9.3	9.1	10.5	10.7
Training costs	664	664	681	631	627
Subsistence allowance	1,152	1,178	1,188	1,156	1,150
Unemployment Compensation					
Unemployment benefits	1,160	1,189	1,185	1,261	1,313
Unemployment assistance	753	721	727	691	713

Source.—Bundesagentur für Arbeit (2001,2002a-2005a), own calculations. Note.—Rows labeled ‘Enrollment Length’ display the average duration of the program in months. Rows labeled ‘Training Costs’, ‘Subsistence Allowance’ and rows under the heading ‘Unemployment Compensation’ contain the average monthly expenditures (in €) divided by the average monthly number of participants/claimants. Expenditures on subsistence allowance, unemployment benefits and unemployment assistance include social security contributions.

Table 2: Sample Sizes, Program Participation, and Transitions by Stratum

	Stratum 1	Stratum 2	Stratum 3
Males			
Unemployed at start of stratum	32,172	20,335	13,241
Short-term training	912	547	662
Classroom further training	389	251	270
Practical further training	86	71	107
Retraining	263	89	99
Other program	1,171	648	755
No program participation	29,351	18,729	11,348
– Exit to employment	9,016	5,488	3,113
– Unemployed until end of stratum	20,335	13,241	8,235
Females			
Unemployed at start of stratum	20,746	13,800	10,096
Short-term training	693	409	497
Classroom further training	344	194	201
Practical further training	88	61	85
Retraining	262	55	72
Other program	671	332	446
No program participation	18,409	12,749	8,795
– Exit to employment	4,609	2,653	1,939
– Unemployed until end of stratum	13,800	10,096	6,856

Note.—Stratum 1 refers to months 1-3 of unemployment, stratum 2 to months 4-6 and stratum 3 to months 7-12. The first row shows the number of people who are still unemployed at the beginning of the stratum and who have not received a program before. The subsequent rows show their transitions into programs and employment and the survivors in unemployment until the end of the stratum.

Table 3: Average Monthly Treatment Effects on the Treated – Employment

	Stratum 1		Stratum 2		Stratum 3	
	Year 1	Year 2	Year 1	Year 2	Year 1	Year 2
Men						
STT vs Searching	-0.002 (0.013)	0.010 (0.016)	0.019 (0.017)	0.034 (0.020)*	0.030 (0.015)**	0.046 (0.018)**
STT vs CFT	0.163 (0.024)***	0.021 (0.038)	0.173 (0.024)***	0.062 (0.041)	0.149 (0.018)***	0.057 (0.035)
CFT vs Searching	-0.190 (0.018)***	-0.019 (0.026)	-0.113 (0.023)***	0.015 (0.032)	-0.078 (0.017)***	0.084 (0.032)***
CFT vs STT	-0.120 (0.027)***	0.028 (0.037)	-0.100 (0.034)***	-0.026 (0.047)	-0.061 (0.028)**	0.076 (0.044)*
Women						
STT vs Searching	0.008 (0.016)	0.037 (0.018)**	0.053 (0.019)***	0.096 (0.025)***	0.050 (0.021)**	0.128 (0.025)***
STT vs CFT	0.143 (0.022)***	0 (0.041)	0.194 (0.027)***	0.021 (0.056)	0.137 (0.023)***	0.067 (0.047)
CFT vs Searching	-0.134 (0.022)***	0.060 (0.031)*	-0.132 (0.029)***	0.063 (0.046)	-0.087 (0.027)***	0.070 (0.038)*
CFT vs STT	-0.141 (0.028)***	-0.011 (0.040)	-0.161 (0.036)***	-0.012 (0.053)	-0.100 (0.038)***	-0.071 (0.054)

Note.—Stratum 1, 2, 3 refer to programs starting in months 1-3, 4-6, 7-12 of unemployment. Columns labeled ‘Year 1’ refer to the average of the monthly treatment effects across the first twelve months since program start, columns labeled ‘Year 2’ to the average of the monthly treatment effects during the second year since program start. Bootstrapped standard errors based on 250 replications in parentheses. *, ** and *** denote significance at the 10%-, 5%- and 1%- level, respectively.

Table 4: Average Monthly Treatment Effects on the Treated – Earnings

	Stratum 1		Stratum 2		Stratum 3	
	Year 1	Year 2	Year 1	Year 2	Year 1	Year 2
Men						
STT vs Searching	-30.382 (30.723)	4.693 (37.668)	-2.509 (34.285)	47.598 (46.249)	42.626 (30.476)	86.533 (40.003)**
STT vs CFT	314.914 (54.905)***	-55.558 (97.928)	300.415 (63.698)***	160.484 (92.818)*	279.830 (37.369)***	123.143 (70.515)*
CFT vs Searching	-424.437 (51.135)***	6.675 (72.886)	-250.291 (58.449)***	-17.567 (81.542)	-175.549 (38.039)***	179.112 (70.706)**
CFT vs STT	-269.484 (59.720)***	139.711 (91.966)	-193.397 (75.756)**	-142.781 (111.560)	-78.360 (56.994)	220.078 (99.181)**
Women						
STT vs Searching	28.555 (30.098)	86.124 (38.009)**	76.040 (33.035)**	163.312 (45.979)***	46.941 (31.321)	174.039 (40.091)***
STT vs CFT	257.059 (40.874)***	-30.683 (83.324)	307.633 (46.511)***	34.625 (102.059)	177.653 (36.686)***	54.631 (79.961)
CFT vs Searching	-190.229 (45.236)***	193.081 (65.128)***	-220.986 (53.101)***	123.872 (88.116)	-128.160 (45.702)***	148.351 (69.050)**
CFT vs STT	-250.685 (59.480)***	-2.396 (85.558)	-250.348 (63.550)***	4.981 (101.255)	-107.283 (57.787)*	-11.134 (97.949)

Note.—Stratum 1, 2, 3 refer to programs starting in months 1-3, 4-6, 7-12 of unemployment. Columns labeled ‘Year 1’ refer to the average of the monthly treatment effects across the first twelve months since program start, columns labeled ‘Year 2’ to the average of the monthly treatment effects during the second year since program start. Bootstrapped standard errors based on 250 replications in parentheses. *, ** and *** denote significance at the 10%-, 5%- and 1%- level, respectively.

Table 5: Results of Outcome Regressions on Treatment Dummy, Personal Characteristics and Interactions

	Stratum 1	Stratum 2	Stratum 3
STT, Males			
χ^2 -Stat. (p -Value)	14.95 (0.780)	22.02 (0.339)	19.75 (0.473)
Corr. ATT with CI	0.005 [-.021, .030]	0.025 [-.009, .060]	0.037 [.007, .067]
Uncorr. ATT (SE)	0.004 (0.013)	0.026 (0.017)	0.037 (0.015)
STT, Females			
χ^2 -Stat. (p -Value)	23.81 (0.251)	16.02 (0.715)	20.38 (0.434)
Corr. ATT with CI	0.021 [-.010, .052]	0.073 [.033, .112]	0.089 [.047, .131]
Uncorr. ATT (SE)	0.023 (0.016)	0.074 (0.020)***	0.088 (0.021)***
CFT, Males			
χ^2 -Stat. (p -Value)	16.22 (0.703)	26.43 (0.152)	15.51 (0.747)
Corr. ATT with CI	-0.109 [-.148, -.070]	-0.054 [-.100, -.007]	-0.006 [-.050, .038]
Uncorr. ATT (SE)	-0.104 (0.019)***	-0.050 (0.025)**	0.002 (0.022)
CFT, Females			
χ^2 -Stat. (p -Value)	8.10 (0.991)	21.96 (0.343)	15.64 (0.739)
Corr. ATT with CI	-0.038 [-.086, .010]	-0.036 [-.102, .031]	-0.016 [-.075, .042]
Uncorr. ATT (SE)	-0.036 (0.024)	-0.035 (0.033)	-0.010 (0.029)

Note.—Stratum 1, 2, 3 refer to programs starting in months 1-3, 4-6, 7-12 of unemployment. The table summarizes the main results from a regression of the average treatment effect on the treated (ATT) averaged over the first 24 months after treatment start on the treatment dummy, a set of personal characteristics and interactions of the treatment dummy and the personal characteristics, see equation (4). Rows labeled ' χ^2 -Stat. (p -Val.)' show the test statistic (p -value) of a Wald-test of joint significance involving the interactions with the treatment dummy ($H_0 : \delta = 0$ in eq. (4)). Rows labeled 'Corr. ATT with CI' show the coefficient on the treatment dummy (γ in eq. (4)) that corresponds to the estimated mismatch corrected ATT and its confidence interval (CI). For comparison the ATT without regression correction and its standard error are shown in the rows labeled 'Uncorrected ATT'. *** denotes statistically significant at 1 %, ** at 5 %, * at 10 %. All inference is based on 250 bootstrap resamples. For additional results on effect heterogeneity, see Section D in the Online-Appendix.

Table 6: Importance of Later Program Participations

	Total Number of Controls	Share of Future Participants (%)						
		Overall	By Time Period	By Program Type				
			Month 4-6	Month 7-12	Year 2	STT	CFT	Other
Males								
Stratum 1	29,351	17.89	5.47	6.44	5.97	6.35	2.33	9.21
Stratum 2	18,729	19.46	—	10.10	9.36	7.04	2.30	10.12
Stratum 3	11,348	15.45	—	—	15.45	5.78	1.39	8.28
Females								
Stratum 1	18,409	17.92	5.69	7.07	5.16	6.98	2.90	8.04
Stratum 2	12,749	17.65	—	10.20	7.44	6.89	2.66	8.10
Stratum 3	8,795	10.79	—	—	10.79	4.38	1.55	4.86

Note.—Stratum 1, 2, 3 refer to programs starting in months 1-3, 4-6, 7-12 of unemployment. All numbers refer to program vs searching comparisons.

Table 7: Percentage of Later Program Participation in Terms of Matching Weights

	Males			Females		
	Stratum 1	Stratum 2	Stratum 3	Stratum 1	Stratum 2	Stratum 3
STT vs Searching	22.42	24.58	20.36	20.69	27.36	22.01
CFT vs Searching	21.65	22.02	15.33	19.83	24.41	16.52

Note.—Stratum 1, 2, 3 refer to programs starting in months 1-3, 4-6, 7-12 of unemployment. Shares are measured in % of all potential comparison units (see column 1, table 6).