UNIL | Université de Lausanne

FACULTÉ DES HAUTES ÉTUDES COMMERCIALES

DÉPARTEMENT D'ÉCONOMÉTRIE ET ÉCONOMIE POLITIQUE

EVIDENCE-BASED ECONOMIC POLICY

THREE ESSAYS IN APPLIED MICROECONOMETRICS

THÈSE DE DOCTORAT

présentée à la

Faculté des Hautes Etudes Commerciales
de l'Université de Lausanne

pour l'obtention du grade de
Docteur en Sciences Economiques, mention « Economie politique »

par

Katharina DEGEN

Directeur de thèse
Prof. Rafael Lalive

Jury

Prof. Thomas von Ungern-Sternberg, Président
Prof. Lorenz Götte, expert interne
Prof. Pierre Cahuc, expert externe
Prof. Andrea Weber, experte externe
Prof. Renate Schubert, experte externe

LAUSANNE
2014

Year : 2014

EVIDENCE-BASED ECONOMIC POLICY – THREE ESSAYS IN APPLIED MICROECONOMETRICS

Katharina Degen

UNIL | Université de Lausanne
HEC Lausanne
Le Décanat
Bâtiment Internef
CH-1015 Lausanne

IMPRIMATUR

Sans se prononcer sur les opinions de l'auteur, la Faculté des Hautes Etudes Commerciales de l'Université de Lausanne autorise l'impression de la thèse de Madame Katharina DEGEN, titulaire d'un master en Sciences en Economie Politique de l'Université de Lausanne, en vue de l'obtention du grade de docteur en Sciences économiques, mention « Economie Politique ».

La thèse est intitulée :

EVIDENCE-BASED ECONOMIC POLICY
-
THREE ESSAYS IN APPLIED MICROECONOMETRICS

Lausanne, le 17 juillet 2014

Le doyen

Thomas von Ungern-Sternberg

HEC Lausanne
Le Décanat
Tél. ++41 21 692 33 40 | Fax ++41 21 692 33 05
www.hec.unil.ch | hecdoyen@unil.ch

Thesis Committee

Professor Rafael Lalive
Supervisor
Professor of Economics at University of Lausanne

Professor Lorenz Goette
Internal Expert
Professor of Economics at University of Lausanne

Professor Pierre Cahuc
External Expert
Professor of Economics at CREST-ENSAE, Ecole Polytechnique

Professor Renate Schubert
External Expert
Professor of Economics at ETH Zurich

Professor Andrea Weber
External Expert
Professor of Economics at University of Mannheim

University of Lausanne
Faculty of Business and Economics

Doctorate in Economics

I hereby certify that I have examined the doctoral thesis of

Katharina DEGEN

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature : R. Lalive _______________ Date : 23.06.2014_____________

Prof. Rafael LALIVE
Thesis supervisor

University of Lausanne
Faculty of Business and Economics

Doctorate in Economics

I hereby certify that I have examined the doctoral thesis of

Katharina DEGEN

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature : Lorenz Götte Date : 17.7.2014

Prof. Lorenz GÖTTE
Internal member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

Doctorate in Economics

I hereby certify that I have examined the doctoral thesis of

Katharina DEGEN

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature : ____________________ Date : 24 June 2014

Prof. Pierre CAHUC
External member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

Doctorate in Economics

I hereby certify that I have examined the doctoral thesis of

Katharina DEGEN

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature : Schubert Date : 30.6.14

Prof. Renate SCHUBERT
External member of the doctoral committee

University of Lausanne
Faculty of Business and Economics

Doctorate in Economics

I hereby certify that I have examined the doctoral thesis of

Katharina DEGEN

and have found it to meet the requirements for a doctoral thesis.
All revisions that I or committee members
made during the doctoral colloquium
have been addressed to my entire satisfaction.

Signature : A. Weber Date : 30.6.2014

Prof. Andrea WEBER
External member of the doctoral committee

Acknowledgements

I would like to express my sincere and warm thankfulness to my excellent supervisor Rafael Lalive. He shaped my research agenda from the very first moment of my PhD, when he offered me to be part of the *ewz* field experiment. This is when I developed my enthusiasm for applied microeconometrics and the state-of-the art methods of causal inference. Rafael Lalive guided me through a successful and pleasant study time and always helped me getting back on track when I was stuck with one of my projects. I benefited greatly from his knowledge, his experience and his enormous reservoir of ideas and creativity.

My special thanks also go to Lorenz Goette. Working with him was always pleasant and insightful. His optimism was motivating and pushed me to always give my best. I would also like to thank the members of my thesis committee – Prof. Pierre Cahuc, Prof. Andrea Weber and Prof. Renate Schubert – for their insightful comments and challenging questions.

Many thanks also go to my friends and colleagues. They were always helpful and willing to give their best suggestions. We shared many good times having lunch at the *Banane*, sharing coffee-breaks at our office and – of course – having inspiring and good discussions.

I also want to thank my two brothers Reto and Matthias for their open ears and hearts. They have always been a role model for how to achieve goals and they always supported my endeavors. Finally, I want to express my deepest gratitude to my parents Rolf and Elsbeth Degen for supporting me in my long study time. They were always happy to share the many good times with me, but helped me getting through difficult periods. When I was to worry, they always helped me to put things into perspective – for which I am truly and deeply thankful.

à mes parents

Contents

General Introduction

> *"Social science should be at the heart of policy making. We need a revolution in relations between government and the social research community – we need social scientists to help to determine what works and why, and what types of policy initiatives are likely to be most effective."* (Blunkett, 2000)

Evidence-based economic policy rests on the idea that the best available evidence should be at the heart of policy development and implementation. Evidence-based policy is – in contrast to opinion-based policy – understood as an approach that helps people make well informed decisions about policies and programs using systematic empirical evidence. Successful economic policy requires a solid understanding of the causal effects of a policy intervention. Over the past two decades, improvements in data quality, development of more robust estimation methods and better research designs have helped to develop a large toolkit of new instruments to identify causal effects of policies. These tools include laboratory experiments, field experiments, or quasi-natural experiments which can be analyzed using various identification strategies.

Although randomized experiments or quasi-natural experiment can be extremely helpful in analyzing policy questions, their results have to be interpreted with caution (Leamer, 2010; Manski, 2013). First, findings from randomized and natural experiments cannot be easily transferred into other contexts. Analyses that are based on randomized or natural experiments often rely on specific, non-representative samples or certain restrictive assumptions, which makes extrapolations to other domains often difficult. Second, credible policy analysis should stress the assumptions that are needed for identification. Policy conclusions are of limited usefulness if they rest on strong or unsupported assumptions. It is therefore important to acknowledge not only the benefits, but also the limits and uncertainty of policy analyses that rest on evaluations of randomized and natural experiments. Not a single analysis, but rather the conclusions reached by a systematic body of research on specific policy questions should be the basis of policy development. The objective of this dissertation is thus not to provide a direct basis for specific policy recommendations, but should rather be viewed as a first step towards a better understanding of certain aspects of a policy question.

Evidence-based policy is a universal approach. The use of causal identification approaches and empirical evidence to assess the effects of a policy and to identify the underlying behavioral margins is not confined to specific policy areas. This thesis consists of three independent chapters in applied microeconometrics. The first two chapters discuss related topics in applied labor economics. Both chapters rely on large administrative databases and use quasi-natural experiments to identify the causal effects of policy interventions using two distinct methodologies. While the first chapter focuses on a discussion of the empirical findings of reducing potential unemployment benefit duration on post-unemployment outcomes, the second chapter develops – inspired by job search theory – a new approach to learn about the relative importance of reservation wages for non-employment duration and survival probabilities. The third chapter is in the field of behavioral environmental economics and discusses the role of information for electricity consumption. This chapter is based on a randomized controlled field experiment and comes next to a purely randomized experiment situation. While the third chapter is not thematically linked to the first two chapters, all three chapters answer policy-relevant questions and contribute to a better understanding of the causal effects of policy interventions.

The first chapter is jointly written with Rafael Lalive and makes use of a policy reform that changed the entitlement rules for the maximum unemployment benefit duration in Switzerland. From policy-perspective, knowing about the medium-run effects of such a policy change on post-unemployment outcomes is of direct importance: If a reduction of potential unemployment benefit duration affects post-unemployment outcomes, then a policy assessment that focuses only on its impacts on the government budget is too narrow. The fiscal benefit of reducing benefit durations comes at a potentially large cost if reductions to benefit durations deteriorate post unemployment job quality. Conversely, reducing benefit durations might carry a double dividend if it improves labor market chances.

The policy reform reduced the maximum potential benefit duration from 24 months to 18 months for job seekers younger than 55 years in Switzerland in 2003. Because the policy change was not applied to all job seekers, we can adopt a difference-in-differences framework to identify the effects of this reduction on medium-run earnings and employment. We find that this reduction in potential benefit duration increases earnings of job seekers aged 50 to 54 years not only in the first 24 months but also up to 50 months after entering unemployment. Effects on employment are also positive but weaker than earnings effects. The positive medium-run effects are concentrated among job seekers who were previously employed in R&D intensive industries and whose previous occupation consisted mainly of manual tasks. Unemployment insurance can affect medium-run labor market outcomes via its effects on skill depreciation or unemployment stigma among older job seekers.

While the first chapter provides only indirect evidence on the underlying behavioral channels, in the second chapter I develop a novel approach that allows to learn about the relative importance of the two key margins of job search – reservation wage choice and search effort. Knowing more about these two margins is important form a policy-perspective: On the one hand, if reduced search effort is the main driving force for prolonged unemployment spells, increases in unemployment insurance generosity could be coupled with stricter search requirement to curtail unwanted disincentive and moral hazard effects. On the other hand, increasing unemployment benefit duration could be welfare improving, if prolonged unemployment spells are mainly driven by reservation wage effects and allow job seekers to accept better job-matches.

This paper develops a new strategy designed to analyze the relative importance of the two margins for the duration of unemployment. To this end, I separately study exits to wage-improving jobs and exits to wage-declining jobs. Unemployment exit hazards to wage-improving jobs are solely determined by search effort, whereas the exit rate to wage-declining jobs is jointly determined by search effort and reservation wages. I test this in the context of a sharp discontinuity in potential benefit duration from 30 to 39 weeks in the unemployment insurance system of Austria using a regression-discontinuity approach. I provide causal estimates for the effects of prolonged benefits on unemployment duration and survival probabilities. Consistent with reservation wage movements, exits to wage-declining jobs account for around 80 % of the overall unemployment effect. Moreover, analyzing treatment effects on survivor functions highlights that the largest contributions are observed in the time period from 30 to 39 weeks. These results suggest an important role of the reservation wage channel in shaping job search behavior.

The third chapter is jointly written with Rafael Lalive and Lorenz Goette and – while not thematically linked to the two other chapters – provides a nice complement from a methodological point of view. The third chapter describes a randomized field experiment which analyzes the electricity saving potentials of information. Tackling climate change has been identified as one of the most important economic policy challenges of the 21^{st} century and fostering residential energy-efficiency has been identified as a promising avenue. In the past, researchers and policymakers have focused on relative prices as main driving force of energy usage. However, pricing instruments aimed at regulating the energy demand have often been found to be short-lived and difficult to enforce politically. This is why recently the focus of energy conservation programs has shifted towards behavioral approaches – such as the provision of information or social norm feedback – which are widely viewed as an effective and relatively cheap tool to increase energy efficiency. Financed by the Swiss Federal Office of Energy and implemented by Zurich's energy supplier *ewz*, this field experiment is a first class example of evidence-based policy analysis: While of great scientific interest in itself, the study design is closely interlinked with the interests and needs of energy providers and policy makers.

In the field experiment we consider three types of information: (i) real-time feedback on one's own current and past electricity consumption using smart metering technology, (ii) personalized electricity savings tips through expert advice, and (iii) social information about one's own and a peer households' electricity consumption. Real-time feedback through smart meters reduces electricity consumption by 3 to 5 % of daily electricity consumption. This roughly corresponds to turning-off four light bulbs for an hour or turning off television for 1 hour and 45 minutes per day. Effects are detected shortly after installation of the in-home displays and persist over the study period. Moreover, we find that the largest savings are realized in peak hours. At the same time, households substitute part of their electricity consumption towards low-tariff hours. Social information reduces electricity consumption by around 1.5 % of daily consumption as long as feedback is frequent enough. Expert advice improves the perception of how easy it would be to improve energy-efficient behavior, but fails to translate into electricity savings. Utility is unaffected by the treatments indicating that the social benefits of improved information may be offset by negative social pressure effects or the costs of behavioral changes.

Chapter 1

How Does a Reduction in Potential Benefit Duration Affect Medium-Run Earnings and Employment?*

Abstract

We study how a reduction of the potential duration of unemployment benefit receipt (PBD) affects medium-run earnings and employment of job seekers. The analysis is based on a reform that reduced PBD from 24 months to 18 months for job seekers younger than 55 years in Switzerland in 2003. Adopting a difference-in-differences framework, we find that this reduction in PBD increases earnings of job seekers aged 50 to 54 years not only in the first 24 months after entering unemployment but also up to 50 months after entering. Effects on employment are also positive but weaker than earnings effects. The positive medium-run effects are concentrated among job seekers who were previously employed in R&D intensive industries and whose previous occupation consisted mainly of manual tasks. Unemployment insurance can affect medium-run labor market outcomes via its effects on skill depreciation or unemployment stigma among older job seekers.

JEL Classification: C41, J64, J65

Keywords: potential benefit duration, unemployment duration, earnings, employment, policy change

*This chapter is jointly written with Rafael Lalive. Contribution: Data preparation, empirical analysis, writing (jointly with Rafael Lalive). Financial support from the State Secretariat for Economic Affairs (SECO) and the Swiss National Science Foundation (SNF-Project No. 100018-146090) is gratefully acknowledged.

1.1 Introduction

The global crisis that erupted in 2008 put around 25 million worker out of a job (ILO, 2012). Unemployment insurance (UI) is the key first safety net to workers and probably the most important program to feather the effects of crises. All OECD member countries currently have a system of unemployment insurance. Yet the details of the unemployment insurance system vary tremendously across the OECD.[1]

This paper studies whether PBD affects earnings and employment of job seekers in the period of four years after entering unemployment. Understanding whether PBD matters for medium-run earnings and employment is important for at least two reasons. First, a policy assessment of changes to PBD that focuses only on its impacts on the government budget is too narrow if PBD also affects job quality. The fiscal benefit of reducing PBD comes at a potentially large cost if reductions to PBD deteriorate post unemployment job quality. Conversely, reducing PBD might carry a double dividend if reduced PBD improves labor market chances. A pure policy assessment therefore requires more information on the post unemployment effects of PBD. Second, existing discussions of the optimality of unemployment insurance ignore its potential effects on post unemployment jobs (Chetty, 2008; Schmieder et al., 2012a). These formulas need to be adapted if PBD affects job quality.

On a theoretical level, it is not clear how longer benefit duration affects post-unemployment outcomes. Standard job search theory predicts that shorter PBD forces job seekers to be less selective and prevents them from waiting for better job offers (Mortensen, 1977; van den Berg, 1990a). This is likely to decrease reemployment wages. Also, job match quality might be reduced and subsequent jobs would then end earlier. In contrast, shortening PBD might even improve wages and earnings in a context where skill depreciation is important. Reductions in PBD improve labor market chances by shortening unemployment duration (Shimer and Werning, 2006). Alternatively, firms may use unemployment duration as a screening device (Gibbons and Katz, 1992). Evidence indicates that prolonged unemployment duration is detrimental to the hiring chances of job seekers (Oberholzer-Gee, 2008; Kroft et al., 2013).

This paper analyzes a reform to Swiss unemployment insurance that reduced PBD from 24 months to about 18 months for job seekers who were younger than 55. This reform, enacted in July 2003, can be used to measure the role of shorter PBD for older workers in a differences-in-differences design. As expected, we find that the reform significantly reduced monthly unemployment benefit receipt by 6.5 percentage points in the period 18 to 24 months after entering unemployment. Job seekers compensate this reduction in benefits by leaving unemployment for jobs thus increasing employment by 3.3 percentage points (pctp) and labor earnings by 3.7 percent. Interestingly, we find that the positive effects of the benefit reduction

[1] For instance, the net replacement rate for a family earning the average production worker wage with two children ranges from 55 percent in New Zealand to 92 percent in Luxembourg in the initial phase of unemployment in 2011. The picture is different for the long-term unemployed (4 to 5 years into the unemployment spell). A two children family earning the average production worker wage sees 41 percent of that wage replaced in Greece but up to 72 percent in Denmark. This shows that both the benefit level and the degree to which benefits are maintained in the course of the spell varies tremendously across OECD members.

persists beyond the period that is insured by UI. Specifically, employment remains 1.5 pctp higher and earnings stay 3.3 percent higher compared to the situation without the reduction in PBD. Subsample analyses indicate that the post-UI effects are especially important for job seekers coming from R&D intensive industries and for individuals whose previous occupation required manual skills. These analyses suggest that the beneficial effects of reduced depreciation of human capital or improvements in non-employment stigma outweigh the negative effects of reduced reservation wages.

This paper is related to at least three strands of literature. The first strand discusses reduced form evidence on the effects of PBD on unemployment duration.[2] Several US studies estimate the effects on the exit rate from unemployment of variations in PBD that take place during recessions.[3] Early studies, including Moffitt and Nicholson (1982), Moffitt (1985), and Grossman (1989) find significantly negative incentive effects. Meyer (1990) and Katz and Meyer (1990) show that the exit rate from unemployment rises sharply just before benefits are exhausted. Such spikes are absent for non-recipients. More recent work by Addison and Portugal (2004) confirms these findings. In contrast, Card et al. (2007b) show that the spike at benefit exhaustion has been over-stated in analyses that focus on registered unemployment duration. Evidence on the effect of PBD in European studies also finds strong effects.[4] A common objection against these studies is policy endogeneity. Benefits are typically extended in anticipation of a worse labor market for the eligible workers. Card and Levine (2000) exploit variation in benefit duration that occurred independently of labor market condition and show that policy bias is substantial. Lalive and Zweimüller (2004a,b) show similar evidence for the Austrian labor market.

The second strand of the literature discusses whether changes to PBD affect post unemployment job quality. Ehrenberg and Oaxaca (1976) were the first to look at the effect of unemployment insurance on post unemployment outcomes and find positive effects of unemployment benefits on post unemployment wages for different age groups and gender. A number of more recent studies find positive effects of increased UI generosity on post-unemployment wages: Addison and Blackburn (2000) provide evidence for a weakly positive effect of unemployment benefits on post unemployment wages in the US. Centeno and Novo (2006) analyze the relationship between the unemployment insurance system and the quality of subsequent wages and tenure over the whole support of the wage and tenure distributions. They find a

[2] The survey of the literature on the effects of PBD on unemployment duration borrows heavily from a similar section in Lalive et al. (2006).

[3] Fredriksson and Holmlund (2006) give a recent overview of empirical research related to incentives in unemployment insurance. See Green and Riddell (1997, 1993), and Ham and Rea (1987) for studies that focus on Canada.

[4] Hunt (1995) finds substantial disincentive effects of extended benefit entitlement periods for Germany. Carling et al. (1996) find a big increase in the outflow from unemployment to labor market programs whereas the increase in the exit rate to employment is substantially smaller. Winter-Ebmer (1998) uses Austrian data and finds significant benefit duration effects for males but not for females. Roed and Zhang (2003) find for Norwegian unemployed that the exit rate out of unemployment increases sharply in the months just prior to benefit exhaustion where the effect is larger for females than for males. Puhani (2000) finds that reductions in PBD in Poland did not have a significant effect on the duration of unemployment whereas Adamchik (1999) finds a strong increase in reemployment probabilities around benefit expiration. van Ours and Vodopivec (2006a) studying PBD reductions in Slovenia find both strong effects on the exit rate out of unemployment and substantial spikes around benefit exhaustion.

positive impact of unemployment benefits on each quantile of the wage and tenure distribution. Caliendo et al. (2013a) exploit discontinuities in benefit entitlement to identify the causal effect of an extended benefit duration on unemployment duration and on post unemployment outcomes using German data. They find that the unemployed who obtain a new job close to benefit exhaustion are more likely to leave subsequent employment and receive lower wages than than their counterparts with extended benefit duration. Centeno and Novo (2009) identify a liquidity effect of the unemployment insurance system and detect a positive impact on job match quality for individuals at the bottom of the wage distribution. A number of recent papers find virtually no or only small effects on job-match quality: Card et al. (2007a) analyze the effects of cash-on hand and extended unemployment benefit duration on unemployment duration and job-match quality. They find significantly lower job finding rates with extended benefits, but no effects on job-match quality. van Ours and Vodopivec (2008) analyze how a change in Slovenia's unemployment insurance law affected the quality of post-unemployment jobs. They find that reducing potential benefit duration has only small effects on wages, on the duration of subsequent employment and on the probability of securing a permanent rather than a temporary job. Schmieder et al. (2012c) analyze the long-term effects of extensions in UI durations taking into account not only the initial, but also all recurrent non-employment spells. They find significant long-run effects of an extension in UI duration on the duration of non-employment up to three years after the start of the initial spell. Schmieder et al. (2013) show that longer potential benefit durations sharply increase in non-employment duration but lower post-unemployment wages. They explain the wage drops through shifting wage offer distributions, possibly due to skill depreciation or stigmatization. Finally, Le Barbanchon (2012) finds a significant and large effect of benefit duration on unemployment exits to work but no effects on wages or employment in France.

The third strand of the literature discusses policy design. Starting from the original insight of Baily (1978), Chetty (2008) uses reduced form evidence to discuss whether the level of unemployment benefits is set so as to maximize welfare.[5] Schmieder et al. (2012a) discuss optimal potential benefit duration over the business cycle. Haan and Prowse (2010) discuss the employment, fiscal and welfare effects of unemployment insurance using a structural life-cycle model allowing for endogenous accumulation of experience. They conclude that from a welfare point of view, reductions of benefit entitlement should be favored over replacement rate reductions.

This paper complements existing studies on the job quality effects of PBD in at least three respects. First, we focus on employment and earnings, outcomes that can be observed for *all* job seekers. In contrast, by focusing on wages and subsequent job tenure, the existing literature analyzes outcomes that are only observed for job seekers who find employment. Interpreting effects on job finders is challenging due to selection into employment. Second, we adopt a longer time window that allows estimating not only short-term immediate effects but

[5]Also, see Chetty (2009) for a general description of the sufficient statistics approach.

also effects that build up over time. For instance, if shortening PBD reduced the depreciation of job seekers' leadership skills, labor market outcomes will improve only in the medium-run when job seekers had time to demonstrate those better leadership skills. Finally, we perform subgroup analyses by industry and occupation of previous job to shed light on the role of reduced human capital and skill depreciation as a potential explanation for positive medium-run effects.

The remainder of this paper is structured as follows. Section 1.2 discusses the institutional background. Section 1.3 provides information on the data sources and a set of key descriptive statistics. Section 1.4 discusses the econometric framework and our main identification strategies. Section 1.5 presents the main results, and section 1.6 provides a summary and implications of our findings.

1.2 Institutional background

This section discusses the relevant background on unemployment insurance, earnings, and employment in Switzerland.[6] Job seekers are entitled to unemployment benefits if they meet two requirements. First, they must have paid unemployment insurance taxes for at least six months in the two years prior to registering at the public employment service (PES). The contribution period is extended to 12 months for those individuals who have been registered at least once in the three previous years. Job seekers entering the labor market are exempted from the contribution requirement if they have been in school, in prison, employed outside of Switzerland or have been taking care of children. Second, job seekers must possess the capability to fulfill the requirements of a regular job - they must be "employable". During the unemployment spell, job seekers have to fulfill certain job search requirements and participate in active labor market programs in order to remain eligible for benefits.[7] Job seekers who are ineligible for unemployment insurance can claim social assistance. Social assistance is means tested and replaces roughly 76 % of unemployment benefits for a single job seeker with no other sources of earnings (OECD, 1999).

Prior to July 1, 2003, job seekers were eligible for 520 daily benefit payments during a two year framework period. Those 520 benefit days are equivalent to two years of potential benefit duration since a calendar year has 260 work days. The replacement ratio is 80 % for workers earning less than 3,536 CHF.[8] prior to unemployment and not caring for children. The replacement rate decreases gradually to 70 % for job seekers who earned between 3,536 CHF and CHF 4,030 and it stays at 70 % thereafter. Benefits insure monthly earnings up to a top cap.[9] Job seekers have to pay all earnings and social insurance taxes except the unemployment insurance tax rate (which stands at about 2 %). This means that the gross

[6]This section borrows from a similar section in Arni et al. (2013).
[7]See Gerfin and Lechner (2002) and Lalive et al. (2008) for detailed background information on and an evaluation of the active labor market programs.
[8]1 CHF = 0.83 EUR.
[9]The cap is currently at 10,500 CHF per month and stood at 8,900 CHF before the reform.

replacement rate is similar to the net replacement rate. Job seekers keep these entitlements during a framework period of two years. For instance, a job seeker who leaves unemployment after 3 months remains eligible for the remaining months of unemployment benefits during the two year framework period.

The July 2003 reform changed a range of aspects of the benefit system. First, the reform now requires everyone to have contributed for at least 12 out of the 24 months prior to registering for unemployment benefits. Second, the reform reduced PBD for individuals below the age of 55 years to 400 daily benefit payments, or to 18.5 months.[10] Job seekers aged 55 years or older who had contributed for at least 18 months prior to entering unemployment remained unaffected by the reform. Yet job seekers aged 55 years or older who had only contributed between 12 and 17 months to UI also experienced a cut in PBD. Third, the reform increased benefit levels somewhat for low to medium earners to reflect inflation adjustment. In order to achieve this objective, the replacement rate was kept at 80 % for job seekers with insured earnings of up to 3,797 CHF and then gradually reduced over the earnings bracket 3'797 to 4,340 CHF.

From an identification point of view, the following issues are crucial. First, there were no concurrent changes to other social insurance programs in the period around the 2003 reform. This ensures that our estimates pick up the specific consequences of the reform rather than changes to other social programs. Second, benefit rules depend on current age of individuals rather than on age at registration. Also, reforms to the UI system apply to all job seekers, not just to those who register after the reform. We will discuss below how we take this into account in our estimation framework. Third, the reform was signed into force around a time when the Swiss labor market situation was deteriorating. The unemployment rate reached a low of slightly over 1.5 % in the first quarter of 2001 and it increased considerably after the bursting of the "dot.com" bubble to a high of 4 % in the last quarter of 2003. Unemployment decreased first slightly then more rapidly to reach a trough of 2.5 % in the second quarter of 2008. The changing macroeconomic environment will not introduce a bias into our estimates if aggregate demand for work varies similarly for the treatment and control groups in our analysis. We assess this key condition further below.[11]

1.3 Data and descriptive statistics

This section discusses the data and provides first descriptive information about treatment and control groups.

[10] A year counts 260 benefit days. A job seeker who is eligible for 400 benefit payments can therefore claim benefits for 18.46 (=400/260 * 12) months.

[11] Note that our analysis identifies a lower bound on the positive effects. As younger worker's unemployment is more sensitive to the cycle than older workers' unemployment (Clark and Summers, 1981) and the average quality of younger unemployed is likely to be lower, the effects on earnings and employment are likely to be negatively biased.

1.3.1 Data

The study is based on two data sources. The first concerns administrative records of the unemployment insurance register (UIR) database covering information on all individuals registering with the public employment service (PES) between 1999 and 2007. This can be job seekers who are eligible for unemployment benefits, but also individuals who ask the public employment service for assistance. The UIR contains the exact date when a job seeker can start a new job – the unemployment start-date.[12] The UIR also contains information on the date when the job seeker starts her or his new job – the job start date. We measure the duration of unemployment as the number of days elapsed between the unemployment start-date and the job start-date if those two pieces of information are available. We use the de-registration date, the date when the file of a job seeker was closed, as a proxy for the unemployment end-date for individuals who do not start a new job. The database also contains socio-demographic characteristics such as gender, age, education, and marital status.

The second data source contains information on unemployment benefit payments, employment and earnings from the Social Security Administration (SSA). This data covers the universe of all individuals who have contributed to the mandatory first pillar retirement pension system between the period between 1982 and 2010. The social security database can be merged to the unemployment insurance register data through a unique person identifier. The data provides monthly information about earnings from employment and some information on transfer income (e.g. unemployment benefits are included but not social assistance). Moreover, for a subsample of around 35 % of the universe of spells we also observe disability and old-age retirement pensions. For each unemployment spell, we extract a history of 50 months before and 50 months after registration from the SSA database.

We impose a number of additional sampling restrictions on the merged database. First, we only consider individuals aged between 50 and 59 years at the start of the spell of unemployment, in order to avoid confounding effects because of early retirement considerations. Second, the sample contains only individuals who contributed to the unemployment insurance for at least 18 of the last 24 months before getting unemployed. This ensures that all job seekers aged 55 or older kept eligibility to two years of benefits. Third, the reform was applied to in-progress spells. This implies that some individuals in the before-treatment regime could actually have experienced a reduction in PBD while unemployed. In order to reduce this potential source of bias, we exclude job seekers who enter unemployment up to 12 months before the reform in July 2003. Fourth, we only consider individuals who are full-time unemployed in the first month of unemployment.[13] The final sample contains 62,563 spells.

[12]The data also contains date of registration and de-registration. The registration date does not correspond to the start date of the unemployment spell because job seekers need to register with the PES the moment they know they will lose a job. This is typically a quarter before they actually lose their job.

[13]Workers who lose one of two part-time jobs are eligible for UI on the job they lost. These job seekers are part-time unemployed. We focus on the full-time unemployed to achieve a homogeneous sample.

1.3.2 Treatment and control groups

Treatment assignment. Table 1.1 provides information on how we defined treatment and control groups. We assign individuals aged below 55 at the start of their unemployment spell to the treatment group, and individuals aged 55 or older to the control group. Because benefit eligibility is based on current age, benefit eligibility is upgraded to 24 months after the 55th birthday of a job seeker. The treatment group thus contains some job seekers who switch from 400 days to 520 days of benefits. Our estimates thus have to be interpreted as lower bounds. Excluding job seekers who were employed for less than 18 months in the last 24 months prior to the start of the unemployment spell ensures that only job seekers in the treatment group are affected by the cut in PBD. Yet, a potential issue could be that the months employed within a two year window prior to unemployment start do not necessarily perfectly coincide with the two year framework period that determines eligibility for benefits. However, over 85 % of our sample claimed unemployment benefits within 3 months after unemployment start. Eligibility issues should therefore not play a major role.[14]

Table 1.1: Treatment assignment

Age	Prior UI contributions	Benefit entitlement before	after	Group
< 55	$\geq$ 18 months	520	400	Treatment
≥ 55	$\geq$ 18 months	520	520	Control

Notes: This table shows the treatment assignment, which is based on the age at unemployment start.

For each individual unemployment spell we observe a history of monthly unemployment benefits, earnings from employment around unemployment start of up to 50 months before, and up to 50 month after unemployment start.[15] We construct a binary indicator on employment that takes the value 1 if the job seeker has generated positive earnings from employment, and zero otherwise. Also, we define a binary variable for benefit receipt that takes the value 1 if unemployment benefits were positive in a month, and zero otherwise. We observe 22,170 spells of job seekers whose unemployment spell started before the reform was implemented on July 1st, 2003 – 9,529 in the treatment group, and 12,641 in the control group (table 1.2). We observe 40,393 unemployment spells starting after July 1st, 2003 – 17,307 spells belong to the treatment group and 23,086 belong to the control group.

[14] One might think that the regression discontinuity (RD) design could also be implemented (Lee and Lemieux, 2010a). Yet note that benefit eligibility does change discontinuously in age. A job seeker who enters unemployment at age 54 years and 11 months will initially be entitled to 18.5 months of benefits but rapidly up-grade to 24 months of benefits once he or she has celebrated her or his 55^{th} birthday. Alternatively, one could think of using the number of contribution months as a running variable. This is challenging for two reasons. Our records indicate that prior contribution months as measured in the SSA are an imperfect predictor of eligibility. We suspect measurement error in prior contribution months. Second, prior contribution months are also unlikely to satisfy the requirement that the running variable can not be manipulated. For these reasons we have adopted a difference-in-differences framework.

[15] Individuals can appear multiple times in our sample. We observe two or more spells for 8 % of the job seekers in the sample.

Descriptive statistics. Table 1.2 presents selected summary statistics for the treatment ($D_i = 1$) and control ($D_i = 0$) group for spells that start before (columns 1 and 2) and after (columns 3 and 4) the reform. The average employment share over the 50 months before entering unemployment is 92 percent for job seekers who entered unemployment before the reform. The employment probability does not differ across the treated and control groups. Employment prior to the unemployment spell was about 1 percentage point higher for job seekers who start a new spell in the period after the reform. Results are similar for earnings. Job seekers in the control group earn about 5,000 CHF per month (about 4,150 EUR) before the reform, and monthly earnings are about the same for treated job seekers and for both groups after the reform.

The table also presents information on two key pieces of information that we will use to learn more about skill obsolescence and depreciation. The first information is *R&D intensity* of the previous employer. We infer R&D intensity of an industry as the average expenditures for R&D for the neighboring countries of Switzerland (Germany, Austria, France and Italy) over the years 2005 to 2008 at the two digit NACE level. We merge this information to each job seeker based on industry prior to losing job. R&D intensive industries are those that have expenditures that exceed the median expenditure, the remaining industries representing the low R&D industries.[16] The share of job seekers from high R&D industries slightly exceeds 50 % for treated and untreated before the reform. After the reform, the proportion of job seekers from R&D intensive industries decreases slightly to around 49 % and 46 % respectively. The second information is related to the task content of the occupation of job seeker. *Cognitive* refers to job seekers whose previous occupation consisted mainly of cognitive tasks. For the classification of occupations into cognitive and manual task content, we adopt an approximation suggested in Acemoglu and Autor (2011). The authors propose a simple classification of occupations into four broad task dimensions: (1) abstract, non-routine cognitive tasks, (2) routine cognitive tasks, (3) routine manual tasks, and (4) non-routine manual tasks. We further condense the first and second category into a "cognitive tasks" group, and the third and fourth into a "manual tasks" group.[17] Before the reform, the proportion is 50 % for untreated and 49 % for treated individuals respectively. After the reform, the proportion of mainly cognitive skilled job seekers in the control and treatment groups decreases to 48 %.

Experience captures the proportion of job seekers with a continuous work experience of at least 24 months prior to their unemployment spell. The proportion of job seekers with a long work history is around three quarters for spells that started before the reform. After the reform, this proportion slightly increases to 82 % for individuals in the control group, and to 78 % for the treatment group. Around 74 % of the individuals in the control group, and roughly 72 %

[16]High R&D industries are for example manufacture of chemicals and pharmaceuticals, manufacture of computer, electronic and optical products, manufacture of machinery, equipment and motor vehicles, or industries in professional, scientific and technical activities.

[17]The most important occupations requiring cognitive skills are engineers, clericals and occupations in administrative support, sales, and education. The most important occupations requiring manual skills are occupations in construction, in production and manufacture of raw materials, and in services and housekeeping.

of the individuals in the treatment group worked in a leader or expert position. The share of female job seekers varies between 42 % and 45 %. The proportion of Swiss citizens is fairly stable for unemployment spells starting before and after the reform, and amounts to 70 % in the treatment group and around 72 % in the control group. There are no large differences between the four groups relative to their marital status: Around two thirds of the individuals are married, one fifth is divorced, roughly 10 % are singles, and around 4 % are widowed. The largest differences between unemployment starts before and after the reform are found for years of schooling: The share of individuals with less than 7 years of schooling, and between 10 and 11 years of schooling remains fairly stable over time and across treatment and control groups, and increases slightly over time for job seekers with between 8 to 9 years of schooling and with more than 14 years of schooling. The share of individuals with 12 to 13 years of schooling, however, increases largely from around 24 to 26 % before to around 36 to 38 % after the reform. At the same time the share of individuals for whom the attained education level is unknown decreases from around 50 % to 30 % over time. Changes in data quality account for this substantial shift in measured education levels. This shift affected treated and untreated individuals in the same way and will not invalidate our identification strategy. Moreover, except for job seekers with 14 years or more of schooling, education levels do not differ statistically significantly between control and treatment groups before and after the reform.

Column 5 of table 1.2 presents difference-in-differences estimates on the control variables. The null hypothesis that the composition of the treated group did not change can not be rejected for most of the variables. A few characteristics show significant differences between the treatment and the control group before and after the reform. We reject the null hypothesis of no change in the composition of the two groups for the share of job seekers in cognitive occupations, prior work experience, Swiss nationality, marital status, and education. Yet note that the resulting changes in sample composition are small. We find below that accounting for these changes in sample composition does not affect results.

Table 1.2: Selected descriptive statistics

	Before reform		After reform		
Treatment status	$D_i = 0$	$D_i = 1$	$D_i = 0$	$D_i = 1$	DiD
A. Dependent variables (prior to unemployment)					
Employment	0.92	0.91	0.94	0.93	0.00
Earnings	5003.82	5075.77	5204.16	5252.01	-24.09
UE benefits	240.78	250.26	167.47	174.73	-2.21
B. Control variables					
R&D intensity	0.53	0.51	0.49	0.46	0.00
Cognitive	0.50	0.49	0.48	0.48	0.02^{**}
Experience	0.75	0.73	0.82	0.78	-0.02^{**}
Leader position	0.74	0.72	0.74	0.73	0.01
Female	0.42	0.43	0.45	0.46	-0.01
Swiss	0.72	0.70	0.75	0.71	-0.03^{***}
Marital status					
Single	0.11	0.12	0.11	0.13	0.01
Married	0.64	0.62	0.63	0.63	0.02^{**}
Widow	0.04	0.03	0.04	0.02	0.00
Divorced	0.22	0.23	0.23	0.21	-0.02^{***}
Years of schooling					
$\leq$ 7 years	0.04	0.04	0.03	0.04	0.00
8-9 years	0.11	0.11	0.15	0.15	0.00
10-11 years	0.04	0.04	0.05	0.05	0.00
12-13 years	0.26	0.24	0.38	0.36	0.00
$\geq$ 14 years	0.06	0.06	0.08	0.10	0.01^{***}
Other	0.50	0.51	0.31	0.30	-0.01
Observations	475,900	631,589	865,350	1,154,300	
No. of spells	9,529	12,641	17,307	23,086	

Notes: This table shows means of selected variables for the treatment and control group for individuals who registered before or after July 1, 2003 respectively. Column 5 shows the differences in differences. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

1.4 Econometric framework

Empirical strategy. This section presents the empirical strategy we employ for the analysis of the effects of PBD on employment and earnings and discusses the underlying identification assumptions. The specific design of the reform creates a natural control group for which benefit entitlement remained unchanged, and a treatment group for which the PBD was reduced from 24 months (520 days) to 18 months (400 days). In order to discuss estimation and identification assumption, let $Y(1)$ be the treated outcome, and $Y(0)$ the non-treated outcome. $D \in \{0,1\}$ is a treatment indicator that is 1 if an individuals receives treatment, i.e. is below 55 years old at unemployment start, and 0 else. Let Y_0 denote the outcome prior to the reform, and Y_1 the outcome after the reform. The observed outcome after the reform can then be written as $Y_1 = DY_1(1) + (1-D)Y_1(0)$. The difference-in-differences estimator is then given by

$$DiD = [E(Y_1 \mid D=1) - E(Y_1 \mid D=0)] - [E(Y_0 \mid D=1) - E(Y_0 \mid D=0)]$$

The difference-in-differences estimator identifies the average treatment effect on the treated by comparing differences in outcomes between the outcomes of the treated and the untreated before and after the reform. The difference-in-differences estimator can be rewritten as

$$DiD = E(Y_1(1) - Y_1(0) \mid D=1)$$

which corresponds to the average treatment effect on the treated.

The main assumption that has to hold for the difference-in-differences estimator to identify the average treatment effect on the treated in repeated cross sections are parallel time trends for the treatment and control group in absence of the treatment, i.e. $E(Y_1(0) - Y_0(0) \mid D=1) = E(Y_1(0) - Y_0(0) \mid D=0)$.[18] This assumption could be violated for at least three reasons. First, repeated cross sections could differ in terms of sample composition. Second, labor market outcomes might evolve differently across treatment and control groups because their outcomes differ with respect to sensitivity to the cycle. Third, the reform might also have changed the incentives to become unemployed thereby changing the composition of the unemployment inflow.

Validity of the identifying assumptions. We now test each of these reasons for failure of the identifying assumption. We have already presented a test for a change in sample composition (see table 1.2, last column). We do find that the test rejects the null of no change in sample composition for a range of background characteristics. But note that the changes in sample composition are fairly small in an economic sense. We further address changing sample composition by discussing the sensitivity of our results to adding observed characteristics.

Second, we assess whether time trends evolve in a parallel fashion across treated and control groups. We focus on unemployment benefit receipt in 22 to 24 months after job seekers, i.e. benefit receipt in the last quarter of a job seeker's framework period of two years after unemployment start. The last quarter of a job seeker's framework period should be mechanically affected by the reform in July 2003. Plotting benefit receipt by quarter of entry into unemployment for groups that were not affected by the reform will provide a visual test of parallel trends. We also visually inspect time trends after the reform was implemented to see whether the effect of the reform is constant and time trends continue to evolve in a parallel fashion after

[18] See also Lee and Kang (2006) for a detailed discussion of the identification assumptions in repeated cross sections.

the reform has been implemented.

Note that the reform was applied to in-progress spells. This means that treated job seekers start to be affected by the cut in PBD even if their spell started before July 1, 2003. Figure 1.1 shows that the treatment group starts to be affected by the cut in PBD for spells that start after July 1, 2001 because the reform gradually removes the final months of benefit eligibility. For instance, a job seeker starting unemployment on January 1, 2002 will be fully affected by the reform since her or his last 6 months of benefit eligibility will be cut by the reform in July 1, 2003. In other words, the effective PBD for the treatment group reduces gradually from 520 to 400 days for entries into unemployment between July 2001 to January 2002. Finally, for spells that started after January 2002, the treated job seekers get a maximum number of 400 days, whereas untreated job seekers still get 520 days of unemployment benefits.

Figure 1.1: Timing of reform

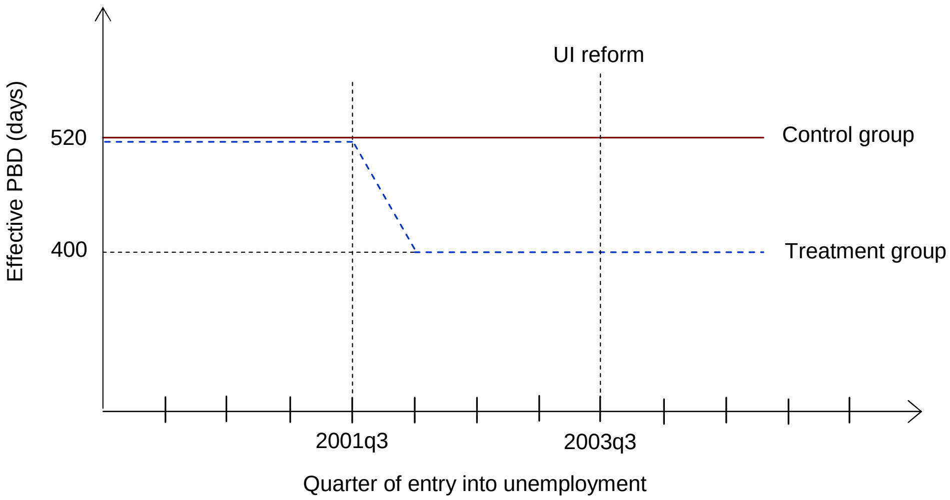

Notes: This figure shows the stylized pattern of effective PBD over the quarter of entry into unemployment for the treatment and the control group respectively.

Figure 1.2 shows unemployment benefit receipt 22 to 24 months after unemployment start of treated and control groups for every quarter between 1999 and 2007. The left hand axis measures the share of job seekers who claim benefits. The right hand axis measures the difference between treatment and control groups. The dashed vertical line in the third quarter of 2001 depicts the first possible date for which effects of the reform are potentially observable. The dashed horizontal line indicates the mean difference between the treated and control group before the reform.

The figure highlights several interesting facts. First, the control group tends to have about 10 percentage points higher benefit receipt than the treated group before the reform because the control group is older than the treated group. Second, benefit receipt varies quite strongly over the period 1999 to 2007 – very much in line with the business cycle. Third, time trends are roughly parallel in the period before the reform, especially so for job seekers entering unemployment between 2000 to the second quarter of 2001. Fourth, the reform led to a substantial reduction in unemployment benefit receipt. This effect can be seen for job seekers entering unemployment in the third quarter of 2001 and later. Finally, the difference in benefit receipt remains approximately constant for all job seekers entering unemployment after the reform. This evidence is therefore consistent with parallel trends in benefit receipt also after the reform.

Figure 1.2: Time trends 22 to 24 months after unemployment start

(a) Benefit receipt

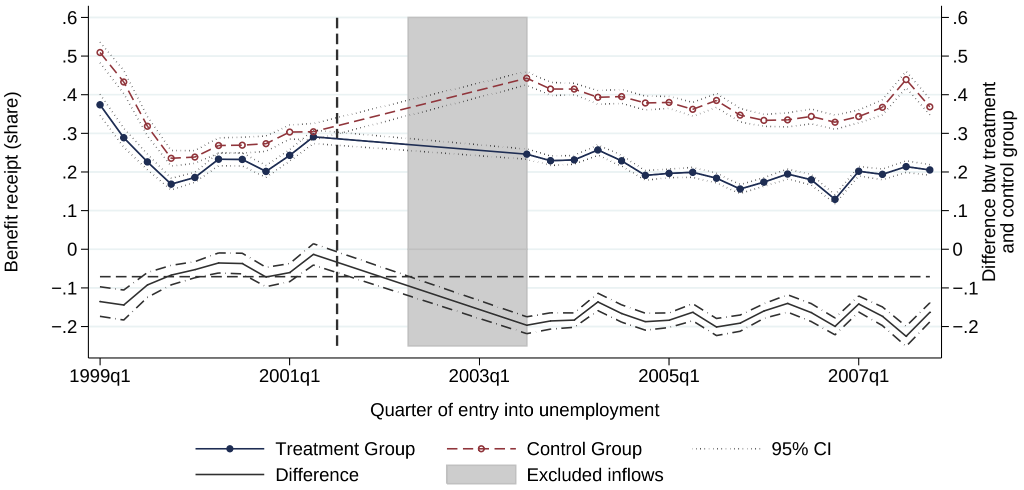

Notes: This figure shows the time trends for benefit receipt together with the 95 % confidence interval. The solid line at the bottom shows the difference between treatment and control group together with the 95 % confidence interval. The dashed vertical line at 2001q3 depicts the first possible date for which treatment effects are detectable. The dashed horizontal line shows the mean difference in benefit receipt between the treated and the control groups. Shaded area indicates that no data is available for that time period (inflow between July 2002 and June 2003 was omitted from the analysis).
Source: Own estimations based on merged UIR-SSA database.

Are trends in employment and earnings also parallel? Figures 1.3a and 1.3b report a similar analysis for employment and earnings. Results indicate that trends are parallel for both outcomes for spells that start before the third quarter of 2001. This evidence suggests trends in outcomes are similar. Moreover, both figures indicate that employment and earnings patterns start to differ from the third quarter of 2001 onwards. These graphs suggest that the assumption of parallel trends is plausible and that the reform effects build up over time as would be expected also for employment and earnings.

Figure 1.3: Time trends 22 to 24 months after unemployment start

(a) Employment

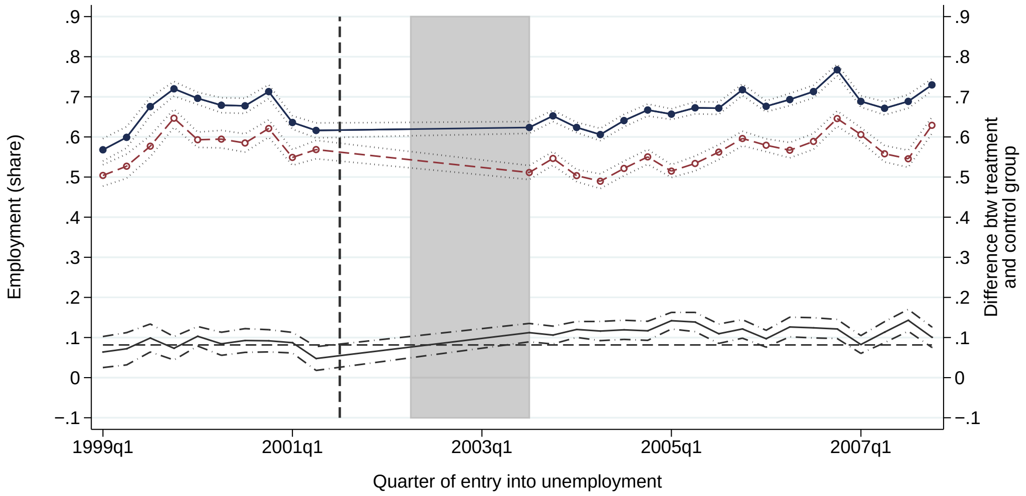

(b) Earnings

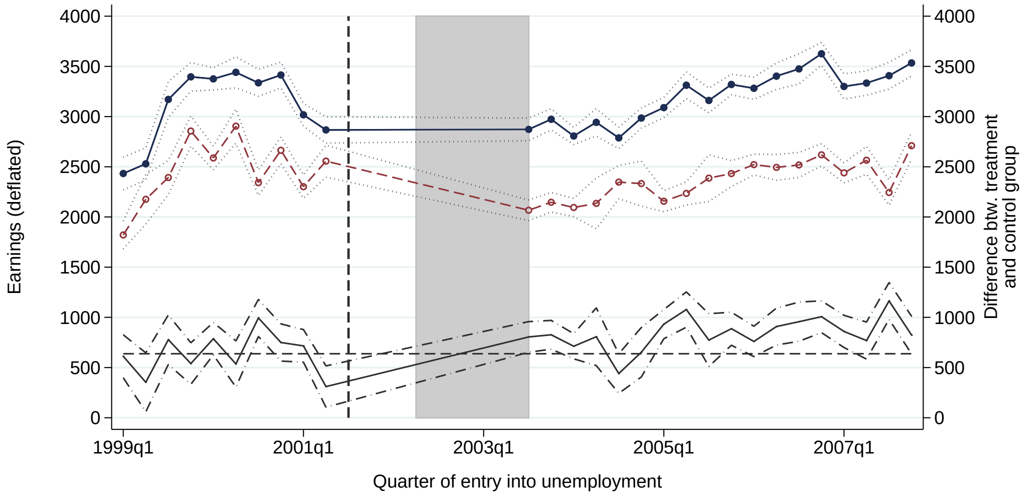

Notes: This figure shows the time trends for employment and earnings together with the 95 % confidence intervals. The solid line at the bottom shows the difference between treatment and control group together with the 95 % confidence interval. The dashed vertical line at 2001q3 depicts the first possible date for which treatment effects are detectable. The dashed horizontal line shows the mean difference in employment and earnings respectively between the treated and the control groups. Shaded area indicates that no data is available for that time period (inflow between July 2002 and June 2003 was omitted from the analysis).
Source: Own estimations based on merged UIR-SSA database.

The third test we implement checks for endogenous entry into unemployment, i.e. if the treated enter unemployment less frequently because they expect a lower benefit duration, this assumption would be violated. Figure 1.4 shows the inflows into unemployment for the treatment and control groups. The left hand axis measures the number of unemployment registrations per quarter. The right hand axis measures the inflow ratio between treatment and control group. If there was endogenous entry into unemployment, we would expect a drop in the number of registrations in the treatment group after the reform relative to the control group.

Graphical evidence however indicates that the inflow ratio does not drop after the reform, but is relatively stable over time.

Figure 1.4: Unemployment inflows (number of registrations)

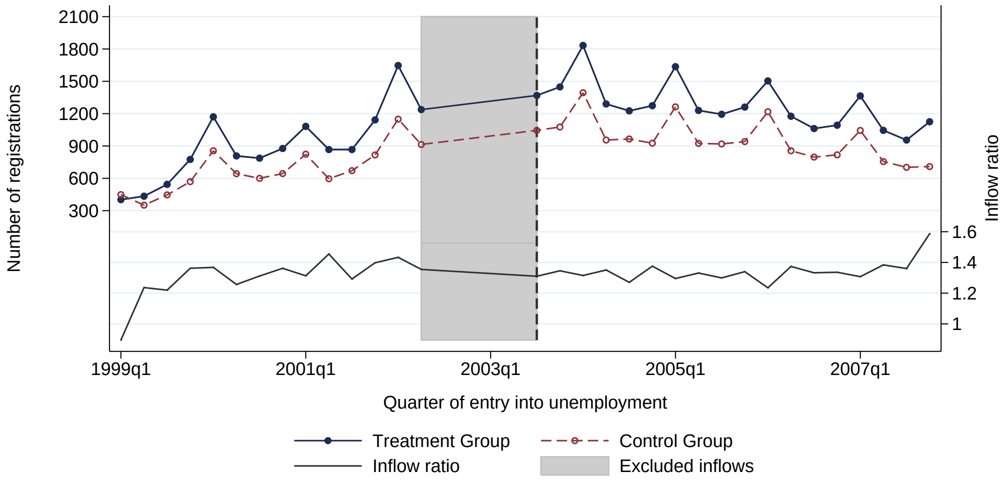

Notes: This figure shows the time trends of the unemployment inflows for treatment and control group. On the right hand axis, the solid line at the bottom shows the inflow ratio between treatment and control group. The vertical line in 2003q3 depicts the date of the reform.
Source: Own estimations based on merged UIR-SSA database.

Table 1.3 presents a formal test of stability of the inflow. It presents a regression of the treatment dummy D_i, the interaction term D_iA_c and a set of quarterly time dummies on the logarithm of the number of registrations per quarter. The reform does not significantly affect the inflow into unemployment in the treated group. This confirms that the reform did not affect the likelihood of entering unemployment. We conclude that the key assumption of parallel trends is likely to be satisfied in the current context.

Table 1.3: DiD estimates for unemployment inflows

	Log(# of registrations)
D_iA_c	0.033
	(0.034)
D_i	0.260***
	(0.032)
Time Fixed Effects	Yes
Observations	64
R-squared	0.986

Notes: The table shows the difference-in-differences estimates for the logarithm of the number of registrations. The regression includes quarterly time dummies. Robust standard errors in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

1.5 Results

This section discusses the estimation results. Subsection 1.5.1 presents graphical evidence, subsection 1.5.2 presents the main estimation, subsection 1.5.3 discusses some sensitivity estimations, and subsection 1.5.4 analyses the issue of heterogeneity in treatment effects. Subsection 1.5.5 relates our results to the existing literature on job-match quality.

1.5.1 Descriptive evidence

Figure 1.5 shows the structure of the data. We distinguish five periods: τ_0 is the period before unemployment start, i.e. 50 to 1 months before unemployment start. τ_1 marks the period 1 to 12 months after unemployment start. In this period, treatment and control group are both entitled to benefits. τ_2 identifies the period 13 to 17 months after unemployment start, where treated and untreated alike are still entitled to unemployment benefits. In this period, anticipation effects start to play a role, because unemployment benefits of the treated will run out soon. τ_3 is the period 18 to 24 months after unemployment start. This is the period in which treated are directly affected by the reform, while untreated are still eligible for benefits. This period captures the direct effect of the reduced PBD. The effect on benefit receipt will be negative and largely mechanic since the reform removes unemployment benefit payments during that period.[19] The effects on employment and earnings will show endogenous responses to the removal of benefits during period τ_3. Finally, τ_4 captures the period 25 to 50 months after unemployment start and allows to identify medium-run effects of the PBD. Period τ_4 is our primary focus since all job seekers have exhausted their framework period after two years. This period allows detecting effects of PBD reductions on medium-run earnings and employment.

Figure 1.5: Data structure

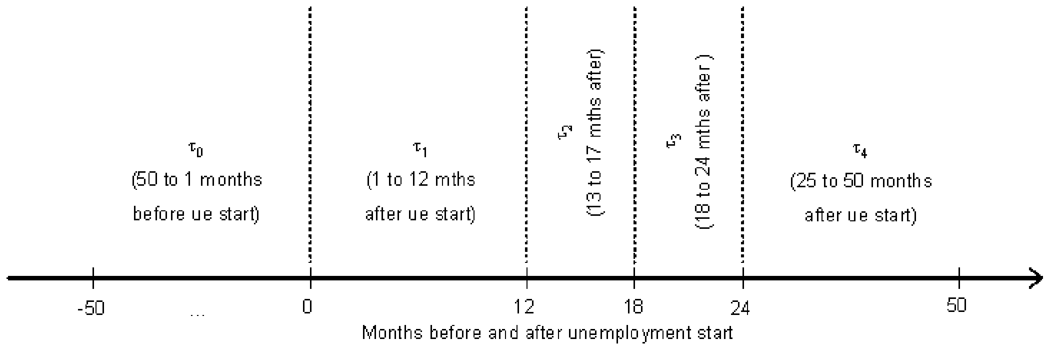

Notes: This figure shows the data structure with its division into τ_0 to τ_4.

There are three issues with this data structure: First, we cannot observe the full history of 50 months after the beginning of unemployment for spells starting after November 2006 since

[19]Some job seekers will keep eligibility to benefits even during period τ_3. These are the job seekers who re-enter after a short employment spell or job seekers who re-establish eligibility to unemployment benefits.

our observation period ends in December 2010 (13 % of all spells). This lack of observation window should, however, not impair our identification strategy, because both treated and untreated groups are affected by this gradual sample reduction in the same way. Second, due to the treatment assignment which is based on age at unemployment start, individuals in the treatment group gradually "grow" into the control group over time. For example, an individual who is 54 years old at the start of his unemployment spell will grow into the control group at most 12 months after the start of unemployment. We therefore potentially underestimate the true effects. Third, the 2003 reform affected both benefit duration and benefit level. However, this fact is unlikely to affect our results because the change to benefit level affected a narrow income bracket earning between 3,500 CHF and 4,300 CHF, and it targeted job seekers without dependents, a minor fraction of our sample.

Unemployment benefit receipt. Figure 1.6 shows average benefit receipt of the treated (50 to 54 years old) and untreated (55 to 59 years old). Benefit receipt is a binary variable which takes the value of one if a job seeker claimed unemployment benefits in a given month. Benefit receipt is shown up to 50 months before and after unemployment start. The vertical line at time 0 identifies the start of unemployment. The vertical line at 18.5 months indicates the benefit exhaustion for the treatment group after the reform, and the vertical line at 24 months marks the old exhaustion date before the reform and the benefit exhaustion date for the control group after the reform respectively. Figure 1.6a depicts benefit receipt for individuals who registered before the policy change in July, 2003 and figure 1.6b shows the same for individuals who registered after the reform in July, 2003.

Figure 1.6: Unemployment benefit receipt before and after the reform

(a) UE start before July 1, 2003 (b) UE start after July 1, 2003

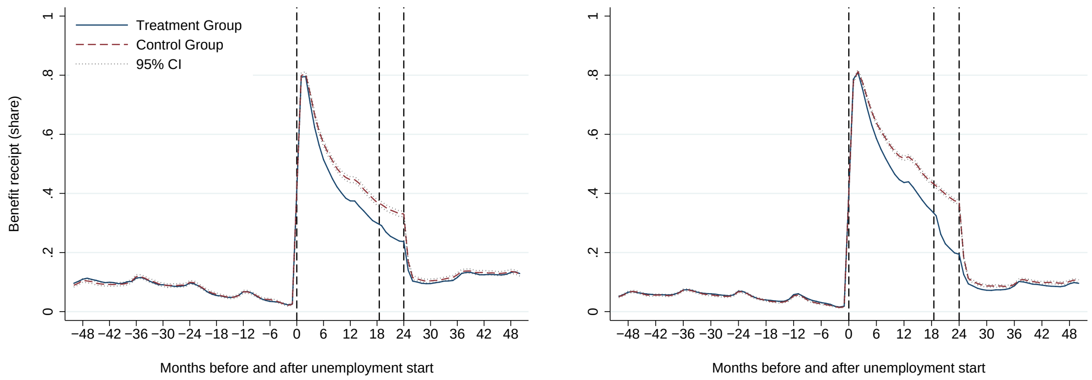

Notes: This figure shows unemployment benefit receipt for treatment and control group 50 months before and 50 months after unemployment start for spells that started before July 1, 2001 (1.6a) and for spells that started after July 1, 2003 (1.6b). The dotted lines around the benefit receipt of the control group indicate the 95 % confidence interval.
Source: Own estimations based on merged UIR-SSA database.

Benefit receipt does not differ between the treated and the untreated before the start of the unemployment spell. Unemployment benefit receipt prior to unemployment start amounts to around 6 % on average. Pre-unemployment benefit receipt is not exactly zero, because there can be spells of unemployment before the one we analyze. After registering at the PES, job seekers can claim unemployment benefits.[20] This is observed in the data by a sharp increase in average benefit receipt to around 80 % in the first month after unemployment start. The share of job seekers claiming unemployment benefits drops as time passes because job seekers gradually re-enter employment or exit the labor force through alternative pathways.

Benefit receipt of treated and untreated starts to diverge after the peak around unemployment start: Job seekers in the treatment group claim on average less unemployment benefits than job seekers in the control group. 12 months after the start of a spell there is a kink for both groups. The kink is due to the benefit exhaustion for job seekers who are exempted from the contribution requirements. They can claim a maximum of 260 days of benefit payments, which is equivalent to 12 months. For the treated group, there is another a kink after 18.5 months (equivalent to 400 days) after the beginning of unemployment: This marks the benefit exhaustion date for the treated group after the reform. A small kink is also observed for job seekers whose spells started *before* the UI policy change (Figure 1.6a). This is because the reform was applied to in-progress spells: Some job seekers in the before-treatment regime are affected by the reform even if their spells started before the 2003 reform. The kink is however much more pronounced in the data covering job seekers who enter *after* the reform, consistent with a larger treatment intensity among this group. After 24 months (equivalent to 520 days), benefits also end for the control group. Benefit receipt sharply drops, and falls back to almost its pre-unemployment level thereafter.

Figure 1.7 highlights the above observations. It shows the difference in differences between the treated and the control group before and after the policy change. In the pre-unemployment period τ_0 (50 to 1 month before unemployment start), benefit receipt has evolved in the same way for treated and control groups, the diff-in-diff estimates are close to zero and not significantly different from zero (except for the period between 7 and 5 months before unemployment start). Around 6 months after the beginning of a spell, the difference in differences starts to turn negative, reaching its minimum in the treatment period τ_3 (18 to 24 months after) where benefit receipt of treated job seekers is on average around 8 percentage points lower compared to the untreated individuals. This is the direct and purely mechanic effect of cutting PBD by 6 months for the below 55 years old job seekers. Beyond 24 months, benefit receipt is no longer affected by the reform, the difference in differences turns not significantly different from zero.

[20] Note that the unemployment start date is defined as the potential entry date for the next job. According to our sample definition, individuals thus fulfill the eligibility for daily benefit payments, conditional on being "employable". Indeed, 85 % of the sample claims unemployment benefits within 3 months after unemployment start.

Figure 1.7: Difference in differences in unemployment benefit receipt

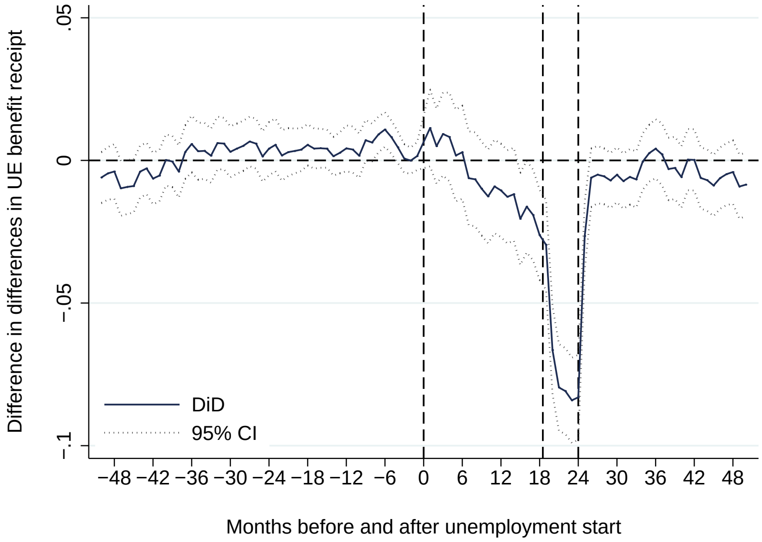

Notes: This figure shows the difference in differences for unemployment benefit receipt for the 50 months before and 50 months after unemployment start. The dotted lines around the difference in differences indicate the 95 % confidence interval.
Source: Own estimations based on merged UIR-SSA database.

Employment. Figure 1.8 replicates the above graphical analysis for the employment rate. Prior to the unemployment spell (50 to 1 months before unemployment) anywhere between 80 % and 98 % of all job seekers are employed. For both the treated and the untreated, employment already starts to fall in the last 12 to 6 months before getting unemployed. In the first month of unemployment, the employment ratio drops to zero. The unemployed start to find new jobs, and the average employment share rises again to around 60 % in the control group and to around 65 - 70 % in the treatment group. The employment patterns of the treated and control groups start to diverge only after the start of the unemployment spell: Average employment of the treated individuals increases more than the average employment of the untreated individuals before (figure 1.8a) and after (figure 1.8b) the reform. This might be due to the fact that the control group is older on average and faces more problems to find a new job. Interestingly, however, the difference in average employment between treated and control group is larger for unemployment spells that started *after* the change in PBD in July, 2003.

Figure 1.8: Employment before and after the reform

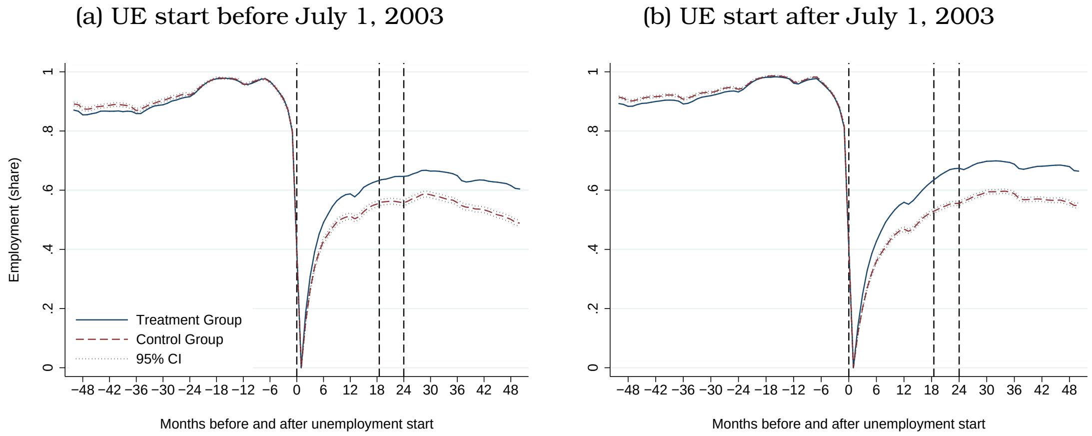

Notes: This figure shows aggregate employment for treatment and control group 50 months before and 50 months after unemployment start for spells that started before July 1, 2001 (1.8a) and for spells that started after July 1, 2003 (1.8b). The dotted lines around the employment share of the control group indicate the 95 % confidence interval.
Source: Own estimations based on merged UIR-SSA database.

Figure 1.9 confirms this observation. In the period before unemployment start, no treatment effect is detectable and the difference in differences is not statistically different from zero. The employment effect rises up to around 3.5 percentage points 20 months after entering unemployment and is statistically different from zero in the anticipation period τ_2 and in the direct treatment period τ_3. In the medium-run period τ_4, the positive employment effects gradually taper off.

Figure 1.9: Difference in differences in employment

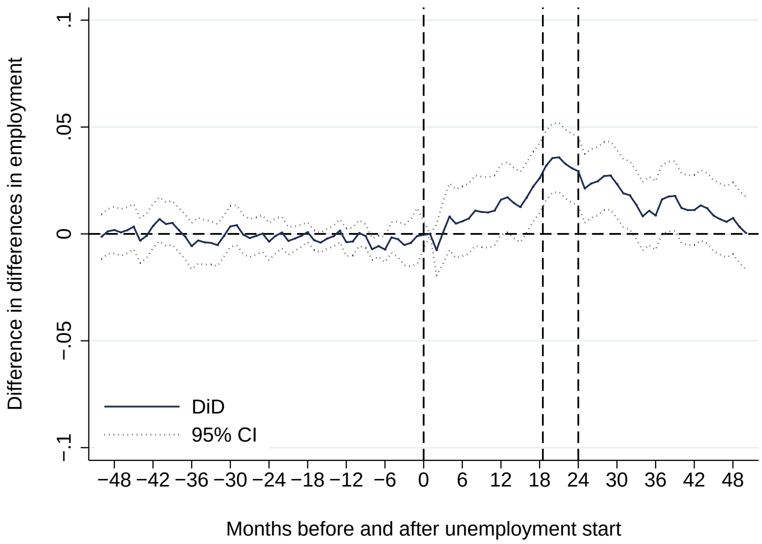

Notes: This figure shows the difference in differences for employment for the 50 months before and 50 months after entering unemployment. The dotted lines around the difference in differences indicate the 95 % confidence interval.
Source: Own estimations based on merged UIR-SSA database.

Earnings. A similar, but more volatile pattern is also observed for earnings. Figure 1.10 shows that pre-unemployment earnings are around 5,000 CHF (about 4,150 EUR), and drop to zero at unemployment start. Like the employment share, earnings rise again, but do no longer reach the pre-unemployment level, and stay at a level of between 2,500 CHF for the control group, and around 3,000 CHF for the treatment group after entering unemployment. Again, although earnings are higher for the treatment group irrespective of whether the start date of a spell was *before* (figure 1.10a) or *after* (figure 1.10b) the reform, earnings increase more for the treated than for the untreated in the *after* reform period.

Figure 1.10: Earnings before and after the reform

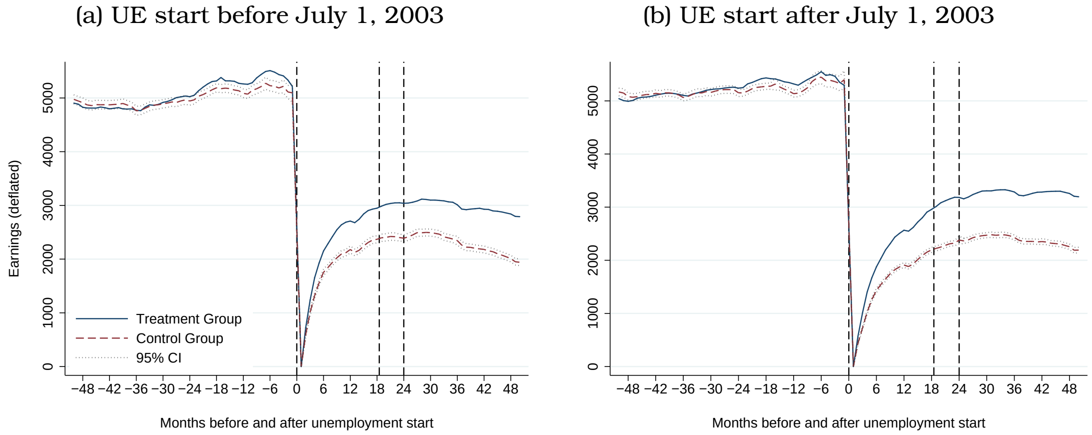

Notes: This figure shows aggregate earnings for treatment and control group 50 months before and 50 months after unemployment start for spells that started before July 1, 2001 (1.10a) and for spells that started after July 1, 2003 (1.10b). The dotted lines around earnings of the control group indicate the 95 % confidence interval.
Source: Own estimations based on merged UIR-SSA database.

The difference in differences graph for earnings completes the picture. Prior to entering unemployment, earnings are balanced across the treatment and control groups (period τ_0). The earnings difference starts to rise significantly after the beginning of a spell to around 200 CHF in the beginning of the treatment period τ_3 (18 to 24 months after unemployment start), and it remains relatively stable and significantly different from zero also in the medium-run period τ_4 (25 to 50 months after unemployment start). In contrast to the result for employment, shortened PBD therefore increases earnings permanently.

Figure 1.11: Difference in differences in earnings

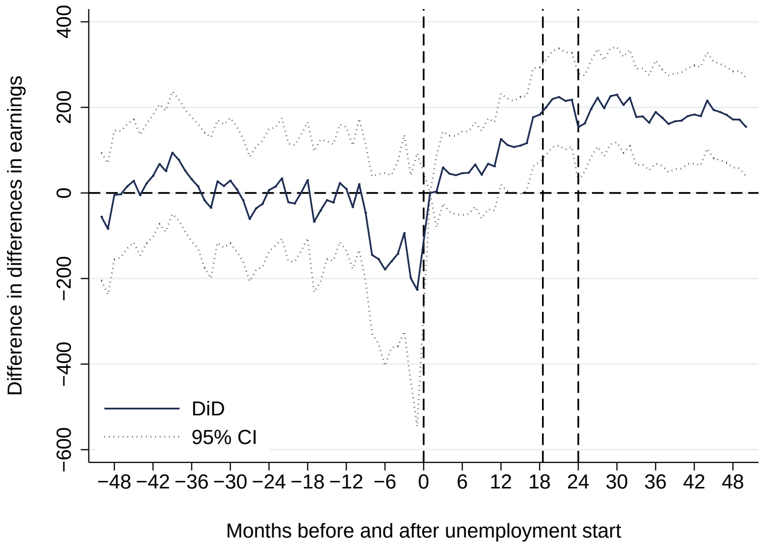

Notes: This figure shows the difference in differences for earnings for the 50 months before and 50 months after unemployment start. The dotted lines around the difference in differences indicate the 95 % confidence interval.
Source: Own estimations based on merged UIR-SSA database.

1.5.2 Main estimates

The difference-in-differences estimator is estimated by the following econometric specification

$$\begin{aligned} Y_{itc} =& \alpha_1 + \alpha_2 \tau_{2t} + \ldots + \alpha_4 \tau_{4t} + \beta_1 \tau_{1t} D_i + \ldots + \beta_4 \tau_{4t} D_i + \gamma_1 \tau_{1t} A_c + \ldots + \gamma_4 \tau_{4t} A_c + \\ & + \delta_1 \tau_{1t} D_i A_c + \ldots + \delta_4 \tau_{4t} D_i A_c + X_i' \eta + \varepsilon_{itc} \end{aligned} \tag{1.1}$$

where Y_{itc} is the outcome variable, that is unemployment benefits, employment, or earnings respectively. i is an indicator for the individual, t indicates the month after unemployment start, and c denotes calendar time. D_i is the treatment dummy which is equal to 1 if an individual belongs to the treatment group and 0 otherwise. A_c is a dummy for unemployment starts after July 1, 2003. τ_{1t} to τ_{4t} are indicators for the different periods after unemployment start, i.e. $\tau_{1t} = \mathbb{1}(1 \leq t < 13$ months$)$, $\tau_{2t} = \mathbb{1}(13 \leq t < 18$ months$)$, $\tau_{3t} = \mathbb{1}(18 \leq t < 24$ months$)$, and $\tau_{4t} = \mathbb{1}(24 \leq t \leq 50$ months$)$ respectively. δ_1 to δ_4 are the coefficients for the interaction effects $\tau_{1t} D_i A_c$ to $\tau_{4t} D_i A_c$, and identify the average treatment effect on the treated. X_i is a vector of control characteristics, such as gender, nationality, marital status (4 categories), professional status (leader/expert function versus non-leader function), and years of schooling (5 categories). As further controls we include a dummy for individuals with a high continuous work experience prior to their unemployment spell, i.e. at least 24 months of continuous employment before their unemployment start, a dummy for individuals whose previous employer is active in a R&D intensive industry, and a dummy for individuals whose task content of previous occupation was mainly cognitive, and all interactions. Finally, we also include the sums of pre-unemployment earnings and benefits, as well as the total number of months spent in employment prior to unemployment start to address the significant difference-in-differences in unemployment benefit receipt during months 7 to 5 prior to the spell we analyze (see figure 1.7). In order to adjust for potential correlation across spells and across time, standard errors of this and all following tables are clustered by person.

Table 1.4 presents the baseline results. In columns 1, 3, and 5, we estimate the treatment effects using equation (1.1) without controls. Columns 2, 4, and 6 show that the estimates of δ_1 to δ_4 remain stable and precisely estimated after the inclusion of covariates. The estimates for unemployment benefit receipt in column 2 indicate that already between 13 and 17 months after unemployment start the treated claim less unemployment benefits than the control group. The treatment effect on benefit receipt amounts to 1.7 percentage points. This treatment effect is interpreted as an anticipation effect. In the period between 18 and 24 months after unemployment start, benefit receipt is on average around 6.5 percentage points lower for the treated. δ_3 quantifies the mechanic effect of reducing benefits for the below 55 years old, but not for the above 55 years old job seekers. In the medium-run, there is no longer any significant difference between treated and untreated in terms of unemployment benefit receipt.

The estimates for employment in column 4 show that we observe an anticipation effect of 1.7 percentage points for the treatment group (13 to 17 months after unemployment start). Already before the actual reform period, the treated re-enter employment more than the untreated. The direct effect of the reform, δ_3 amounts to 3.3 percentage points. This effect is not large enough to compensate for the reduction in benefit receipt. Yet employment is also 1.5 percentage points higher for the treated in the medium-run. This positive effect compensates somewhat for lost benefit months among treated job seekers. We will explore below whether the compensation is sufficient to undo the removal of benefits.

Earnings are normalized by average earnings 3 months prior to unemployment start. In column 6, we observe a significant anticipation effect of around 2.2 percentage points. The direct effect for earnings amounts to 3.7 percentage points, and the medium-run effect stays at about the same magnitude with 3.3 percentage points. The significant medium-run coefficients δ_4 for employment and earnings show that reducing PBD does not have a purely mechanic effect, but that the positive earnings effect and to some smaller extent the employment effect persist in the medium-run.

These baseline findings suggest that the beneficial effects of a reduced human capital and skill depreciation or improvements in the non-employment stigma seem to outweigh the negative effects of reduced reservation wages. Baseline results could, however, still be spurious. We now turn to discussing the sensitivity of these baseline findings.

Table 1.4: DiD estimates for unemployment benefit receipt, employment and earnings

	Benefit receipt		Employment		Earnings	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect in the period ...						
... 1-12 mths after	-0.001	-0.002	0.006	0.007	0.009	0.008
	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
... 13-17 mths after	-0.016**	-0.017**	0.017**	0.017**	0.023**	0.022**
	(0.008)	(0.008)	(0.008)	(0.008)	(0.010)	(0.009)
... 18-24 mths after	-0.064***	-0.065***	0.032***	0.033***	0.037***	0.037***
	(0.007)	(0.007)	(0.008)	(0.008)	(0.010)	(0.009)
... 25-50 mths after	-0.005	-0.006	0.014*	0.015**	0.034***	0.033***
	(0.004)	(0.004)	(0.007)	(0.007)	(0.010)	(0.009)
Controls	No	Yes	No	Yes	No	Yes
Avg. of dep. var.	0.81	0.81	0.91	0.91	5'387.37	5'387.37
R-squared	0.22	0.22	0.05	0.09	0.03	0.18
Observations	3,073,557	3,073,557	3,073,557	3,073,557	3,073,557	3,073,557
Clusters	57,429	57,429	57,429	57,429	57,429	57,429

Notes: This table shows the baseline difference-in-differences estimates for unemployment benefit receipt (columns 1 and 2), employment (columns 3 and 4) and earnings (columns 5 and 6). Regressions with controls include also the interactions of all controls. Earnings are relative to average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

1.5.3 Sensitivity analyses

We first implement a placebo analysis to check whether trends are indeed parallel before the reform was implemented. To this end, we simulated a UI reform in July 2000 and used only inflows before July 2001. If the treatment effects in the reform periods get significant although there was no treatment in that period, this could be an indication for unequal time trends for the treated and the untreated. Table 1.5 reports difference-in-differences estimates of the treatment effects on this placebo reform. The estimated placebo treatment effects are not significant with the exception of a marginally significant employment effect 1 to 12 months after unemployment start (t-statistic of 1.68). In the actual reform period between 18 and 24 months all estimates are however non-significant. We therefore argue that the assumption of equal time trends is not violated in July 2000.[21]

[21] Note that the power of the placebo analysis to detect departures from a null effect is smaller than in the main analysis (since standard errors are two times larger). If we adopt the standard errors from the main analysis and test for significance of the effects in the placebo analysis more placebo estimates are significant. Nonetheless, the magnitude of the effects are smaller than in the main analysis.

Table 1.5: DiD estimates for a placebo reform in July 2000

	Benefit receipt	Employment	Earnings
Treatment effect in the period ...			
... 1-12 mths after	-0.004	0.021*	0.023
	(0.012)	(0.013)	(0.015)
... 13-17 mths after	0.014	0.003	0.014
	(0.015)	(0.016)	(0.020)
... 18-24 mths after	0.009	-0.003	0.005
	(0.014)	(0.016)	(0.020)
... 25-50 mths after	0.000	-0.017	-0.022
	(0.009)	(0.015)	(0.019)
Avg. of dep. var.	0.78	0.91	5'136.36
R-squared	0.17	0.08	0.18
Observations	604,550	604,550	604,550
Clusters	11,681	11,681	11,681

Notes: This table shows the baseline difference-in-differences estimates for unemployment benefit receipt (column 1), employment (column 2) and earnings (column 3) for a placebo reform in July 2000. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

A widely discussed potential concern when looking at a sample of older job seekers is that the effects of reducing PBD could be biased because of early retirement considerations and/or disability retirement as an alternative way to exit the labor force after unemployment.[22]

Table 1.6 discusses how the cut in PBD affected disability retirement pensions. A cut in PBD could affect disability pensions in mainly two ways: First, reducing PBD could amplify the adverse health effects of job-loss[23] and thereby increase disability pensions, and second, reducing PBD could induce a substitution of unemployment benefits with disability pensions. Table 1.6 shows the effects of reducing PBD on disability retirement pensions.[24] Point estimates are negative and only marginally significant. Estimates suggest that the use of disability pensions decreases between 20 and 30 % compared to average disability pension benefits before the unemployment spell. In contrast to the concerns mentioned above, these results suggest that the positive employment effects of reducing PBD also lower the need for disability pension claims.

Do reductions in PBD affect old-age pensions? Old-age retirement pensions are never observed for the treated group, and we start to observe them for the control group 26 months after unemployment start at the earliest for females, and 36 months after unemployment start for males respectively. This is because women are eligible for early retirement at the age of 62 years and and men are eligible for old-age pensions at the age of 63 years. This suggests that reductions in PBD do not affect the claiming of old-age pensions.

Nonetheless, age could be an issue because job seekers in the treatment group "grow" into the control group. This will end up reducing our estimates of the treatment effects for the

[22] Inderbitzin et al. (2012) study a regional extended benefit program in Austria and find substantial early retirement through disability insurance triggered by the unemployment benefit reform.
[23] Kuhn et al. (2009) find important health effects of job loss, particularly for men.
[24] Disability pension data is only available for a random subsample of around 35 % of job seekers.

Table 1.6: DiD estimates for disability retirement

	Disability pensions
Treatment effect in the period ...	
... 1-12 mths after	-0.206*
	(0.120)
... 13-17 mths after	-0.185
	(0.147)
... 18-24 mths after	-0.163
	(0.155)
... 25-50 mths after	-0.293*
	(0.166)
Avg. of dep. var.	93.59
R-squared	0.06
Observations	1,153,356
Clusters	21,463

Notes: This table shows the difference-in-differences estimates for disability pensions normalized by the average disability pension 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

periods τ_2 to τ_4. To address this concern, we estimate a model that excludes the oldest age cohorts of the treatment and the control group. That is, we exclude the 54 years old individuals in the treatment group, and the 59 years old individuals in the control group. Table 1.7 reports the estimates for this restricted sample. Excluding the oldest age cohorts in each group does not affect the estimates drastically: Compared to our main estimates, the treatment effects are virtually unchanged for employment and earnings, and slightly stronger for unemployment benefit receipt. Statistical significance decreases somewhat, because in the restricted sample around one fifth of all observations is lost. The overall picture however is unchanged.[25]

[25]We have also explored whether our results are sensitive to how we define the start date of the unemployment spell. Overall, results are similar to the baseline result when we use the date a job seeker registers at the employment center as the start date of her or his unemployment spell. Registration dates are, however, not ideal as unemployment start dates because job seekers need to register at the job center as soon as they are informed that their employment spell ends. This leads to a situation where the effects on outcomes in different phases of the spell get blurred since the timing is not quite correct. Results are available upon request.

Table 1.7: DiD estimates by age

	Baseline			50-53 vs. 55-58 years old		
	Benefit receipt	Employ-ment	Earnings	Benefit receipt	Employ-ment	Earnings
Treatment effect in the period . . .						
. . . 1-12 mths after	-0.002	0.007	0.008	-0.002	0.007	0.008
	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)
. . . 13-17 mths after	-0.017**	0.017**	0.022**	-0.018**	0.017**	0.023**
	(0.008)	(0.008)	(0.009)	(0.008)	(0.008)	(0.010)
. . . 18-24 mths after	-0.065***	0.033***	0.037***	-0.073***	0.030***	0.037***
	(0.007)	(0.008)	(0.009)	(0.008)	(0.008)	(0.010)
. . . 25-50 mths after	-0.006	0.015**	0.033***	-0.008*	0.012	0.034***
	(0.004)	(0.007)	(0.009)	(0.004)	(0.008)	(0.010)
Avg. of dep. var.	0.81	0.91	5’387.37	0.81	0.91	5’419.94
R-squared	0.22	0.09	0.18	0.22	0.09	0.18
Observations	3,073,557	3,073,557	3,073,557	2,532,295	2,532,295	2,532,295
Clusters	57,429	57,429	57,429	47,919	47,919	47,919

Notes: This table shows the difference-in-differences estimates for subsamples split by age. Columns 1 to 3 replicate the baseline estimates, and columns 4 to 6 include only 50 to 53, and 55 to 58 years old individuals respectively. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

1.5.4 Treatment effect heterogeneity

This section analyzes whether the effects of a reduction in potential benefit duration differ between subgroups of job seekers with different previous industry and occupation.

A first sample split discusses the role of human capital depreciation due to skill obsolescence as a possible driving force of the positive medium-run effects. We split the sample in two groups which likely differ in terms of the speed at which industry specific skills become obsolete: job seekers from industries with high R&D intensity versus job seekers from industries with low R&D intensity (see section 3 for a discussion of how we define R&D intensity of the industry). Skill depreciation is assumed to play a more important role for individuals working in fast-evolving, highly R&D intensive industries, because a job-loss disconnects the unemployed faster from rapid technological change in those industries. An shortened period of unemployment is therefore expected to be more beneficial for job seekers in highly R&D intensive industries.

Table 1.8 presents estimates for the sample split by R&D intensity of previous industry.[26] Columns 1 to 3 of these two tables reproduces the baseline estimates for the sake of comparison. Columns 4 to 6 report estimates for job seekers coming from above median R&D intensive industries, and columns 7 to 9 for job seekers from industries with below median R&D intensity respectively. For both subsamples we observe a negative effect on benefit receipt in the reform period from 18 to 24 months after unemployment start. The effects on earnings and employment, however, differ considerably between the two groups. The effects are much stronger

[26]All sub-group regressions include the same list of control variables as the baseline regression analysis.

for job seekers who left R&D intensive industries than for job seekers who left industries with little expenditure on R&D. In the former group we observe strong and significant anticipation effects in the period from 13 to 17 months after unemployment start. The direct effects of the reform (18 to 24 months after unemployment start) lead to a 4.6 percentage points increase in employment and to a earnings effect of around 6.8 percentage points. The effects persist also in the medium-run (25 to 50 months after unemployment start): Reducing PBD by 6 months leads to a 2.2 percentage points increase in terms of employment and it boosts earnings by 5.2 percentage points in the medium-run. In contrast, treatment effects for job seekers leaving industries with low R&D intensity are mostly absent except for an employment effect in the period from 18 to 24 months after unemployment start directly induced by the cut in benefits over that period. Other point estimates are close to zero for employment, and even negative for earnings, but none of them are statistically significant.[27]

A second subgroup analysis discusses the importance of job-specific human capital depreciation due to lack of use of skills (atrophy). In this analysis, we assess whether the task content of the previous occupation matters for the medium-run effect of benefit reductions on earnings and employment. As above, we split the sample into two subgroups: One of them contains job seekers who worked in occupations that use primarily cognitive skills, and the other subgroup contains job seekers who worked in occupations that use mainly manual skills (see section 3 for the definition of these two groups). We expect that skill depreciation differs between the two groups. However, whether skill depreciation would be stronger in mainly cognitive or mainly manual occupations is not clear. One argument holds that cognitive skills depreciate faster because extended periods of unemployment generate adverse effects on mental health. This would indicate that occupations with cognitive tasks suffer more from job loss than occupations with manual tasks. Conversely, one might also argue that occupations with cognitive skills might be insulated from atrophy because they are used to maintain those skills better than occupations with manual skills.

Table 1.9 shows the estimates for this sample split by task content of previous occupation. Columns 1 to 3 repeat the baseline estimates, columns 4 to 6 contains the subgroups of job seekers with mainly cognitive skills and columns 7 to 9 the subgroups of job seekers with mainly manual skills. The reform led to a decrease in benefit receipt of around 6.5 percentage points for both groups in the reform period from 18 to 24 months after unemployment start (and a small anticipation effect for job seekers with manual skills). Employment effects are also quite similar between the two groups in the first 24 months (columns 5 and 8). Reductions in PBD tend to increase employment in the reform period (and also a bit earlier for job seekers with manual skills.) Employment effects over the medium-run period differ strongly between the two groups. Employment is significantly higher for job seekers with manual skills 25 to 50

[27] Job seekers in industries with high R&D expenditure have higher mean earnings than job seekers in industries with low R&D expenditure. Yet the difference in the effects of PBD is not simply related to the difference in earnings. When we split the sample by previous earnings, we find positive medium-run effects for both subsamples (see table 1.A3).

Table 1.8: DiD estimates by R&D intensity of previous industry

	Baseline			High R&D intensity			Low R&D intensity		
	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings
Treatment effect in the period ...									
... 1-12 mths after	-0.002	0.007	0.008	-0.011	0.020**	0.024***	0.006	-0.004	-0.011
	(0.006)	(0.006)	(0.007)	(0.009)	(0.008)	(0.009)	(0.009)	(0.009)	(0.010)
... 13-17 mths after	-0.017**	0.017**	0.022**	-0.030***	0.036***	0.055***	-0.006	0.002	-0.017
	(0.008)	(0.008)	(0.009)	(0.011)	(0.011)	(0.012)	(0.011)	(0.011)	(0.013)
... 18-24 mths after	-0.065***	0.033***	0.037***	-0.077***	0.046***	0.068***	-0.056***	0.021**	-0.001
	(0.007)	(0.008)	(0.009)	(0.010)	(0.011)	(0.012)	(0.010)	(0.011)	(0.014)
... 25-50 mths after	-0.006	0.015**	0.033***	-0.007	0.022**	0.052***	-0.004	0.008	0.011
	(0.004)	(0.007)	(0.009)	(0.006)	(0.010)	(0.011)	(0.006)	(0.010)	(0.013)
Avg. of dep. var.	0.81	0.91	5'387.37	0.82	0.92	6'168.21	0.79	0.90	4'642.51
R-squared	0.22	0.09	0.18	0.25	0.10	0.19	0.20	0.08	0.17
Obs.	3'073'557	3'073'557	3'073'557	1'503'425	1'503'425	1'503'425	1'570'132	1'570'132	1'570'132
Clusters	57'429	57'429	57'429	28'927	28'927	28'927	29'473	29'473	29'473

Notes: This table shows the difference-in-differences estimates for subsamples split by innovative pace of industries (as measured by R&D intensity of industries) for unemployment benefit receipt, employment, and earnings respectively. Columns 1 to 3 replicate the baseline estimates, columns 4 to 6 include industries with an R&D intensity above the median, and columns 7 to 9 include industries with an R&D intensity below the median. R&D intensity of industries is measured as an average of R&D in percentage of GDP for the neighboring countries Germany, Austria, France and Italy over the years 2005 to 2008. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** $P<0.01$ ** $P<0.05$ * $P<0.1$.
Source: Own estimations based on merged UIR-SSA database.

months after the start of the spell whereas employment is not affected among job seekers with cognitive skills. Earnings patterns also differ strikingly between the two subgroups (columns 6 and 9). Job seekers from occupations with largely manual skill content enjoy significantly higher earnings already from the start. The effect is small (1.6 percent) immediately after the spell starts but it builds up to a sizable 5.1 percent differential in the medium-run period. There is also a positive earnings effects for job seekers in manual occupations but this effect is concentrated in the reform period and comparably small (2.8 percent). Taken at face value, these results suggests that skill depreciation affects occupations with manual skill content more strongly than occupations with cognitive skill content. This result is consistent with direct evidence on skill atrophy. Li (2013), for instance, finds strong human capital depreciation for some manual occupations such as sales and production workers, or, conversely, human capital appreciation for cognitive occupations like education professionals. Görlich and de Grip (2009) find that skill depreciation rates are higher for low-skilled workers than for high-skilled workers.[28]

Finally, we take a closer look at post-unemployment earnings. Reduced PBD can increase earnings either by increasing employment, or by increasing earnings of employed individuals, or both. Table 1.10 presents estimates of earnings effects conditional on employment. These results need to be interpreted with caution because earnings of employed individuals are only observed for job seekers who have found employment. The causal effects of PBD may therefore be masked by selection into employment effects.[29]

Results present a clear picture. Reductions in potential benefit duration increase earnings of employed individuals from the first month of unemployment onwards (Table 1.10 column 1). Earnings of treated job seekers are about 2 percent higher than they would have been without the reduction of potential benefits. Earnings are particularly positively affected for job seekers leaving industries that spend a lot on R&D (column 2) – earnings gains are between 4 to 5 percent of post-unemployment earnings. In contrast, job seekers who leave industries with low R&D expenditure tend to have lower earnings when employed, particularly so in the period 18 to 24 months after the unemployment spell started.[30] Interestingly, results by task-content of the occupation display positive point estimates, both for occupations with mainly cognitive task content as well as occupations with mainly manual task content. Point estimates are on the order of 2 percent of post-unemployment earnings, and significant in 4 out of 8 cases. In sum, results in table 1.10 suggest that reductions in PBD have a positive effect on medium-run earnings because of both, somewhat higher employment and higher earnings while employed.

[28]Table 1.A2 presents subgroup estimations split by gender. We find that the effects are very similar among male and female job seekers. We have also investigated results by the extent of routine or non-routine tasks involved. These estimates (not shown) indicate that occupations with a routine tasks have higher medium-run earnings and employment with reduced PBD. No such effect is present for occupations involving non-routine tasks.

[29]Note, however, our results give a lower bound on the earnings effects if job seekers select into employment based on ability or earnings potential. Reductions in PBD lead job seekers to accept jobs earlier, so more job seekers are observed with reduced PBD than with extended PBD.

[30]Note that this negative effect has been documented also by Caliendo et al. (2013a).

Table 1.9: DiD estimates by task content of previous occupation

	Baseline			Cognitive Tasks			Manual Tasks		
	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings
Treatment effect in the period ...									
... 1-12 mths after	-0.002	0.007	0.008	0.003	0.008	0.005	-0.008	0.009	0.016*
	(0.006)	(0.006)	(0.007)	(0.009)	(0.008)	(0.009)	(0.009)	(0.008)	(0.009)
... 13-17 mths after	-0.017**	0.017**	0.022**	-0.010	0.020*	0.015	-0.024**	0.016	0.033***
	(0.008)	(0.008)	(0.009)	(0.011)	(0.011)	(0.013)	(0.011)	(0.011)	(0.012)
... 18-24 mths after	-0.065***	0.033***	0.037***	-0.067***	0.035***	0.028**	-0.064***	0.030***	0.048***
	(0.007)	(0.008)	(0.009)	(0.010)	(0.011)	(0.013)	(0.010)	(0.011)	(0.012)
... 25-50 mths after	-0.006	0.015**	0.033***	-0.008	0.007	0.019	-0.003	0.021**	0.051***
	(0.004)	(0.007)	(0.009)	(0.006)	(0.010)	(0.012)	(0.006)	(0.010)	(0.012)
Avg. of dep. var.	0.81	0.91	5'387.37	0.82	0.91	6'270.97	0.79	0.91	4'551.34
R-squared	0.22	0.09	0.18	0.24	0.10	0.18	0.20	0.09	0.18
Obs.	3'073'557	3'073'557	3'073'557	1'494'298	1'494'298	1'494'298	1'579'209	1'579'209	1'579'209
Clusters	57'429	57'429	57'429	28'631	28'631	28'631	29'955	29'955	29'955

Notes: This table shows the difference-in-differences estimates for subsamples split by the task content of the previous occupation for unemployment benefit receipt, employment, and earnings respectively. Columns 1 to 3 replicate the baseline estimates, columns 4 to 6 include occupations with primarily cognitive task content, and columns 7 to 9 include occupations with mainly manual task content. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

Table 1.10: DiD estimates for employed individuals

	All	R&D intensity		Task content	
		High	*Low*	*Cognitive*	*Manual*
Treatment effect in the period ...					
... 1-12 mths after	0.019**	0.042***	-0.006	0.014	0.028**
	(0.009)	(0.013)	(0.012)	(0.013)	(0.011)
... 13-17 mths after	0.023**	0.056***	-0.016	0.025*	0.022*
	(0.010)	(0.014)	(0.014)	(0.014)	(0.013)
... 18-24 mths after	0.014	0.048***	-0.025*	0.018	0.011
	(0.010)	(0.013)	(0.013)	(0.014)	(0.012)
... 25-50 mths after	0.022**	0.045***	-0.007	0.025**	0.018
	(0.009)	(0.012)	(0.013)	(0.012)	(0.012)
Avg. of dep. var.	5'916.48	6'705.41	5'148.71	6'861.50	5'015.72
R-squared	0.32	0.32	0.31	0.32	0.31
Obs.	1'702'374	822'642	879'732	861'405	840'957
Clusters	48'188	24'139	24'918	24'682	24'549

Notes: This table illustrates the difference-in-differences estimates for the effect of reduced unemployment benefit duration on earnings for individuals conditional on employment. The effects are shown for the baseline specification in column 1, for the sample splits by R&D intensity in columns 2 and 3, and for the sample splits by task content of previous occupation in columns 4 and 5. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

1.5.5 Relation to existing literature

The existing literature mainly focuses on outcomes that capture job-match quality for job seekers who find jobs after their unemployment spell. We now discuss what happens if we analyze the effects of the PBD reduction on these direct measures of job-match quality.

Panel A of table 1.11 reports the effects of reducing the PBD on unemployment duration. Unemployment duration is defined as the number of months spent in unemployment until the next job. If there is no next job observed in the data, unemployment duration is right censored by the last observed date. On average job seekers spent around 14.6 months in unemployment.[31] Reducing the PBD by 6 months lowers the time spent in unemployment for the treatment group by about 1 month for the entire sample. The effect of benefit reductions is considerably larger in the subgroup of job seekers from R&D intensive industries: treated job seekers leave unemployment about 1.7 months earlier than they would have without a cut in benefits. Interestingly, there is a negative yet insignificant effect on the duration of unemployment of job seekers in industries with low R&D intensity. The effect is less strong on job seekers who leave R&D industries because they do not exhibit an anticipation effect in the period 13 to 17 months after the start of the unemployment spell (see table 1.8). When looking at the subgroup analysis that splits job seekers into cognitive and manual skilled groups respec-

[31]The average duration spent in unemployment for individuals who actually found a job within the observed time period is almost cut by half, with around 8.4 months.

tively, we find that treated job seekers with mainly cognitive skills are on average 0.88 months less unemployed than their counterparts in the control group. Job seekers whose occupation predominantly require manual skills leave unemployment about 1.16 months earlier than they would without the benefit reduction. Effects are similar across the two groups of job seekers.

Panel B presents the difference-in-differences estimates for post-unemployment earnings measured in the second month after reemployment. This analysis is based on all spells where job seekers left unemployment and stayed in their job for at least two months.[32] Results indicate no significant effect in the overall sample. Interestingly, point estimates are positive for the subgroups with large skill depreciation. The effects are significant for job seekers in occupations with high manual task content, and positive but insignificant for job seekers leaving R&D intensive industries. Point estimates are negative for job seekers with low skill depreciation. The effect is significant for job seekers leaving industries with low R&D intensity and insignificant for job seekers in occupations with cognitive task content. These results reinforce the interpretation that reductions in PBD may improve the lot of job seekers who face rapid skill depreciation or skill obsolescence. In contrast, reducing PBD tends to hurt job seekers who do not face human capital depreciation.

[32]We focus on earnings in the second month after reemployment because the first month after reemployment is the month when job seekers leave unemployment. If a job seeker starts her job in the middle of this month, earnings do not reflect full-time monthly earnings. (We do not observe number of days worked on the job so we can not adjust for this.)

Table 1.11: Effects on unemployment duration, subsequent earnings, and job loss

	All	R&D intensity		Task content	
		High	*Low*	*Cognitive*	*Manual*
A. Unemployment duration (months)					
D_iA_c	-1.022***	-1.737***	-0.374	-0.882**	-1.160***
	(0.287)	(0.402)	(0.410)	(0.391)	(0.418)
Avg. of dep. var.	14.61	15.35	13.90	14.22	14.98
R-squared	0.08	0.09	0.08	0.07	0.10
Obs.	3,073,557	1,503,425	1,570,132	1,494,298	1,579,259
Clusters	57,429	28,927	29,473	28,631	29,956
B. Monthly earnings (CHF)					
D_iA_c	-21.121	115.620	-138.611**	-131.590	104.986*
	(52.225)	(81.413)	(65.243)	(84.173)	(60.302)
Avg. of dep. var.	3'777.91	4'041.25	3'529.98	4'071.17	3'484.18
R-squared	0.28	0.30	0.24	0.29	0.26
Obs.	2,399,823	1,163,748	1,236,075	1,200,875	1,198,948
Clusters	44,763	22,395	23,123	23,008	22,672
C. Job loss within 12 months					
D_iA_c	0.003	0.010	-0.007	-0.004	0.010
	(0.009)	(0.013)	(0.013)	(0.013)	(0.013)
Avg. of dep. var.	0.41	0.38	0.44	0.37	0.45
R-squared	0.05	0.04	0.06	0.04	0.06
Obs.	2,555,279	1,238,717	1,316,562	1,272,023	1,283,256
Clusters	47,507	23,777	24,562	24,326	24,200
D. Job loss within 24 months					
D_iA_c	-0.004	-0.001	-0.009	0.000	-0.006
	(0.009)	(0.013)	(0.013)	(0.013)	(0.013)
Avg. of dep. var.	0.53	0.51	0.56	0.50	0.57
R-squared	0.05	0.04	0.06	0.04	0.06
Obs.	2,555,279	1,238,717	1,316,562	1,272,023	1,283,256
Clusters	47,507	23,777	24,562	24,326	24,200

Notes: This table shows difference-in-differences estimates for unemployment duration and a number of job-match quality measures together with their means. Panel A shows the estimates for number of months spent in unemployment. Panel B illustrates the estimates for earnings in the second month of re-employment. Panels C and D focus on job loss within 12 and 24 months.Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

Table 1.11 also looks at the job loss probabilities 12 (Panel C) and 24 months (Panel D) after reemployment respectively. Estimations include only observations which we observe for at least 24 months after reemployment.[33] The probability of losing the job after reemployment varies between roughly 37 % for a job loss within 12 months and 57 % for a job loss within 24 months after reemployment. Overall, reducing the PBD does not affect the duration of

[33]A total of 52,795 of all job seekers leave unemployment for a job. Out of those, 47,507 job seekers start their new jobs at least two years before the end of the observation period. We exclude about 10 % of re-employed job seekers whose employment durations are observed for less than two years.

employment spells.

All in all, results in table 1.11 confirm the general pattern of findings of the existing literature, which finds only small or no effects of UI policy changes on job-match quality. The findings however also support the view that the beneficial effects of reduced human capital depreciation and improvements in non-employment stigma outweigh the negative effects of reduced reservation wages, leading to positive overall effects on earnings and employment in the medium-run.

1.5.6 Effects on income

We have documented that a reduction in PBD reduces benefit receipt but increases employment and earnings of job seekers. We now assess the effects on income. Income is the sum of labor earnings and unemployment benefits, i.e. income from social assistance or other transfer programs is not counted. Effects on total income provide information on how disposable income – a key component of individual welfare – is affected by reductions of PBD.[34]

Results in table 1.12 indicate that reductions in PBD do not lower total income, not even in the period when benefits are cut (18 to 24 months after the start of the unemployment spell). This is surprising considering that benefit receipt goes down considerably in the reform period. Yet loss of benefits is more than compensated by increased earnings. The effect of PBD once the framework period has ended is even positive with income increasing by 2.9 percent. The average effect of reducing PBD on income remains positive and amounts to 1.7 percent of income. Thus, reducing PBD tends to increase income *on average*.

How does this look in subsamples? Results by R&D intensity of the previous industry indicate that reductions in PBD lower income during the period when benefits are withheld only for job seekers in low R&D industries (income drops by 3.9 percent). The effect is positive and even significant for job seekers in high R&D industries. Income of job seekers leaving industries with high R&D expenditure increases considerably, by 4.7 percent, in the medium-run period. No corresponding effect can be detected for job seekers leaving low R&D industries. As a result, job seekers from high R&D industries on net enjoy a 3.6 percent higher income in the system with reduced PBD on average. There is no effect of reduced PBD on average income for job seekers from low R&D industries.

Results by task content of the previous occupation also indicate important differences. Interestingly, none of the two occupation groups suffers a significant reduction in income in the reform period when benefits are removed. Job seekers with manual occupations benefit from a significant increase in income in the medium-run (income increases by 4.8 percent); there is no corresponding effect for job seekers with cognitive tasks. On net, job seekers with manual occupations have a 2.6 percentage points higher net income in a system with lower PBD. Job seekers whose occupations entail mainly cognitive tasks do not fare worse in a system with reduced PBD.

[34]Note, however, that assessing individual welfare would imply accounting for a number of additional aspects (leisure, discounting, general equilibrium effects). Assessing these aspects is beyond the scope of the current analysis.

Table 1.12: DiD estimates for total income

	All	R&D intensity		Task Content	
		High	*Low*	*Cognitive*	*Manual*
Treatment effect in the period ...					
... 1-12 mths after	0.013*	0.025***	-0.003	0.014	0.011
	(0.007)	(0.009)	(0.010)	(0.009)	(0.010)
... 13-17 mths after	0.012	0.037***	-0.019	0.012	0.010
	(0.008)	(0.011)	(0.013)	(0.011)	(0.012)
... 18-24 mths after	-0.007	0.018*	-0.039***	-0.006	-0.010
	(0.008)	(0.011)	(0.013)	(0.012)	(0.011)
... 25-50 mths after	0.029***	0.047***	0.010	0.015	0.048***
	(0.009)	(0.011)	(0.013)	(0.012)	(0.012)
Total	0.017**	0.036***	-0.004	0.010	0.026***
	(0.007)	(0.009)	(0.010)	(0.009)	(0.009)
Avg. of dep. var.	5,453.29	6,238.86	4,703.92	6,338.44	4,615.68
R-squared	0.24	0.25	0.23	0.24	0.24
Obs.	3,073,557	1,503,425	1,570,132	1,494,298	1,579,209
Clusters	57,429	28,927	29,473	28,631	29,955

Notes: This table illustrates the difference-in-differences estimates for the effect of reduced unemployment benefit duration on total income. The effects are shown for the baseline specification in column 1, for the sample splits by R&D intensity in columns 2 and 3, and for the sample splits by task content of previous occupation in columns 4 and 5. Income is the sum of labor earnings and unemployment benefits relative to its mean three months before the unemployment spell starts. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

1.6 Conclusions

We discuss the effects of shortening potential benefit duration (PBD) for job seekers aged 50 to 54 years. Shortening PBD pushes job seekers into jobs during the period when benefit payments are cut. But these jobs may be of lower quality than the jobs that job seekers would have found with longer PBD. Conversely, inciting job seekers to leave unemployment more quickly help them find jobs before their human capital depreciates or before they acquire the stigma of long-term unemployment.

We find strong evidence for the job push effect. Interestingly, we also find that the initial push into jobs carries longer lasting benefits. Job seekers who find employment more quickly because of a reduction in PBD tend to earn more not only during the period when benefits are removed but up to 2 years later on. The medium-run benefits are especially strong for job seekers who left R&D intensive industries and basically absent for job seekers in low R&D intensive industries. We find similar discrepancies for job seekers whose occupation necessitate manual skills and no medium-run benefits for job seekers with occupations rich in cognitive skills. Moreover, when we assess the effects on total income, we find that reduced PBD raises total income of job seekers who enjoy medium-run benefits and has no effect on income of job seekers where such medium-run effects are absent.

The evidence we find is consistent with unemployment insurance having an important role in human capital depreciation, especially for subgroups that face rapid skill depreciation. Reductions in PBD can improve earnings and employment of job seekers in these subgroups whereas extensions of PBD could probably also lead to reductions in labor market outcomes. Should benefit duration be reduced across the board? We believe this conclusion is premature for a number of reasons. First, we have seen that the effects of reducing PBD differ by task content and previous industry. Second, reducing PBD carries a cost in terms of reduced protection against economic shocks. This cost should be weighed against the potential benefits we have isolated. Third, reducing PBD to zero will, arguably, have more detrimental effects than removing 6 months out of 24 months. Fourth, human capital depreciation and long-term unemployment stigma might be more important for old job seekers than for younger ones. These issues should be explored in further research.

1.A Additional tables

Table 1.A1: DiD estimates for unemployment benefits, employment and earnings

	Benefit receipt		Employment		Earnings	
	(1)	(2)	(3)	(4)	(5)	(6)
$\tau_1 D_i A_c$ (1-12 mths after)	-0.001	-0.002	0.006	0.007	0.009	0.008
	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
$\tau_2 D_i A_c$ (13-17 mths after)	-0.016**	-0.017**	0.017**	0.017**	0.023**	0.022**
	(0.008)	(0.008)	(0.008)	(0.008)	(0.010)	(0.009)
$\tau_3 D_i A_c$ (18-24 mths after)	-0.064***	-0.065***	0.032***	0.033***	0.037***	0.037***
	(0.007)	(0.007)	(0.008)	(0.008)	(0.010)	(0.009)
$\tau_4 D_i A_c$ (25-50 mths after)	-0.005	-0.006	0.014*	0.015**	0.034***	0.033***
	(0.004)	(0.004)	(0.007)	(0.007)	(0.010)	(0.009)
τ_2	-0.175***	-0.175***	0.149***	0.149***	0.131***	0.131***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)
τ_3	-0.244***	-0.244***	0.183***	0.183***	0.160***	0.160***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
τ_4	-0.466***	-0.466***	0.170***	0.170***	0.136***	0.136***
	(0.005)	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)
$\tau_1 D_i$	-0.048***	-0.049***	0.057***	0.061***	0.070***	0.070***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$\tau_2 D_i$	-0.075***	-0.075***	0.078***	0.083***	0.108***	0.108***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.008)	(0.007)
$\tau_3 D_i$	-0.084***	-0.085***	0.081***	0.085***	0.115***	0.115***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.008)	(0.007)
$\tau_4 D_i$	-0.007**	-0.008**	0.095***	0.099***	0.131***	0.131***
	(0.003)	(0.003)	(0.006)	(0.006)	(0.008)	(0.007)
$\tau_1 A_c$	0.055***	0.053***	-0.054***	-0.060***	-0.046***	-0.046***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$\tau_2 A_c$	0.074***	0.072***	-0.038***	-0.044***	-0.042***	-0.041***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
$\tau_3 A_c$	0.054***	0.052***	-0.015**	-0.022***	-0.022***	-0.021***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.008)	(0.007)
$\tau_4 A_c$	-0.024***	-0.026***	0.029***	0.023***	0.020***	0.021***
	(0.003)	(0.003)	(0.006)	(0.006)	(0.007)	(0.006)
Sum of pre-reg. benefits		0.000***		0.000***		0.000***
		(0.000)		(0.000)		(0.000)
Sum of pre-reg. earnings		-0.000***		0.000***		0.000***
		(0.000)		(0.000)		(0.000)
Mths employed before reg.		-0.000		0.007***		0.002***
		(0.000)		(0.000)		(0.000)
$\geq$ 24 mths of work exp.		0.016***		-0.094***		-0.079***
		(0.002)		(0.004)		(0.005)
R&D intensive industry		0.018***		-0.030***		-0.011***
		(0.002)		(0.003)		(0.004)
Cognitive task		0.022***		-0.001		0.008**
		(0.002)		(0.003)		(0.004)
Female		0.001		0.009***		-0.066***
		(0.002)		(0.003)		(0.006)
Swiss		-0.022***		0.095***		0.066***
		(0.002)		(0.004)		(0.005)
Leader position		0.002		0.046***		0.081***
		(0.003)		(0.004)		(0.005)
Marital status (reference group are singles)						
Married		-0.025***		0.027***		0.033***
		(0.003)		(0.005)		(0.005)
Widowed		-0.020***		0.006		-0.017
		(0.007)		(0.011)		(0.012)
Divorced		-0.015***		0.027***		0.031***
		(0.003)		(0.005)		(0.006)
Education (reference group is "8-9 years of schooling")						
$\leq$ 7 years		0.010		-0.067***		-0.034**

Table 1.A1 – continued

	Benefit receipt		Employment		Earnings	
	(1)	(2)	(3)	(4)	(5)	(6)
		(0.009)		(0.014)		(0.013)
10-11 years		-0.002		0.013		0.015*
		(0.005)		(0.008)		(0.008)
12-13 years		-0.018***		0.058***		0.067***
		(0.003)		(0.005)		(0.006)
$\geq$ 14 years		-0.017***		0.051***		0.165***
		(0.006)		(0.009)		(0.013)
Other		-0.044***		0.050***		0.086***
		(0.003)		(0.005)		(0.006)
Avg. of dep. var.	0.81	0.81	0.91	0.91	5'387.37	5'387.37
R-squared	0.22	0.22	0.05	0.09	0.03	0.18
Obs.	3'073'557	3'073'557	3'073'557	3'073'557	3'073'557	3'073'557
Clusters	57'429	57'429	57'429	57'429	57'429	57'429

Notes: This table shows the baseline difference-in-differences estimates for unemployment benefit receipt (columns 1 and 2), employment (columns 3 and 4) and earnings (columns 5 and 6). Regressions with controls include also the interactions of all controls. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** $P<0.01$ ** $P<0.05$ * $P<0.1$.
Source: Own estimations based on merged UIR-SSA database.

Table 1.A2: DiD estimates by gender

	Baseline			Female			Male		
	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings
Treatment effect in the period . . .									
. . . 1-12 mths after	-0.002	0.007	0.008	-0.007	0.008	0.001	0.002	0.006	0.013
	(0.006)	(0.006)	(0.007)	(0.009)	(0.009)	(0.010)	(0.008)	(0.008)	(0.008)
. . . 13-17 mths after	-0.017**	0.017**	0.022**	-0.029**	0.023**	0.023	-0.007	0.013	0.023**
	(0.008)	(0.008)	(0.009)	(0.012)	(0.012)	(0.014)	(0.010)	(0.010)	(0.011)
. . . 18-24 mths after	-0.065***	0.033***	0.037***	-0.082***	0.035***	0.035**	-0.051***	0.030***	0.038***
	(0.007)	(0.008)	(0.009)	(0.011)	(0.011)	(0.014)	(0.009)	(0.010)	(0.012)
. . . 25-50 mths after	-0.006	0.015**	0.033***	-0.010*	0.010	0.035**	-0.002	0.019**	0.034***
	(0.004)	(0.007)	(0.009)	(0.006)	(0.011)	(0.014)	(0.005)	(0.009)	(0.011)
Avg. of dep. var.	0.81	0.91	5'387.37	0.81	0.91	3'881.60	0.80	0.91	6'593.96
R-squared	0.22	0.09	0.18	0.24	0.10	0.18	0.20	0.09	0.16
Obs.	3,073,557	3,073,557	3,073,557	1,366,336	1,366,336	1,366,336	1,707,221	1,707,221	1,707,221
Clusters	57,429	57,429	57,429	25,796	25,796	25,796	31,633	31,633	31,633

Notes: This table shows the difference-in-differences estimates for subsamples split by gender for unemployment benefit receipt, employment, and earnings respectively. Columns 1 to 3 replicate the baseline estimates, columns 4 to 6 include only females, and columns 7 to 9 include only males. Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own estimations based on merged UIR-SSA database.

Table 1.A3: DiD estimates by previous earnings

	Baseline			Low ($\leq$ 4,395 CHF per month)			High ($>$ 4,395 CHF per month)		
	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings	Benefit receipt	Employment	Earnings
Treatment effect in the period . . .									
. . . 1-12 mths after	-0.002	0.007	0.008	0.000	0.003	-0.007	-0.004	0.010	0.010
	(0.006)	(0.006)	(0.007)	(0.009)	(0.009)	(0.012)	(0.008)	(0.008)	(0.008)
. . . 13-17 mths after	-0.017**	0.017**	0.022**	-0.010	0.029**	0.017	-0.021**	0.007	0.019*
	(0.008)	(0.008)	(0.009)	(0.011)	(0.011)	(0.016)	(0.010)	(0.010)	(0.011)
. . . 18-24 mths after	-0.065***	0.033***	0.037***	-0.054***	0.035***	0.030*	-0.073***	0.029***	0.034***
	(0.007)	(0.008)	(0.009)	(0.010)	(0.011)	(0.017)	(0.010)	(0.010)	(0.011)
. . . 25-50 mths after	-0.006	0.015**	0.033***	-0.004	0.019*	0.024	-0.007	0.010	0.032***
	(0.004)	(0.007)	(0.009)	(0.006)	(0.011)	(0.016)	(0.005)	(0.009)	(0.010)
Avg. of dep. var.	0.81	0.91	5'387.37	0.78	0.88	2'750.17	0.83	0.93	7'520.64
R-squared	0.22	0.09	0.18	0.21	0.08	0.10	0.24	0.10	0.14
Obs.	3,073,557	3,073,557	3,073,557	1,373,671	1,373,671	1,373,671	1,699,886	1,699,886	1,699,886
Clusters	57,429	57,429	57,429	25,723	25,723	25,723	32,398	32,398	32,398

Notes: This table shows the difference-in-differences estimates for subsamples split by earnings for unemployment benefit receipt, employment, and earnings respectively. Columns 1 to 3 replicate the baseline estimates, columns 4 to 6 include low earnings ($\leq$ 4,395 CHF per month), and columns 7 to 9 include high earnings ($>$ 4,395 CHF per month). Earnings are normalized by the average earnings 3 months prior to unemployment start. Standard errors clustered by individual in parentheses. *** $P<0.01$ ** $P<0.05$ * $P<0.1$.
Source: Own estimations based on merged UIR-SSA database.

Chapter 2

Winning versus Losing: How Important are Reservation Wages for Non-employment Duration?

Abstract

Standard job search theory predicts that extending unemployment benefit durations prolongs non-employment, an effect that has been confirmed empirically in numerous studies. However, little is known about the empirical relevance of the two key margins - reservation wages and search effort - in determining job seekers optimal response to changes in unemployment benefit duration. This paper develops a new strategy designed to analyze the relative importance of the two margins for the duration of non-employment. To this end, I separately study *exits to wage-improving jobs* and *exits to wage-declining jobs*. Exit hazards to wage-improving jobs are solely determined by search effort, whereas the exit rate to wage-declining jobs is jointly determined by search effort and reservation wages. I test this in the context of a sharp discontinuity in potential benefit duration from 30 to 39 weeks and provide causal estimates for the effects of prolonged benefits on non-employment duration and survival probabilities. Consistent with reservation wage movements, exits to wage-declining jobs account for virtually all of the overall non-employment effect. Moreover, analyzing treatment effects on survivor functions highlights that the largest contributions to the overall unemployment effect are observed in the time period from 30 to 39 weeks. These results suggest an important role of the reservation wage channel in shaping job search behavior.

JEL Classification: J64, J65, C41

Keywords: potential benefit duration, unemployment duration, reservation wage, search effort

2.1 Introduction

Unemployment insurance is the most important policy tool to feather the negative effects of unemployment and to provide income replacement during job search. Around 25 million worker lost their job in the global crisis that erupted in 2008 (ILO, 2012). Unemployment insurance (UI) is the key first safety net to workers and probably the most important program to feather the effects of crises. Not surprisingly, a country's UI generosity plays an important role for the job search behavior of the unemployed. While designed to ease liquidity constraints of the unemployed during job search, UI can also generate reemployment disincentives. A standard prediction of job search theory asserts that extending UI benefit durations lowers search effort and raises reservation wages (Mortensen, 1977). Both behavioral margins – a job seekers' search effort and his reservation wage choice – tend to prolong non-employment duration.

A large empirical literature analyzed the effects of UI generosity on unemployment duration and the finding that increased UI generosity prolongs non-employment duration is one of the most robust results in labor economics (Cahuc and Zylberberg, 2004; Tatsiramos and Van Ours, 2014). However, there is an ongoing discussion about whether extending UI generosity has beneficial effects on job match quality or whether it only subsidizes unproductive search. In order to contribute to this discussion, separating the two behavioral margins is essential. So far only relatively little is known about the relative importance of search effort and reservation wages and how they shape labor market transitions in response to UI changes. Analyzing the relative importance of these two margins is difficult in practice, because reservation wages and search effort are rarely directly observed. Knowing more about these channels is nevertheless important from a policy perspective: On the one hand, if reduced search effort is the main driving force for prolonged non-employment spells, changes in UI could be coupled with stricter search requirement to curtail unwanted disincentive and moral hazard effects. On the other hand, increasing UI benefit duration could be welfare improving, if prolonged non-employment spells are mainly driven by reservation wage effects and allow job seekers to accept better job-matches.

This paper sheds light on the relative importance of the two behavioral margins of job search. I propose a new approach which allows to infer the relevance of the reservation wage choice for non-employment duration without actually observing reservation wages directly. Overall effects of extended UI benefit durations on non-employment duration and survival probabilities can be decomposed into contributions from *exits to wage-improving ("winning") jobs* and *exits to wage-declining ("losing") jobs*. Wage-improving exits are exits to jobs with reemployment wages that exceed pre-unemployment wages. Wage-declining exits are exits to jobs that are worse paid relative to a job seekers' previous job. I show that the likelihood of exits to wage-declining jobs is jointly affected by search effort and reservation wages, while the likelihood of exits to wage-improving jobs is solely determined by search effort. While reservation wages for exits to wage-improving jobs should not be binding, they may be binding for exits to wage-declining jobs. Thus, decomposing the overall effect of extending benefit durations into its contributions from winning and losing job exits allows to shed light on the relative importance of the reservation wage channel for the duration of non-employment. If reservation wages matter for job search, we would expect benefit extensions to primarily affect exits to wage-declining jobs.[1]

The empirical analysis relies on quasi-experimental variation in benefit duration that allows to identify the causal effect of extended benefit duration on non-employment duration and survivor functions. The paper exploits a sharp discontinuity in potential benefit duration (PBD) from 30 to 39 weeks around the age of 40 in Austria to analyze the effects of prolonged benefit duration. Overall, increasing PBD by 9 weeks prolongs non-employment duration by around 0.6 weeks. Consistent with search theory, the largest contributions to the prolonged non-employment duration are observed between 30 and 39 weeks, which corresponds to the period between the two benefit exhaustion dates. By decomposing the overall effects on non-employment and survival probabilities into contributions from exits to wage-improving and wage-declining jobs, I show that mainly exits to wage-declining jobs are affected. Exits to wage-declining jobs contribute account for virtually all of the total effect on non-employment, with its largest contributions in the time period from 30 to 39 weeks. These findings support the view that reservation wages play an important role for job search behavior.

The paper is related to several strands of literature. First, a large body of empirical studies analyzes the impact of UI policy changes on unemployment or non-employment durations, wages and other job characteristics. These studies confirmed the finding that extending UI generosity unambiguously prolongs unemployment duration. Starting from the observation that European countries with relatively generous UI systems have suffered much larger and much more persistent increases in unemployment in the 1980's than the United States, Katz and Meyer (1990) identify the potential unemployment benefit duration as a key driver of these cross-country differences and investigate the effect of potential benefit duration on unemployment duration. A number of studies, including Moffitt and Nicholson (1982), Moffitt (1985), and Grossman (1989) find significantly negative incentive effects. Winter-Ebmer (1998) uses Austrian data and finds significant benefit duration effects for males but not for females. In more a recent work, Lalive et al. (2006) use changes in the Austrian UI law as a natural experiment to examine the impact of the policy changes on the unemployment duration. They find that both an increase in the earnings replacement rate and a prolonged benefit duration lead to longer unemployment duration. Lalive (2008) studies the causal effects of a unique regional benefit extension on the unemployment duration in Austria and finds positive effects of the extension on unemployment duration for both men and women. Schmieder et al. (2012b) discuss the effects of extended PBD for benefit duration and non-employment duration over the business cycle in Germany. Moreover, a number of papers shows that the effects of unemployment benefit changes on the unemployment duration is not homogeneous over the unemployment spell. Meyer (1990), Katz and Meyer (1990) or Addison and Portugal (2004) find spikes in the

[1]This holds under the assumptions that job seekers set reservation wages strictly below their pre-unemployment wage and that search is not directed towards higher paying jobs. Both assumptions will be discussed in section 2.2.

unemployment exit rates just before benefit exhaustion. van Ours and Vodopivec (2006b) study the effects of PBD reductions in Slovenia and find both strong effects on the unemployment exit hazards and also substantial spikes around benefit exhaustion. Roed and Zhang (2003) find for Norwegian job seekers that the exit rate out of unemployment increases sharply close to benefit exhaustion and that the effects are stronger for females than for males.

Job search theory provides less guidance regarding the effects of prolonged benefit durations on reemployment wages and other job characteristics. On the one hand, more generous unemployment insurance policies such as longer benefit durations allow liquidity constrained unemployed to be more selective and to wait for better job offers. This is likely to improve reemployment wages. Also, job match quality can be improved and subsequent jobs should last longer. On the other hand, prolonging benefit durations can have negative effects on reemployment wages if the wage offer distribution is declining over the spell. Human capital and skill depreciation (Ljungqvist and Sargent, 1998) or stigmatization (Blanchard and Diamond, 1994) are possible causes for that. Empirical findings are mixed: Ehrenberg and Oaxaca (1976) are the first to look at the effect of unemployment insurance on post-unemployment outcomes and find positive effects of unemployment benefits on post-unemployment wages for different age groups and gender. A number of more recent studies also find positive effects of increased unemployment insurance generosity on post-unemployment wages (e.g. Addison and Blackburn (2000) and Centeno and Novo (2006) for the US, Centeno and Novo (2009) for Portugal, or Caliendo et al. (2013b) for Germany). Other studies, among them van Ours and Vodopivec (2008) for a Slovenian context or Card et al. (2007a) and Lalive (2007) for Austria, find either small or no effects on wages and/or job-stability.

A second strand of related literature is concerned with the role of reservation wages for labor market transitions. Most of this literature is based on survey evidence on self-reported reservation wages. Feldstein and Poterba (1984) find a relatively large elasticity of reservation wages with respect to unemployment benefit levels and conclude that reducing net unemployment insurance benefits could significantly lower the average unemployment duration through the reservation wage channel. DellaVigna and Paserman (2005) find that the self-reported reservation wage is positively correlated with a dummy for benefit receipt and is an important predictor for the actual reemployment wage. They conclude that reservation wages reflect an important aspect of job search behavior. Krueger and Mueller (2013) use high-frequency longitudinal survey data about self-reported reservation wages to provide evidence on the behavior of reservation wages over the spell of unemployment. They find – in accordance with the theoretical predictions of a non-stationary job search model – that reservation wages decline over the duration of the spell, though only at a modest rate. Moreover they find that the reservation wage have more predictive power than pre-displacement wages, suggesting that reservation wages contain useful information about workers' future decisions and thus play an important role for job search.

A small literature uses quasi-experimental variation in unemployment benefit eligibility to analyze the effects of increased PBD on reemployment wages and provide indirect evidence on reservation wages. Schmieder et al. (2013) analyze the causal effect of unemployment duration on reemployment wages in Germany. They decompose the effect of increased PBD into a component which is due to reservation wages and into a component which is due to shifts in the wage offer distribution over the duration of the spell. They find a negative and significant overall effect on reemployment wages and argue that this effect can be solely attributed to a declining wage offer distribution over the spell and not to reservation wages. Similar evidence is also found by Lalive et al. (2013) for Austria in the context of the regional benefit extension program which increased unemployment benefit duration drastically for a subset of workers in selected regions. Nekoei and Weber (2013) also analyze reemployment wages in Austria using a similar approach. They argue that the UI effect on expected wages is determined by two counteracting effects: On the one hand, UI increases reservation wages and tends to raise subsequent wages. On the other hand, job opportunities decrease due to the prolonged time spent in unemployment, which tends to decrease subsequent wages. Which of the two effects prevails depends on the importance of the agent's job seeking effort relative to the job selectiveness. They also exploit discontinuity around the age of 40 in Austria and find a statistically significant and positive effect of extended benefit duration on reemployment wages. Their finding suggests that job selectiveness (through the setting of a minimum acceptable target wage) plays an important role for job search behavior.

This paper complements the existing literature on the effects of UI changes on unemployment in several respects. First, I propose a novel approach how to study the role of the reservation wage channel for job search behavior. I show in a simple job search framework how the exit rate to jobs can be decomposed into contributions from exits to wage-improving jobs and exits to wage-declining jobs. Under two assumptions – job seekers set reservation wages below previous wages and do not direct search towards higher paying jobs – this decomposition is informative on the relevance of reservation wages for the duration of non-employment. Second, exploiting a sharp age discontinuity in unemployment benefit eligibility together with information on previous wages and reemployment wages, I show that reservation wages are an important factor of a worker's job search behavior. In doing so, this paper contributes to a small but growing literature which indirectly infers about reservation wages using only information about previous wages, reemployment wages and a quasi-experimental variation in benefit duration. Finally, the paper serves as a middle ground between the reduced form literature that analyzes the effects of UI changes on non-employment without inferring about the behavioral margins of job search and the structural literature which explicitly models reservation wages and search intensity from a job search model, but relies on untestable distributional assumptions in order to estimate the model parameters.

The remainder of the paper is structured as follows: Section 2.2 discusses how and under what assumptions the overall effects of extended benefit durations on non-employment and survival probabilities can be decomposed in order to be informative on the reservation wage channel. Section 2.3 describes the institutional background and the data and section 2.4 discusses the econometric framework. In section 2.5, I present the main results of how extending benefit duration affects non-employment duration and survival probabilities. Moreover, these findings are decomposed into its contributions from exits to wage-improving jobs and exits to wage-declining jobs. I also check the robustness of the findings in a number of sensitivity analyses. Section 2.6 concludes.

2.2 Conceptual framework

Partial-equilibrium, non-stationary job search models with endogenous search and stochastic wage offers[2] predict that UI extensions affect non-employment through reservation wages ρ_t and through search effort s_t (van den Berg, 1990b; Schmieder et al., 2013). These models predict that reservation wages decline and search effort increases over the duration of the spell until benefit exhaustion P and stay constant thereafter. The exit rate from non-employment is given by $\theta_t = s_t Pr(w_t \geq \rho_t) = s_t(1 - F(\rho_t; t))$ and is increasing until benefit exhaustion and stays flat thereafter.

Extending unemployment benefit durations increases the value of being non-employed prior to benefit exhaustion. Because the prolonged period of unemployment benefits b_t allows job seekers to maintain a higher consumption level for a longer time and reduces the pressure to find a new job quickly, job seekers lower their search effort, and maintain higher reservation wages and, consequently, exit later to jobs. Moreover, note that there is a direct relationship between the non-employment hazard rate θ_t, survivor functions $S(t)$ and the expected unemployment duration T: $S(t) = exp(-\int_0^t \theta_s \; ds)$ and $E(T) = \int_0^\infty S(u) \; du$. The negative effect on non-employment exit rates thus directly translates into higher survival probabilities and longer expected non-employment durations.

Overall changes in the observed non-employment duration and survival probabilities, however, cannot be directly mapped into changes due to search effort and/or reservation wages. In order to decompose the overall effects of extended benefit durations in a way which is informative on the two behavioral margins of job search – reservation wages and search effort – it is helpful to rewrite the non-employment exit rate as

$$\begin{aligned} \theta_t &= s_t\Big[\big(1 - F(w_0;t)\big) + \big(F(w_0;t) - F(\rho_t;t)\big)\Big] \\ &= s_t\big(1 - F(w_0;t)\big) + s_t\big(F(w_0;t) - F(\rho_t;t)\big) \\ &= \theta_t^W + \theta_t^L \quad . \end{aligned} \tag{2.1}$$

The first term inside the square brackets of equation (2.1) denotes the probability that the reemployment wage is above the previous wage, that is $w_t \geq w_0$, and the second term denotes

[2] Wage offers are assumed to be drawn from a i.i.d. distribution $F(w;t)$ that is allowed to decline over the duration of the spell for example due to stigmatization or skill depreciation.

the probability that the reemployment wage is below the previous wage but still acceptable, that is $w_0 > w_t \geq \rho_t$. Thus, we can learn about the importance of the reservation wage channel by decomposing the non-employment hazard into "winning" and "losing" job exit destinations. A "winning" exit destination, θ_t^W, is defined as an exit to a wage-improving job and a "losing" exit destination, θ_t^L, is an exit to a wage-declining job. Wage-improving jobs are accepted wage offers that pay wages *above* previous wages w_0 and wage-declining jobs are accepted wage offers that pay wages which are *below* previous wages w_0. The exit rate to wage-declining jobs is determined by the job seeker's search effort s_t and his reservation wage choice ρ_t, whereas exits to wage-improving jobs are influenced by the job seekers choice of search effort s_t only.

This decomposition is informative on the relevance of reservation wages for the unemployment duration and survivor functions under two assumptions: *First*, job seekers set reservation wages below previous wages. This assumption can be justified both from a theoretical and an empirical perspective: There are a number of theoretical reasons why one should expect wage reductions after an involuntary job loss: Loss of job-specific human capital, a deterioration of the value of an employee or incomplete information about the skills of a new employee (Feldstein and Poterba, 1984). Empirical evidence suggests that job seekers who lost their jobs involuntarily anchor reservation wages to previous wages and reduce reservation wages with increasing time spent in unemployment (Krueger and Mueller, 2013).[3] The *second* assumption is, that search is undirected. In order to separate the exit hazard into "winning" and "losing" exit destinations, we have to assume that search is not directed towards higher paying jobs, that is $s_t^W = s_t^L = s_t$. This assumption is justified if there is no wage posting. In Austria, wage posting is compulsory only from March 2011. Prior to this date, wages were not posted in job ads which makes the assumption of undirected search realistic for the empirical analysis.[4]

Clearly, if the wage offer distribution $F(w;t)$ is declining, the probability of exits to wage-improving jobs mechanically declines and the probability of exits to wage-declining jobs mechanically increases over the duration of the unemployment spell. However, assuming that an exogenous change in PBD does not have a direct effect on the wage offer distribution, i.e. that $\frac{\partial F(w;t)}{\partial P} = 0$, observed changes in unemployment duration and survivor functions have to be either due to search effort and/or reservation wages and the decomposition into "winning" and "losing" exit destinations is informative on the reservation wage channel. Formally, the exit

[3] However, the variability across workers is substantial: Feldstein and Poterba (1984) examine reservation wage choices of a large sample of unemployed job seekers in the United States in 1976 and find that a non-negligible fraction of job seekers sets reservation wages above previous wages. If a subset of workers who exit to wage-improving destinations set reservation wages above previous wages, decomposing overall effects becomes less informative on reservation wages: for the subset of workers who exit to wage-improving jobs and set reservation wages above previous wages their reservation wages choice might also have been binding. This issue will be further discussed in section 2.5.4.

[4] Krueger and Mueller (2011a) provide some evidence that a subset of workers is engaged in directed search. If search was directed towards high paying jobs, we would expect the search effort to wage-improving jobs to exceed the search effort for exits to wage-declining jobs, that is $s_t^W \geq s_t^L$. Noting that $\psi(.)$ is a convex and twice differentiable function, it holds that $\psi''(s_t^W) \geq \psi''(s_t^L)$ and thus $\frac{\partial s_t^L}{\partial b_P} \leq \frac{\partial s_t^W}{\partial b_P} < 0$. Thus, with directed search and wage posting, the exit hazard response of wage-declining exit destinations would be even more negative compared to wage-improving exit destinations. Consequently, if a subset of job seekers was engaged in directed search, the differential impact on unemployment exit hazards between winning and losing exit destinations could not be fully attributed to the reservation wage channel, but part of the difference would come through the search intensity channel.

hazard to wage-improving jobs is given by $\log \theta_t^W = \log s_t + \log(1 - F(w_0))$. A marginal increase in PBD on the log hazard rate to wage-improving jobs is characterized by

$$\frac{\partial \log \theta_t^W}{\partial b_P} = \frac{\partial \log s_t}{\partial b_P} \quad ,$$

and allows for a direct mapping from changes in the hazard rate into changes in search intensity. The log hazard rate to wage-declining jobs is given by $\log \theta_t^L = \log s_t + \log(F(w_0) - F(\rho_t^L))$ and a marginal increase in PBD is calculated as

$$\frac{\partial \log \theta_t^L}{\partial b_P} = \frac{\partial \log s_t}{\partial b_P} - \frac{f(\rho_t^L)}{F(w_0) - F(\rho_t^l)} \frac{\partial \rho_t^L}{\partial b_P} \quad .$$

The exit rate to wage-declining jobs is a product of changes in search intensity and changes in reservation wages. If reservation wages play a role, we would expect a stronger response to an increase in PBD for exits to wage-declining jobs, that is:

$$\frac{\partial \log \theta_t^L}{\partial b_P} < \frac{\partial \log \theta_t^W}{\partial b_P} < 0 \quad .$$

Due to the direct relationship between hazard rates and survivor functions we would expect that benefit extensions affect survival probabilities of wage-declining jobs more than wage-improving jobs. What is more, if reservation wages matter, the major contribution to the overall effect of PBD on average non-employment duration should come from exits to wage-declining jobs.

2.3 Institutions and data

Subsection 2.3.1 discusses the relevant institutional details and subsection 2.3.2 describes the data and the sampling procedure.

2.3.1 Institutional background

The empirical analysis uses administrative records for the universe of job seekers from Austria. Although virtually all private sector jobs are covered by collective bargaining agreements at the region and industry level, the Austrian labor market is relatively flexible and is characterized by a low unemployment rate (Card et al., 2007c; EIROnline, 2013). Over the period from 1993 to 2005, the average unemployment rate was around 4.2 %.

Job seekers in Austria are entitled to a limited period in which they can draw regular unemployment benefits. Voluntary quitters and workers discharged for misconduct are subject to a four-week waiting period. UB recipients must be employable and willing to work. Recipients are expected to search actively for a new job that should be within the scope of the claimant's qualifications.

Eligibility for unemployment benefits depends on prior unemployment insurance contributions and on age. In terms of UI generosity, Austria is comparable to the US: The replacement rate in Austria is rather low and replaces around 55 % of previous after-tax earnings. Job seekers qualify for unemployment benefits if they have worked at least for 52 weeks in the 2 years prior to their unemployment spell.[5] Job seekers who have worked fewer than 156 weeks in the past 5 years before the start of their unemployment spell can claim up to 20 weeks of unemployment benefits. Individuals with more than 36 months of unemployment insurance contributions in the past 5 years are eligible for 30 weeks of benefits. Since August 1989, the potential benefit duration also depends on age: Job seekers above 40 years old with more than 156 weeks within 5 years and more than 312 weeks within 10 years of work experience can claim benefits for 39 weeks and individuals aged 50 or more with at least 468 weeks of employment in the last 15 years before the start of the unemployment spell are eligible for a maximum benefit duration of 52 weeks.

Moreover, employees are protected by a firing regulation which obliges firms to pay a lump-sum severance pay equal to 2 months of salary for individuals who were laid off after at least three years of service. After exhaustion of their regular unemployment benefits, job seekers can claim unemployment assistance ("Notstandshilfe"), which is a means-tested, infinite secondary benefit. Because unemployment assistance benefits are reduced euro for euro by any other source of family income, Card et al. (2007a) calculate that the average unemployment assistance is around 38 % of the unemployment benefit level in the population.

2.3.2 Data description

I use data from two different sources to analyze the effects of extended UI benefit duration on the duration of non-employment. The first data source is the Austrian social security database (ASSD), which contains detailed information about individuals labor market histories from 1972 to 2010 for the private sector employees and the unemployed.[6] The database contains daily labor market states, yearly earnings and a limited set of demographic variables, such as month and year of birth, gender, state of origin, and some information about the employers, such as industry affiliation or geographical location of the firm. The second data source is the Austrian unemployment register (AMS) which is available from 1987 to 1998. From this data I extract education and marital status of the last recorded unemployment spell. For individuals, whose only unemployment spell started after 1998, these variables are missing.[7]

The main outcome is non-employment duration. It is measured as time elapsed between the end of the last job to the start of a new job.[8] Non-employment duration is right-censored

[5] For job seekers below the age of 25 at registration, the minimum work requirements prior to unemployment are 26 weeks within one year.

[6] The database does not include self employed and civil servants. Card et al. (2007c) report that around 10 % of the labor force were self employed and around 7 % were civil servants in 1996, so that the ASSD contains labor market histories of roughly 85 % of the total workforce.

[7] By using the information of the last recorded unemployment spell, I can still assign around 75 % of the information for spells that started after 1998.

[8] Using *registered* unemployment as main outcome – a measure that is based on the time elapsed between the

at two years. Less than 5 % of observations are censored. From the universe of individuals in ASSD, I only consider layoffs that ended in non-employment between August 1993 and December 2005. Focusing on this inflow window ensures that estimations are not affected by the regional extended benefits program (REBP) which was abolished in August 1993. Also, during that period there were no major reforms in the UI system which could bias estimates. I further focus on individuals aged between 30 and 50 years at the date of their unemployment registration. Moreover, I focus on job seekers with a continuous work history to ensure that individuals are eligible for at least 30 weeks of benefits: Only individuals with at least one year (52 weeks) of work experience out of 2 years prior to the start of the non-employment spell, with 3 (156 weeks) out of the last 5 years, and with at least 6 (312 weeks) out of the last 10 years prior to the start of the non-employment spell are retained in the sample. The sample may contain multiple spells per job seekers, but excludes spells with less than 7 days of length. Finally, in order to minimize the influence of seasonal workers in the sample, I exclude job seekers from the construction and tourism sector and also drop recalls to the previous firm. More than 40 % of all job seekers belong to the construction and tourism sector and around 27 % of the remaining individuals are recalled. The final sample counts 183,001 individuals and 258,337 spells.

2.4 Econometric framework

Subsection 2.4.1 presents the empirical specifications used to identify the causal effects of extended PBD on non-employment duration and survivor functions. Subsection 2.4.2 discusses the validity of the RD approach.

2.4.1 Empirical specification

Estimating the effect of extended UI duration on non-employment. As discussed above, the Austrian legislature for unemployment benefits contains sharp discontinuities with respect to age, which can be exploited to analyze how extending benefit duration affects non-employment. As described in section 2.3.1, benefit entitlement discontinuously changes around the threshold of 40 years. Job seekers below 40 years old are entitled to 30 weeks of benefits whereas job seekers above 40 years old are entitled to 39 weeks. The regression discontinuity approach allows to identify causal effects around this cut-off age.

Following Hahn et al. (2001), let $D_i \in 0, 1$ denote a binary treatment variable, indicating whether an individual is above the cut-off c of 40 years ($D_i = 1$) or below ($D_i = 0$). Because of exact knowledge of treatment assignment, D_i is a deterministic function of the forcing variable

registration and de-registration at the unemployment office – would be misleading and could lead to purely mechanical effects of changes in potential benefit duration, if job seekers de-register from the unemployment office once benefits are exhausted irrespectively of whether they found a job or not. Results using unemployment duration as outcome are available upon request.

age, A_i, that is:

$$D_i = \mathbb{1}(A_i \geq c)$$

Furthermore, let T_{i1} denote the outcome that occurs under treatment, and T_{i0} the outcome if not exposed to the treatment. The observed outcome T_i can be written as $T_i = T_{i0} + D_i(T_{i1} - T_{i0})$. Under some continuity assumptions, i.e. if $E[T_{i0}|A_i = a]$ and $E[(T_{i1} - Y_{i0})|A_i = a]$ are continuous in a at c and under a weak conditional independence assumption, the average treatment effect at the cut-off c can be written as

$$E[T_{i1} - T_{i0} \mid A_i = c] = \lim_{\varepsilon \downarrow 0} E[T_i \mid A_i = c + \varepsilon] - \lim_{\varepsilon \uparrow 0} E[T_i \mid A_i = c + \varepsilon]$$

Under the above assumptions, the average treatment effect for the job seekers at the cut-off can be obtained by estimating the discontinuity at the cut-off using the following empirical regression function:

$$T_i = \alpha + \beta D_i + f^-(A_i - c) + f^+(A_i - c) + \eta X_i + \varepsilon_i \quad .$$

The parameter β identifies the average causal effect of increasing PBD by 9 weeks on non-employment duration T_i at the threshold. $f^-(A_i - c)$ and $f^+(A_i - c)$ capture a possibly non-linear trend relationship between age and the duration of non-employment, which is allowed to differ on both sides of the age threshold. X_i is a set of control variables, such as year and month fixed effects, state and industry fixed effects, and a number of sociodemographic characteristics. Including control covariates is not needed for identification but might improve the precision of the estimates (Lee and Lemieux, 2010b). A crucial issue in a RD framework is the correct specification of the trend relationship between the outcome Y_i and the forcing variable A_i. Falsely assuming a linear relationship between non-employment and age might lead to the identification of discontinuities where there are none. Another relevant issue is the choice of the bandwidth: In a RD framework, there is an inherent trade-off between precision and bias. The main estimates are estimated using the data-driven asymptotically optimal bandwidth as proposed by Imbens and Kalyanaraman (2012). I perform a number of sensitivity tests in section 2.5.4 in order to test the sensitivity of results to the bandwidth choice and the order of the polynomial in age.

The main dependent variable is non-employment duration T_i, measured as the time elapsed between the end of the previous job until the start of the new job. I use this definition of non-employment rather than registered unemployment, because changes in PBD may affect registered unemployment duration in a purely mechanical way, if job seekers de-register from unemployment as soon as benefits exhaust irrespective of whether they found a job or not (Card et al., 2007c). Another potential issue with analyzing non-employment duration is the following: Analyzing average non-employment duration might be misleading if a lot of spells are right-censored. Right-censoring is however not an issue in this study, because less than 5

% of non-employment spells are right-censored.

Estimating the effect of extended UI duration on survivor functions. In a second part of the empirical analysis, I decompose the overall effect on non-employment into contributions to its change as a function of duration. In order to study the effects of UI changes on labor market transitions, it is useful to decompose the total effect of extended PBD on non-employment duration in the following way: The expected non-employment duration is defined as $E(T) = \int S(u)du$, where $S(t) = exp(-\int \theta(s)ds)$ is the survivor function and $\theta(.)$ is the non-employment exit hazard. Thus, analyzing treatment effects on survivor functions allows to study how the total non-employment effect is decomposed over the duration. For the analysis of the survivor functions, I calculate the probability that an non-employment spell lasts longer than t, that is $Pr(T > t)$, and estimate treatment effects on survivor functions period by period, that is

$$Pr(T_i > t) = \alpha + \beta D_i + f^-(A_i - c) + f^+(A_i - c) + \eta X_i + \varepsilon_i$$

for each period $t \in \{0, \ldots, 102\}$.

Winning versus losing: decomposition of overall effects. In order to learn about reservation wages, I additively decompose the overall effects on non-employment and survivor functions into its contributions from exits to wage-improving jobs and exits to wage-declining jobs. Let $W_i = \mathbf{1}(w_t > w_0)$ be an indicator for an exit to a wage-improving job and $L_i = 1 - W_i = \mathbf{1}(w_t \leq w_0)$ an indicator for an exit to a wage-declining job. Non-employment duration (and survival probabilities respectively) can thus be additively decomposed into $T_i = W_i \times T_i + (1 - W_i) \times T_i = W_i \times T_i + L_i \times T_i$. Then, I can estimate the contributions from the winning and the losing exits separately using $W_i \times T_i$ and $L_i \times T_i$ as dependent variables. The overall effect on non-employment then additively decomposes into contributions from exits to winning and exits to losing jobs. In the same way, the survival probabilities $Pr(T_i > t)$ are decomposed into $Pr(T_i > t) \times W_i + Pr(T_i > t) \times L_i$ and estimated separately for the two components.

2.4.2 Validity of the RD approach

Identification in a RD framework mainly rests on the assumption of continuity of the potential outcomes around the cut-off with respect to age. In other words, the RD approach is suitable if treatment is as good as randomly assigned around the threshold. This assumption could be violated if individuals are able to influence treatment assignment. Treatment assignment depends on age and prior work experience. Work experience as well as age at registration can be influenced by job seekers to some extent, because job seekers could wait with unemployment registration until they reach a certain age threshold or work experience requirement. The extent to which job seekers can manipulate the start date of unemployment is however limited, because employers or job seekers have to announce their unemployment spell at the latest the day after the end of the job in order to avoid cuts in benefit payments.

As a test of such strategic behavior I examine the density of the running variable around the threshold. If individuals sort themselves into treatment, then one should observe bunching around the threshold. In other words, in an appropriate RD design the marginal density of age over the population should be continuous. McCrary (2008) proposes a formal test for manipulation of the running variable. Figure 2.1a shows the inflows into unemployment as a function of age. The vertical line at the age of 40 indicates the threshold above which job seekers can claim 39 instead of 30 weeks of benefits. The figure does not show any evidence that job seekers manipulate their age at unemployment entry. Figure 2.1b shows an undersmoothed histogram together with the local linear density estimates proposed by McCrary (2008). There is no discontinuity in the density around the age threshold. A formal test of continuity around the cut-off value fails to reject the null hypothesis of continuity with a t-value of 0.44.

Figure 2.1: Density around cut-off

(a) Inflows around age threshold   (b) Local linear density around age threshold

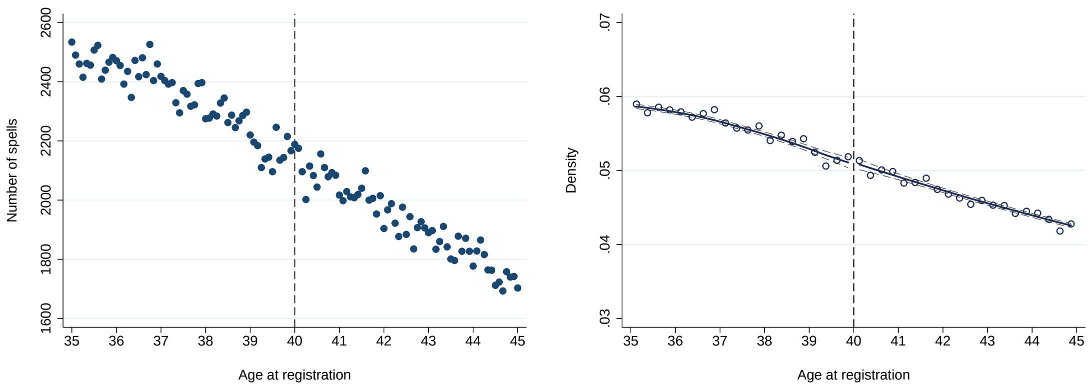

Notes: Figure 2.1 shows the density of the running variable around the threshold value of 40 years. The x axis shows age at registration. A window of 5 years around the threshold is shown. Subfigure 2.1a shows the inflows into unemployment around the cutoff value, and subfigure 2.1b shows an undersmoothed histogram together with the local linear density estimates proposed by McCrary (2008).
Source: Own calculations based on ASSD.

A second analysis for the validity of the identification assumption is a test of continuity of observable characteristics. Discontinuous variation of the observables around the threshold would be a strong indication for a failure of the identifying assumption. In figure 2.B1 I examine a range of characteristics above and below the threshold. Individuals are very similar above and below the threshold and none of the characteristics exhibit a jump at the threshold. Table 2.1 shows a formal discontinuity test for all covariates. The formal test confirms the graphical evidence: Most of the characteristics do not vary statistically significantly around the threshold. Although we reject continuity of covariates in a few cases, such as the occurrence of past non-employment spells, university degree and region, the differences are economically very small. Overall, the analysis of the covariates around the threshold suggests that the assumption of as good as random assignment around the threshold fails to be rejected.

Table 2.1: Covariates discontinuity test

	Overall	Winning	Losing
A. Labor market history			
Mean past earnings	4.596	-4.527	9.123
	(6.271)	(7.671)	(8.132)
Mean past wage	0.155	-0.133	0.288
	(0.191)	(0.261)	(0.272)
Past unemployment spell	-0.002	0.002	-0.004**
	(0.002)	(0.002)	(0.002)
Work exp. in past 15 years (in weeks)	-1.331	-4.355*	3.024
	(0.815)	(2.558)	(2.786)
Tenure (in weeks)	1.949	0.310	1.640
	(1.489)	(0.932)	(1.635)
Severance pay	-0.001	0.000	-0.001
	(0.003)	(0.002)	(0.003)
B. Worker Characteristics			
Female	0.004	0.002	0.002
	(0.004)	(0.003)	(0.003)
Austrian	0.001	-0.007*	0.008**
	(0.003)	(0.004)	(0.004)
Married	-0.003	-0.004	0.001
	(0.004)	(0.004)	(0.004)
Education			
Less than elementary school	0.001	-0.000	0.001
	(0.001)	(0.001)	(0.001)
Elementary school	-0.001	-0.004	0.003
	(0.004)	(0.003)	(0.003)
Apprenticeship/High School	-0.003	-0.004	0.001
	(0.004)	(0.003)	(0.003)
University	-0.002*	0.000	-0.002**
	(0.001)	(0.001)	(0.001)
Other	0.005**	0.003	0.003
	(0.003)	(0.002)	(0.002)
Previous industry			
Manufacture	-0.000	-0.003	0.003
	(0.003)	(0.003)	(0.003)
Wholesale and retail trade	0.002	0.001	0.002
	(0.003)	(0.002)	(0.003)
Financial, insurance activities, extraterritorial bodies	-0.004	-0.002	-0.002
	(0.003)	(0.002)	(0.002)
Transportation	0.002	-0.001	0.003**
	(0.002)	(0.001)	(0.002)
Health and social activities	0.001	0.001	0.000
	(0.002)	(0.001)	(0.002)
Other	-0.001	-0.000	-0.001
	(0.002)	(0.002)	(0.001)
Region			
Vienna	-0.007*	-0.008***	0.001
	(0.004)	(0.002)	(0.003)
Lower Austria	0.007***	0.002	0.005**
	(0.003)	(0.002)	(0.003)
Upper Austria	-0.001	0.001	-0.002
	(0.003)	(0.002)	(0.002)

Table 2.1 – continued

	Overall	Winning	Losing
Burgenland	-0.003***	-0.002**	-0.002**
	(0.001)	(0.001)	(0.001)
Carinthia	0.000	-0.000	0.000
	(0.002)	(0.001)	(0.001)
Salzburg	0.002	-0.000	0.002
	(0.002)	(0.001)	(0.002)
Styria	0.002	-0.001	0.003*
	(0.002)	(0.002)	(0.002)
Tyrol	0.000	0.002	-0.002
	(0.002)	(0.001)	(0.001)
Vorarlberg	-0.000	0.001	-0.001
	(0.001)	(0.001)	(0.001)
Unknown	0.000	0.001	-0.000
	(0.001)	(0.001)	(0.001)

Notes: This table presents first-order polynomial RDD estimates for the covariate controls with a bandwidth of 5 years. Standard errors clustered by age in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

2.5 Results

This section discusses the empirical results. Subsection 2.5.1 presents descriptive evidence for the movements of reservation wages over the non-employment spell. Subsection 2.5.2 shows the causal effects of extending benefits on non-employment duration and survivor functions. Subsection 2.5.3 decomposes the overall effects into its contributions from exits to wage-improving and wage-declining jobs, and subsection 2.5.4 discusses some sensitivity analyses.

2.5.1 Descriptive evidence

In subsection 2.2 I showed how exit hazards to jobs are separable into exits to wage-improving jobs and exits to wage-declining jobs. By decomposing the exit rate in this way I can isolate the reservation wage channel from the search intensity channel and learn about reservation wage movements over the non-employment spell. The underlying idea is that reservation wage choices directly affect the likelihood of exits to wage-declining jobs, but not the likelihood of exits to wage-improving jobs, whereas a job seekers choice of search effort affects both exit destinations likewise.

In a non-stationary job search model, wage offers are drawn from a random wage offer distribution $F(w;t)$ and thus - by chance - generate wage offers which are above previous wages for some workers and below previous wages for other workers. Figure 2.2 shows the distribution of reemployment wages relative to previous wages. Around 56 % of the spells are exits to wage-declining jobs. A relatively large proportion of job seekers accepted wage offers which are relatively close to their previous wages. Only around 25 % of job seekers have a ratio of reemployment to previous wages below 0.8. Even though recalls were excluded from the sample, around 10 % of job seekers gain exactly the same reemployment wage as they had prior to their unemployment spell. A possible explanation for this could be that these wages were bargained by unions. Only relatively few job seekers have a ratio of reemployment to previous wage considerably above 1. The ratio of reemployment to pre-unemployment wages is 1.01 or below for more than 75 % of job seekers, and equal to or below 1.23 for around 90 % of job seekers.

Figure 2.2: Distribution of reemployment wages relative to previous wages

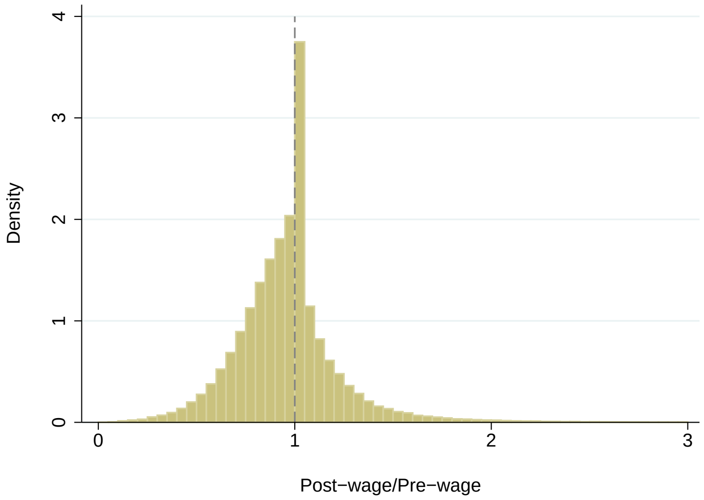

Notes: This figure shows the distribution of the ratio of reemployment wages relative to previous wages.
Source: Own calculations based on ASSD.

One major prediction of the non-stationary job search model is that reservation wages fall over the duration of the spell until benefit exhaustion and stay constant thereafter and that search effort is increasing over the spell until benefit exhaustion and is flat thereafter. A simple test to see whether the decomposition into the two exit destinations is informative on reservation wage choices is to look at the ratio of the hazard rates $\frac{\theta_t^L}{\theta_t^W}$ over the spell duration:

$$\frac{\theta_t^L}{\theta_t^W} = \frac{s_t Pr(w_0 > w_t \geq \rho_t)}{s_t Pr(w_t \geq w_0)} = \frac{F(w_0;t) - F(\rho_t;t)}{1 - F(w_0;t)}$$

Assume for the moment that the wage offer distribution is constant over the spell: Because search is assumed to be undirected, search effort s_t cancels out of the ratio of the hazard rates. Movements of the hazard ratio are then informative of reservation wage movements over the spell duration. If reservation wages matter for exits to wage-declining jobs, the ratio of the hazard rates should be increasing until benefit exhaustion P and be constant thereafter.

Figure 2.3 shows the smoothed hazard rates for exits to wage-improving and wage-declining jobs destinations. Subfigure 2.3a shows hazard rates for job seekers with 30 weeks of benefit entitlement and subfigure 2.3b for job seekers with 39 weeks of benefits respectively. Subfigures 2.3c and 2.3d depict the corresponding hazard ratios. Under both benefit regimes, hazard rates to wage-declining jobs are relatively flat or increasing prior to benefit exhaustion. After benefit exhaustion hazard rates decline. In contrast, hazard rates of wage-improving job destinations are declining over the whole spell duration. Comparing winning and losing exit rates thus highlights a diverging pattern of winning and losing hazards prior to benefit exhaustion, which becomes parallel after benefit exhaustion.

The observed pattern aligns well with the theoretical predictions: In a non-stationary job search model we would expect hazard rates to increase until benefit exhaustion and to be flat after benefits ran out. Because exits to wage-improving job destinations are affected by search effort only and exits to wage-declining job destinations are affected by search effort and reservation wages, we would expect a steeper slope for the exit hazard to wage-declining jobs prior to benefit exhaustion and parallel movements after benefit exhaustion.[9]

Subfigures 2.3c and 2.3d highlight the differential evolution of the hazard ratios before and after benefit exhaustion. The hazard ratio for job seekers with 30 weeks of benefits increases until week 30. At week 30 there is a kink and the ratio is increasing at a much lower rate thereafter. The same pattern is observed for the hazard ratio of job seekers with 39 weeks of benefits, but the kink is now observed around week 39. This pattern is consistent with the predictions of the non-stationary job search model. Because reservation wages decrease over the duration of unemployment, the numerator of the hazard ratio, $F(w_0;t) - F(\rho_t;t)$ increases with the duration of unemployment. After benefit exhaustion, the environment becomes stationary and the slope in the hazard ratio becomes much flatter. The observation that the ratio of the hazard rates is not completely flat after benefit exhaustion, might be due to the fact that the wage offer distribution is not constant over the duration of unemployment, but rather declining.

Clearly, if the wage offer distribution is declining over the spell, then $\frac{\theta_t^L}{\theta_t^W}$ increases mechanically over the spell duration. A key insight, however, is that there is no reason why the wage offer distribution should become stationary at benefit exhaustion - rather one would expect the wage offer distribution do decrease continuously around benefit exhaustion. Therefore, with a declining wage offer distribution, we would expect $\frac{\theta_t^L}{\theta_t^W}$ to be increasing after benefit exhaustion P. If reservation wages however matter, we would expect to see a steeper slope prior to benefit exhaustion. What is more, we would expect to find a kink in the slope of $\frac{\theta_t^L}{\theta_t^W}$ at benefit exhaustion, because the reservation wage path becomes constant after that.

[9] The fact that the hazard rates are not flat after benefit exhaustion but rather decreasing could stem from a declining wage offer distribution: If $[1 - F(w_0)]$ is decreasing with increasing time spent out of employment, then the overall hazard rate can be decreasing.

Figure 2.3: Hazard ratios for winning and losing job exits

(a) Exit hazards with 30 weeks of benefits

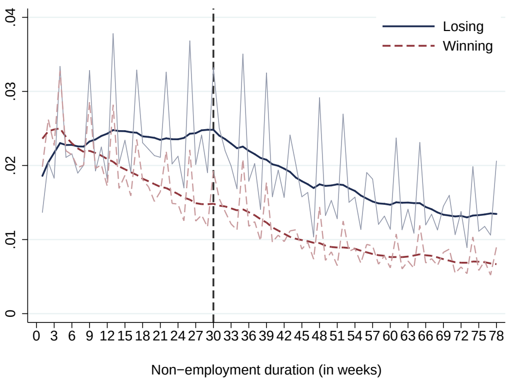

(b) Exit hazards with 39 weeks of benefits

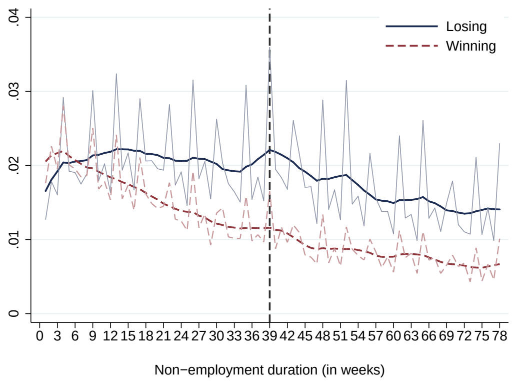

(c) Ratio of hazard rates

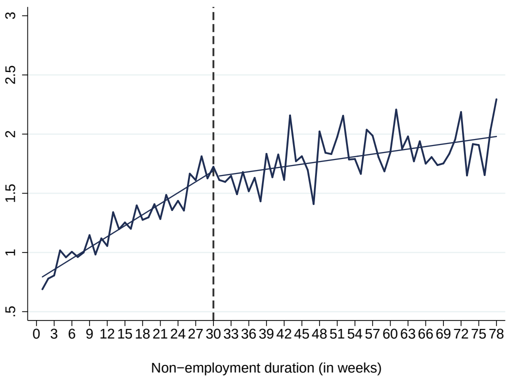

(d) Ratio of hazard rates

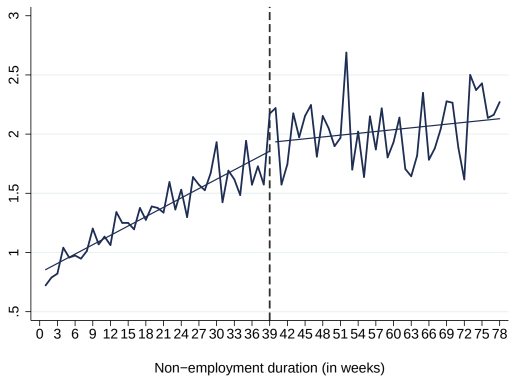

Notes: This figure shows the exit hazards (subfigures 2.3a and 2.3b) and the ratio of the hazard rates of wage-declining exits to wage-improving exits over the duration of non-employment (subfigures 2.3c and 2.3d) for job seekers with 30 or 39 weeks of benefit entitlement respectively.
Source: Own calculations based on ASSD.

While the divergent patterns in the hazard ratio prior and after benefit exhaustion and the kink at benefit exhaustion in figure 2.3 provide a first descriptive evidence on the role of reservation wage choices for job search behavior, it is still difficult to disentangle reservation wage effect from effects due to a declining wage offer distribution.[10]

[10]Figures 2.B2 and 2.B3 in the appendix provide additional evidence for a declining wage offer distribution over the duration of non-employment. First note, that with a declining wage offer distribution over the spell, receiving a wage-improving offer gets less and less likely the more time spent non-employed. This leads to a mechanical increase in the likelihood of exits to wage-declining jobs over the duration of the spell. Figure 2.B2 shows how the probability of exit to wage-declining jobs given exit evolves over the first 78 weeks of the spell. Early in the spell, up to 6 weeks since the start of the non-employment spell, the probability of exiting to wage-declining jobs is below 50 % - thus it is more probable to exit to a wage-improving job than to a wage-declining job. Then, with increasing spell duration, the probability of exit to wage-declining jobs is increasing to over 65 %. After around 52 weeks, the curve eventually starts to level off at around 70 %. Exits to wage-declining jobs thus get more and more probable with increasing time spent unemployed. Figure 2.B3 provides another indication for a declining wage offer distribution over the duration of the spell. The figure plots average *previous* wages over the duration of the spell. With a declining wage offer distribution, the likelihood of exits to wage-improving jobs mechanically declines over the duration of the spell. Because it is getting increasingly difficult to exit to wage-improving jobs with relatively well paid previous jobs, we should consequently observe a decline in average *previous* wages for exits to wage-improving jobs, but not for exits to wage-declining jobs. The figure shows that average previous wages indeed fall over the spell duration for exits to wage-improving jobs. For exits to wage-declining jobs, average previous wages first also decline somewhat in the first around 50 weeks of the spell, but much less. Both figures thus point towards a declining wage offer distribution over the duration of the spell.

However, assuming that there is no direct effect of an exogenous increase in UI benefit duration on the wage offer distribution, observed changes in non-employment duration and survivor functions have to be either due to so search effort and / or reservation wage considerations. If reservation wages matter, extending UI benefit durations should affect the likelihood of exits to wage-declining jobs more than the likelihood of exits to wage-improving jobs. The major contribution to the prolonged non-employment duration should then come from exits to wage-declining jobs. We can exploit a quasi-experimental variation in UI benefit eligibility around the age of 40 in Austria to empirically investigate these theoretical predictions.

2.5.2 The effects of extended benefit duration

Overall effects on non-employment duration. I start by estimating the overall effects of extended UI durations on non-employment duration. In doing so I replicate the well-known finding that increasing UI duration prolongs non-employment. Figure 2.4 shows observed non-employment duration (in weeks) as a function of age at unemployment registration. The vertical line at the age of 40 years indicates the cut-off value, above which job seekers can claim 39 weeks of benefits. To the left of the threshold job seekers are eligible for 30 weeks of benefits. Each dot represents average unemployment duration for job seekers in age bins of one month. The fit of a linear regression which allows for a discontinuity at the age threshold and for different age trends on both sides of the cut-off is superimposed together with the 95 % confidence interval.

Average non-employment duration is roughly 25 weeks below the threshold, and around 28 weeks above the age threshold and increases with age. A discontinuity at the threshold can be interpreted as a first descriptive evidence of a causal effect of extended PBD, as long as the assumption of continuity of potential outcomes around the cut-off is satisfied. The observed discontinuity at the threshold shows that increasing benefit entitlement by 9 weeks increases average non-employment duration by around half a week. Because I use elapsed time to the new job rather than time to unemployment de-registration as main outcome, the jump around the age threshold reflect pure behavioral changes due to the extension of UI duration.[11]

[11] Using *registered* unemployment rather than time to new job as outcome variable shows larger jumps around the age threshold. These effects may, however, be partly mechanical, if job seekers de-register from the unemployment insurance system at benefit exhaustion irrespective of whether they found a job or not. Graphs based on registered unemployment are available upon request.

Figure 2.4: Duration of non-employment around the age threshold

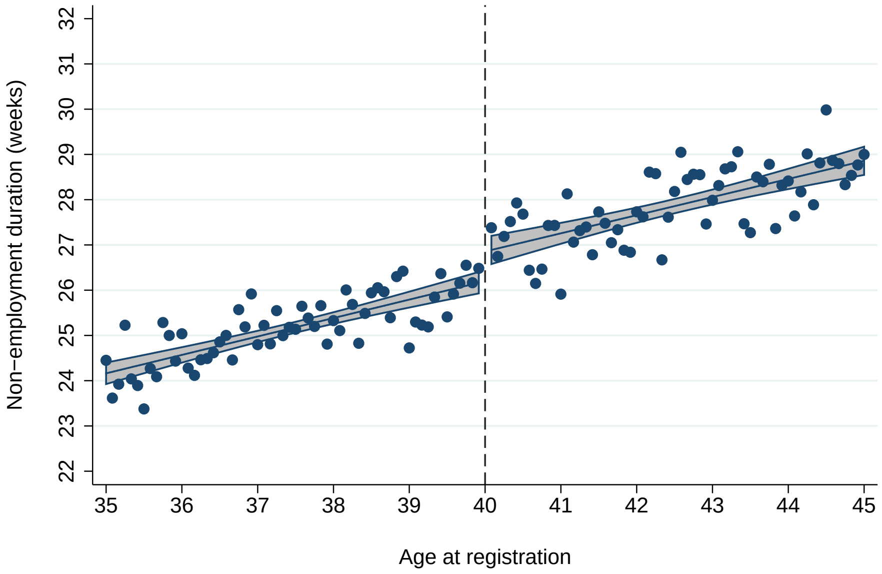

Notes: Figure 2.4 shows non-employment duration (measured in weeks) as a function of age at registration. Each dot corresponds to a bin size of one month. A window of 5 years around the threshold is shown. The vertical line indicates the threshold value at 40 years old.
Source: Own calculations based on ASSD.

Table 2.2 shows the treatment effect of increasing PBD by 9 weeks on non-employment durations. The regressions were estimated without control covariates and were obtained using a linear specification in age and the asymptotically optimal bandwidth proposed by Imbens and Kalyanaraman (2012). The optimal bandwidth is 3.456 years. A large number of robustness analyses testing the sensitivity of the estimates with respect to variations in the econometric specification will be discussed in section 2.5.4.

The causal effect of extending UI benefits on non-employment is 0.614 and is estimated at the 5 % significance level. Extending UI benefits by 9 weeks prolongs non-employment by roughly 4.5 days. One additional month of UI benefits thus prolongs non-employment by 0.27 weeks.

How do these findings fit into the existing literature? One caveat of the regression discontinuity approach is external validity. This analysis is based on a sample of prime-wage workers with long and stable labor market histories. Displaced workers with a stable labor attachment might however react differently to benefit extensions than younger job seekers with relatively little work experience. Also, the benefit extension analyzed in this paper is relatively modest and it is not clear *a priori* how results would change if benefit extensions would be more important. However, I find that the estimates fit very well in the existing literature. Schmieder et al. (2013) find around 0.3 additional months of unemployment per month of increased UI duration and 0.15 months of additional non-employment - magnitudes that are slightly larger than the ones found in this paper. With 0.12 weeks of non-employment per additional month of UI the effect sizes are slightly smaller in Nekoei and Weber (2013).

Table 2.2: RD estimates

D_i	0.614**
	(0.243)
$(A_i - c)$	0.316***
	(0.068)
$D_i(A_i - c)$	0.193
	(0.117)
Bandwidth	3.456
Observations	177,821
R-squared	0.002

Notes: This table presents RDD estimates for non-employment duration (in weeks) Standard errors clustered by age in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

Overall effects on survivor functions. In a second step I decompose the overall effects on non-employment into its contributions over the spell by estimating the effect of an extended UI duration on survival probabilities in each period of a spell. Clearly, extending UI durations increases survival in non-employment. Figure 2.5 shows descriptive evidence on the survivor functions for the group of job seekers with 30 weeks of benefits (dashed line) and for the group of job seekers with 39 weeks of benefits (solid line). Both control and treated job seekers exit non-employment relatively fast: Around 50 % of the job seekers found a job after 15 weeks. After 30 weeks, roughly 75 % of job seekers found a job, and after 52 weeks, only less than 20 % are still non-employed. Treated job seekers with 30 weeks of benefits, however, leave non-employment earlier than their counterparts with 39 weeks of benefits. The survivor functions of treated and controls start to differ immediately, but the gap between the two functions widens with increasing duration. The maximum difference is observed in the time period between 30 and 39 weeks. After that, the difference between the two survivor functions shrinks back.

Figure 2.5: Survivor functions

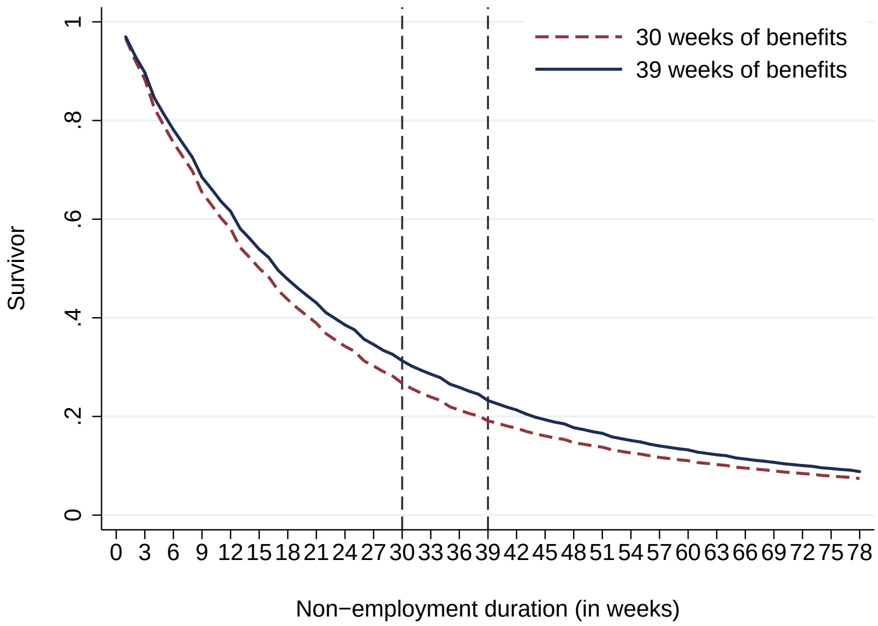

Notes: This figure shows survivor functions for non-employment for job seekers aged between 35 and 40 years old at registration (eligible for 30 weeks of benefits) and job seekers aged between 40 and 45 years old (eligible for 39 weeks of benefits) respectively. *Source:* Own calculations based on ASSD.

Figure 2.6 shows the estimated treatment effects on the survivor functions for non-employment estimating the period by period effects of extended UI duration on the survival probability, that is on $P(T_i > t)$ for each $t \in \{0, \dots, 104\}$. The difference between the two survivor functions is not significantly different from 0 in the first 26 weeks of the spell. The survival probabilities of job seekers with an extended benefit duration of 39 weeks starts to differ from those of job seekers with only 30 weeks of benefits after 26 weeks. The difference between survival probabilities peaks in the time period between 30 and 39 weeks. Extending the benefit duration by 9 week has a maximum contribution of around 1.5 percentage points during the peak period and shrinks back close to zero thereafter. The maximum contribution to the overall non-employment effect therefore comes from the period between the two exhaustion dates at 30 and 39 weeks respectively.

Figure 2.6: Estimated effects for survivors

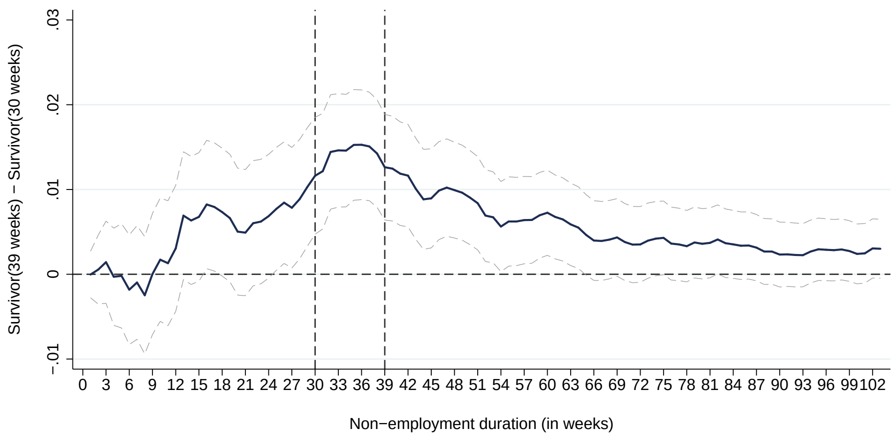

Notes: Figure 2.6 shows estimated treatment effects on survivors functions for non-employment.
Source: Own calculations based on ASSD.

2.5.3 Decomposition of overall effects by exit destination

Table 2.3 shows how the overall effect of extended UI durations on total non-employment duration decomposes into its contributions from job seekers who exit to wage-improving jobs and job seekers who exit to wage-declining jobs respectively. Table 2.3 presents the estimation results. Column 1 replicates the overall treatment effect, columns 2 and 3 show the contributions of the two exit destinations respectively. Calculations are based on a linear specification in age and uses the asymptotically optimal bandwidth of 3.456 years. The estimates suggest that mainly exits to wage-declining jobs are affected by extended benefit durations: The overall non-employment effect of around 0.614 weeks splits up into a contribution of around 0.597 weeks due to exits to wage-declining jobs and 0.018 weeks due to exits to wage-improving jobs. Thus, losing exits account for more than 95 % of the overall non-employment effect, whereas the contribution of job seekers who exit to wage-improving jobs is small and not statistically significant. For non-employment, almost over 95 % of the overall effect is accounted for by exits to wage-declining jobs. Absent a direct effect of UI extensions on the wage offer distribution, these estimates suggest an important role for reservation wages in the job search behavior of displaced workers.

Table 2.3: RD estimates decomposed by exit destination

	Overall	Winning	Losing
D_i	0.614**	0.018	0.597**
	(0.243)	(0.149)	(0.285)
Bandwidth	3.456	3.456	3.456
Observations	177,821	177,821	177,821

Notes: This table presents decomposed RD estimates into exits to wage-improving (winning) jobs and exits to wage-declining (losing) jobs. Standard errors clustered by age in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

Figure 2.7 decomposes the overall effects on survivor functions into contributions from winning and losing exits respectively. The dashed red line displays the period-per-period effects of extended UI durations on the survival probabilities of exits to wage-improving jobs, i.e. $P(T_i > t) \times W_i$, and the solid blue line depicts the period-per-period effects for losing exits, that is $P(T_i > t) \times L_i$. The thin gray line replicates the overall effects of figure 2.6. The figure shows that the major contribution to the positive treatment effects on the survivor functions comes from the exits to wage-declining jobs. Even early in the spell the most important contribution comes from job seekers who exit to wage-declining jobs. In the period from 30 to 39 weeks, the contribution from exit to wage-declining jobs is around two thirds. In the later periods, especially after 52 weeks, overall contributions to the total non-employment duration decrease. However, virtually all of the remaining effect is accounted for by losing exits.

Figure 2.7: Estimated effects on survivor functions for winning and losing exits

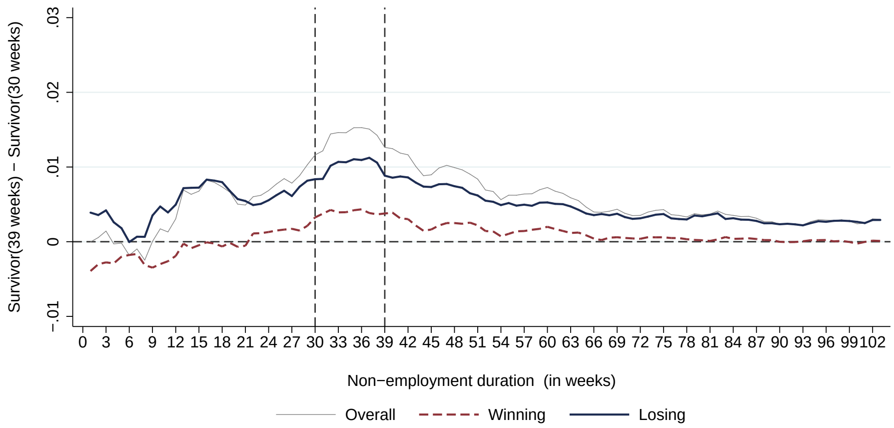

Notes: This figure shows the decomposition of the estimated treatment effects on non-employment hazard rates and on survivors functions for wage-improving and wage-declining exits.
Source: Own calculations based on ASSD.

2.5.4 Sensitivity analyses

Robustness of RD estimates. In this section a number of model variations are discussed. A first set of robustness regressions is concerned with the sensitivity of the results to the econometric specification. Table 2.4 presents a number of regressions which test the sensitivity of the estimates with respect to the bandwidth choice and the order of the polynomial in age. Misspecification of the trend relationship between non-employment and age might result in detecting discontinuities where there are none. Estimates using quadratic (cubic) terms in $(A_i - c)$ are presented in the second (third) row. The estimates from a local linear regression which uses Epanechnikov weights is displayed in the fourth row. Columns 2 to 4 use bandwidths of 2, 5, and 7 years respectively and provide sensitivity tests to the data-driven bandwidth in column 1.

The estimation results are quite robust to changes in the econometric specification. Most estimates range between 0.413 and 0.737 and are very close to the baseline estimates for most specifications. The optimal order of the polynomial is 1 for all specifications according to Akaike's information criterion.

In order to investigate how sensitive the estimations are to the inclusion of additional control covariates, I re-estimated the baseline specification of table 2.2 including a range of individual characteristics such as sex, marital status, nationality, education, previous industry and geographical region. Furthermore, I include a number of covariates that pick up job seekers previous labor market histories, such as past earnings and wages, past unemployment incidence, past work experience, tenure, and an indicator of whether a job seeker is eligible for severance pay. Moreover, dummies for the registration years and months are included to pick up seasonal variations and variations over the business cycle. Adding observed characteristics should not affect the identification of the treatment effects, but might pick up some random variation and improve precision of the estimates. Table 2.C1 shows that including control covariates does not change the point estimates in an important way, but improves precision somewhat.

Estimating the effect of extended UI duration on hazard rates. Previous analyses discussed effects of extended PBD on non-employment duration and survivor functions. The most direct test of the theoretical predictions, however, would be to analyze hazard rates. Appendix 2.A provides a graphical evidence of the effects of extended UI duration on job finding hazards.

One difficulty that arises when analyzing hazard rates, however, is that there might be dynamic selection over the course of the non-employment spell (Ham and LaLonde, 1996). As good as random assignment to treatment at unemployment registration does not guarantee that treatment incidence is independent of unobservables in later periods of the spell. If there

Table 2.4: RD estimates with different bandwidth and polynomial order

	Opt B.	B=2	B= 5	B=7
p=1	0.614**	0.825***	0.586***	0.584***
	(0.243)	(0.295)	(0.201)	(0.170)
p=2	0.664**	0.413	0.696**	0.623**
	(0.324)	(0.392)	(0.292)	(0.250)
p=3	0.629	-0.108	0.609*	0.666**
	(0.432)	(0.528)	(0.343)	(0.320)
LLR	0.631**	0.737**	0.620***	0.595***
	(0.252)	(0.295)	(0.212)	(0.181)
Order of p	1	1	1	1
Bandwidth	3.456	2.000	5.000	7.000
Observations	177,821	100,587	254,100	354,157

Notes: This table presents RDD estimates for non-employment duration (in weeks) for different orders of the polynomial *p* and different bandwidth choices *B*. LLR are estimates from a local linear regression using Epanechnikov kernel weights. Standard errors clustered by age in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

is dynamic selection, the pool of workers on both sides of the threshold may no longer be comparable in later periods of the spell even though it was at the start of the spell. Analyzing effects of extended UI durations on hazard rates thus remain suggestive and cannot be identified without making additional assumptions about the distribution of the unobserved heterogeneity.

For the analysis of the transitions out of unemployment, I estimate a piecewise-constant exponential model with unobserved heterogeneity which is assumed to follow a gamma distribution $G(\nu)$:

$$\theta(t \mid D_i, A_i, X_i) = \exp\left(\sum_{k=1}^{5} \lambda_k I(\tau_{k-1} < t < \tau_k)\right) \nu \quad , \tag{2.2}$$

with $\lambda_k = \alpha_k + \beta_k D_i + \gamma_k(A_i - c) + \delta_k D_i(A_i - c) + \eta X_i$ for $k = 1, \ldots, 5$. τ_{k-1} and τ_k define the cut-points of the different intervals. Within each interval the baseline hazard is restricted to be constant. The cut-points are set at 15, 28, 31, 37 and 40 weeks in order to capture differences in behavior early in the spell, differences around the benefit exhaustion dates, as well as differences in later periods in the spell. The β_k coefficients identify the average treatment effect around the age threshold under the assumption that the hazard rates of the two groups would be continuous in age at each duration interval. This duration-dependent specification allows to identify treatment effects at different lengths of non-employment.

Table 2.5 shows estimates from a piecewise-constant exponential, Gamma distributed shared frailty model as presented in equation 2.2. The estimates are consistent with the graphical evidence in appendix 2.A. Consider first the overall effect of extended UI duration on non-employment hazards. The ratio of the hazard rates is close to one and not statistically significantly different from one in the first 28 weeks. The largest effects in terms of hazard ratios is estimated in the period from 29 to 31 weeks: The hazard ratio drops to 0.876, meaning that the hazard rate of the treated is on average around 12 % lower than the one of the untreated during that period. This results shows that job seekers with 30 weeks of benefits exit non-employment significantly more than their counterparts with 39 weeks of benefits in the period from 16 to 28 weeks and even more so in the period from 29 to 31 weeks. The strong negative effect in the period from 29 to 31 weeks captures the spike at benefit exhaustion for job seekers with 30 weeks of benefits. The ratio of the hazard rates is also around 0.87 in the period between 32 and 37 weeks, and around 1.05 in the period between 38 and 40 weeks. This positive effect captures the spike around the benefit exhaustion for job seekers with 39 weeks of benefits. After 40 weeks, the ratio of the hazard rates is again close to one and not significant.

Considering the decomposition of effects into exits to wage-improving and wage-declining jobs (columns 2 and 3), which are modeled as independent competing risks, the theoretical predictions of section 2.2 are supported: There are no large differences between the hazard rate responses of the two exit destinations in the first interval from 0 to 15 weeks. The ratio of the hazard rates for losing job seekers slightly exceeds that of the winning job seekers in the period from 16 to 28 weeks, but differences are small. In the period between 29 and 31 weeks, the ratio of the hazard rates of losing job seekers amounts to 0.875, whereas the ratio of the hazard rates of winning job seekers is around 0.893. Thus, the exit hazard to wage-declining jobs with 39 weeks of unemployment benefits is only around 87 % of the exit rate of their counterparts with 30 weeks of benefits during the period from 29 to 31 weeks. For exits to wage-improving jobs, the corresponding measure amounts to around 89 %. The differences in hazard ratios of the two exit destinations disappear in the later periods. These findings suggest that - especially in the "treatment" period from 29 to 31 weeks - exits to wage-declining jobs are somewhat stronger affected by the UI extension in terms of hazard rates.[12]

Thus, also when looking directly at exit rates, the theoretical hypothesis in section 2.2 is supported: Theory predicts stronger effects for exits to wage-declining jobs, because the exit rate to wage-declining jobs is composed of both a search effort component and a reservation wage component, whereas the hazard rate to wage-improving jobs is determined only by the search effort channel. Although the differences between exits to wage-declining and wage-improving jobs are not large and significantly different from each other in terms of hazard rates, these differences aggregate up to considerable effects for survivor functions and average non-employment, as documented in section 2.5.3.

[12] Results are more pronounced for *registered* unemployment hazards and are available upon request.

Table 2.5: Job exit hazard ratio estimates

	Overall	Winning	Losing
Treatment effect in the interval from ...			
... 0 to 15 weeks	0.986	0.994	0.989
	(0.015)	(0.023)	(0.022)
... 16 to 28 weeks	0.989	0.992	0.998
	(0.023)	(0.037)	(0.030)
... 29 to 31 weeks	0.876**	0.893	0.875**
	(0.045)	(0.078)	(0.057)
... 32 to 37 weeks	0.876***	0.868*	0.890**
	(0.039)	(0.064)	(0.050)
... 38 to 40 weeks	1.056	1.040	1.073
	(0.067)	(0.112)	(0.086)
... 41 weeks and more	0.974	0.994	0.979
	(0.032)	(0.056)	(0.039)
Obs.	461,162	461,162	461,162
Individuals	132,134	132,134	132,134

Notes: This table presents RD estimates of the job exit hazard ratios. Treatment effects are estimated for different intervals and account for unobserved heterogeneity using a gamma distributed shared-frailty model. Estimates are calculated using a linear specification in age and include a set of control characteristics. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

Relaxing the underlying assumptions. A last sensitivity check is concerned with the theoretical assumption that job seekers set reservation wages strictly below their previous wage. Even though there are several theoretical reasons why one should expect wage reductions after an involuntary job loss, existing empirical evidence suggests that a subset of job seekers set reservation wages above previous wages. Feldstein and Poterba (1984) examine the ratio of reservation wage to previous wage by duration of non-employment and find that the ratio is slightly above 1 in the first 25 weeks of the spell. Using a large set of self-reported reservation wages of unemployed job seekers, Krueger and Mueller (2011b) provide empirical evidence that on average reservation wages are essentially the same as previous wages. However, they report a substantial variability across workers: The 25th percentile reservation wage ratio is 0.7, the median is around 0.91, and the 75th percentile is 1.17. The theoretical implication is, that for the subset of job seekers who set reservation wages above previous wage, the optimal response to extended UI durations might be determined - like for exits to wage-declining jobs - by both the reservation wage and the search effort channel even if they exit to wage-improving jobs.

In order to account for the variability in reservation wage choices, the assumption that $\rho_t < w_0$ for all job seekers is relaxed by multiplying the previous wage by a scaling factor α, so that $w_0^* = \alpha w_0$. By varying the scaling factor α, the probability that job seekers set reservation wages above previous wages can be controlled. Increasing α reduces $P(\rho_t \geq w_0^*)$

and the subset of job seekers setting reservation wages above w_0^* decreases. By doing so, exits to wage-improving jobs should be influenced less and less by job seekers whose reservation wage choices are potentially binding.

Table 2.6 presents RD estimates for different scaling factors. Columns 1 and 2 present estimates using a scaling factor of 0.85, columns 3 and 4 replicate the estimates from table 2.3 using a scaling factor of $\alpha = 1$, and columns 5 and 6 use a scaling factor of $\alpha = 1.15$. A scaling factor of 1.15 for example means that only job seekers with reemployment wages above 1.15 times previous wages are classified as exits to wage-improving jobs. Thus only job seekers who exit to jobs with considerable wage improvements relatively to their previous jobs are classified as winning job seekers and reservation wages are unlikely to have played an important role for them. Accordingly, a scaling factor of 0.85 means that for a relatively large fraction of job seekers who are classified as exiting to wage-improving jobs, reservation wage choices might have been played a role. The estimates in table 2.6 support these hypotheses: In columns 1 and 2, for a scaling factor of $\alpha = 0.85$ exits to wage-improving and exits to wage-declining jobs contribute both in the same order of magnitude to the overall non-employment effect. This can be explained by the fact that both wage-declining and wage-improving exits are similarly affected by both reservation wage and search effort considerations. In line with the theoretical predictions, columns 5 and 6 show that by lowering the share of job seekers who exit to wage-improving jobs the decomposition between winning and losing exits becomes even more clear cut and virtually all of the effect on non-employment comes from exits to wage-declining jobs. Thus, for increasing values of α the gap between wage-declining and wage-improving non-employment estimates tends to widen and confirms the theoretical hypothesis.

At the same time, the set of job seekers who exit to wage-improving jobs is decreasing in α. So there is a inevitable trade-off with this kind of sensitivity test: On the one hand, increasing α helps to get an empirical specification which is closer to the theoretical predictions in section 2.2. On the other hand that specification is estimated with less and less observations in the group of winning exits. If α is set to one (the baseline case), around 44 % of the job seekers exit to wage-improving jobs (see figure 2.2). If α is set to 1.15, only around 15 % of spells are classified as exits to wage-improving jobs, and if α is set to 0.85 around 70 % of spells are classified as exits to wage-improving jobs. Nevertheless, this sensitivity analysis provides a valuable test which highlights once more the importance of the reservation wage channel for the job search behavior.

Table 2.6: RD estimates using threshold scaling

	Winning	Losing	Winning	Losing	Winning	Losing
	$\alpha = 0.85$	$\alpha = 0.85$	$\alpha = 1$	$\alpha = 1$	$\alpha = 1.15$	$\alpha = 1.15$
D_i	0.304*	0.311	0.018	0.597**	-0.039	0.653**
	(0.177)	(0.230)	(0.149)	(0.285)	(0.130)	(0.261)
Bandwidth	3.456	3.456	3.456	3.456	3.456	3.456
Observations	177,821	177,821	177,821	177,821	177,821	177,821

Notes: This table presents decomposed RD estimates into exits to wage-improving (winning) jobs and exits to wage-declining (losing) jobs for different sacling factors α. Standard errors clustered by age in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

2.6 Conclusion

Reservation wages and search effort are the key drivers of job search behavior. The finding that extended UI generosity prolongs non-employment duration is one of the most robust results in labor economics. However, there is an ongoing discussion about whether extending UI benefits has beneficial effects in terms of job match quality or whether it only subsidizes unproductive search. To contribute to this discussion, separating the two behavioral margins of job search is essential.

This paper developed a novel approach which allows to make inference about the relative importance of reservation wages for job search without actually observing reservation wages directly. I show that the exit hazard to jobs can be decomposed into two exit destinations - a component due to exits to wage-improving jobs and a component due to exits to wage-declining jobs. This decomposition allows to isolate the reservation wage channel from the search effort channel, because reservation wages should not affect the likelihood of exits to wage-improving exits, whereas search effort affects wage-improving and wage-declining exits likewise. If reservation wages matter, extended benefit durations should affect exits to wage-declining jobs more than exits to wage-improving jobs.

The empirical analysis relies on a quasi-experimental variation in benefit duration in Austria that allows to identify causal effects of extended benefit duration on non-employment duration and survivor functions. Overall, extending PBD by 9 weeks prolongs non-employment by around 0.6 weeks. Analyzing the effect of increases PBD on survivor functions reveals that the maximum contributions to the prolonged non-employment duration are observed between 30 and 39 weeks, the period between the two benefit exhaustion dates. The decomposition of the overall effects into contributions from exits to wage-improving and wage-declining jobs shows that the major contributions come from exits to wage-declining jobs. Exits to wage-declining jobs contribute to virtually all of the total effect on non-employment, with its largest effects in the time period from 30 to 39 weeks. Investigating job finding hazards directly while controlling for unobserved heterogeneity points in the same direction: I find larger responses in terms of exit hazards to wage-declining jobs. Absent a direct effect of extended PBD on the wage offer distribution, the differential effects between exits to wage-improving and exits to wage-declining jobs are likely to be driven by the reservation wage choices of job seekers who exit to wage-declining jobs.

The empirical results show that reservation wages play an important role for job search behavior. This has potentially important implications for the optimal provision of UI and for the analysis of welfare. In an influential paper, Shimer and Werning (2007) showed that a worker's after-tax reservation wage encodes all the relevant information about his welfare. Extending benefits can be beneficial for aggregate welfare if pre-tax reservation wages are very responsive to unemployment benefits. Moreover, if reservation wage considerations play a role for job search, policies that help job seekers to form realistic beliefs about expected wage offers could help them to make informed reservation wage choices. This could reduce average non-employment duration and potentially leads to positive effects on total welfare.

2.A The effect of extended UI duration on unemployment and job finding hazards

Figure 2.A1 provides first descriptive evidence of the effect of extended UI duration on unemployment hazards and job finding hazards for the first 78 weeks of the unemployment spell. The figure shows the empirical unemployment hazard rates and the job finding hazards for a group of job seekers with 30 weeks of benefits (dashed red line) and for a group of job seekers with 39 weeks of benefits (solid blue line) respectively. For the exit rates from unemployment, clear spikes around benefit exhaustion are discernible. But analyzing job finding rates rather than unemployment exit rates shows that spikes nearly disappear. This suggests that a considerable part of the spikes around benefit exhaustion is due to mechanical de-registration from UI after benefit exhaustion rather than from exits to jobs. Card et al. (2007c) discuss this issue in more detail.

Figure 2.A2 shows the regression-discontinuity estimates of extended UI duration on (discrete) unemployment hazard rates. The unemployment hazard rate is estimated as the probability of exiting unemployment in period t given that unemployment lasted until period t. The effect of extending UI duration by 9 weeks is estimated for each period $t \in \{0, \ldots, 78\}$. Subfigure 2.A2a reports the estimated difference between the unemployment hazard rates of the treated and the untreated. Extending benefit duration does not affect hazard rates significantly in the first 26 weeks. After that, the difference in the hazard rates starts to drop below zero reaching a trough in week 30. The RD estimates highlight that job seekers with 39 weeks of benefits exit unemployment significantly less than their counterparts with 30 weeks. The trough in period 30 corresponds to the spike around benefit exhaustion for those job seekers with a maximum of 30 weeks of benefits. The estimated difference stays significantly below zero in the period from 30 to 38 weeks and switches to a significantly positive effect in the periods 39 and 40. The peak in period 39 captures the spike around benefit exhaustion for job seekers with 39 weeks of benefits. For durations beyond 40 weeks, the difference between the hazard rates fluctuates around zero. Because later periods contain less and less observations, estimates are less and less precisely estimated with increasing duration. Estimating effects rather for job finding rates (subfigure 2.A2b) instead of unemployment exit rates shows that the peaks around benefit exhaustion mostly disappear.

Figure 2.A1: Empirical hazard rates

Notes: Figure 2.A1 shows unemployment exit and job finding hazards for job seekers aged between 35 and 40 years old at registration (eligible for 30 weeks of benefits) and job seekers aged between 40 and 45 years old (eligible for 39 weeks of benefits).
Source: Own calculations based on ASSD.

Figure 2.A2: Estimated effects for hazard rates

(a) Unemployment

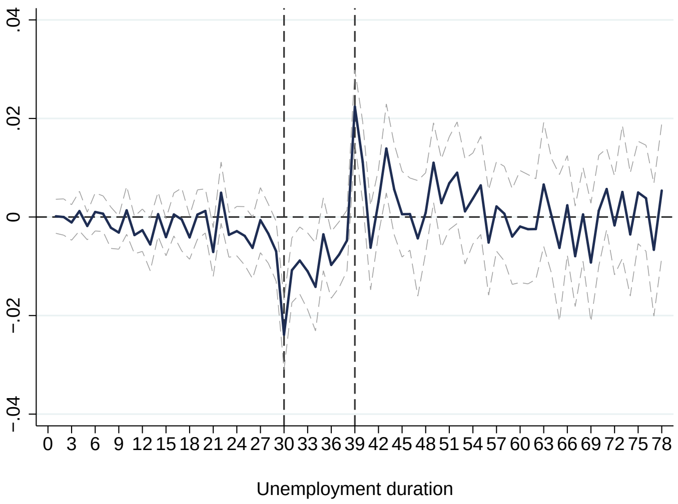

(b) Non-employment

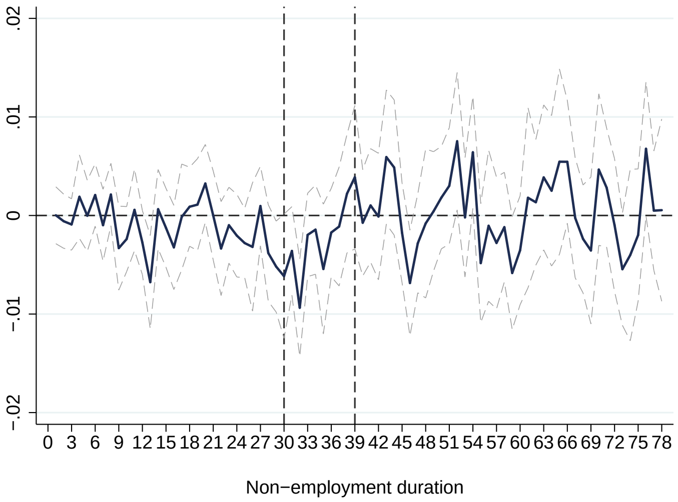

Notes: Figure 2.A2 shows estimated treatment effects for exit rates from unemployment.
Source: Own calculations based on ASSD.

2.B Additional figures

Figure 2.B1: Covariates around the age threshold

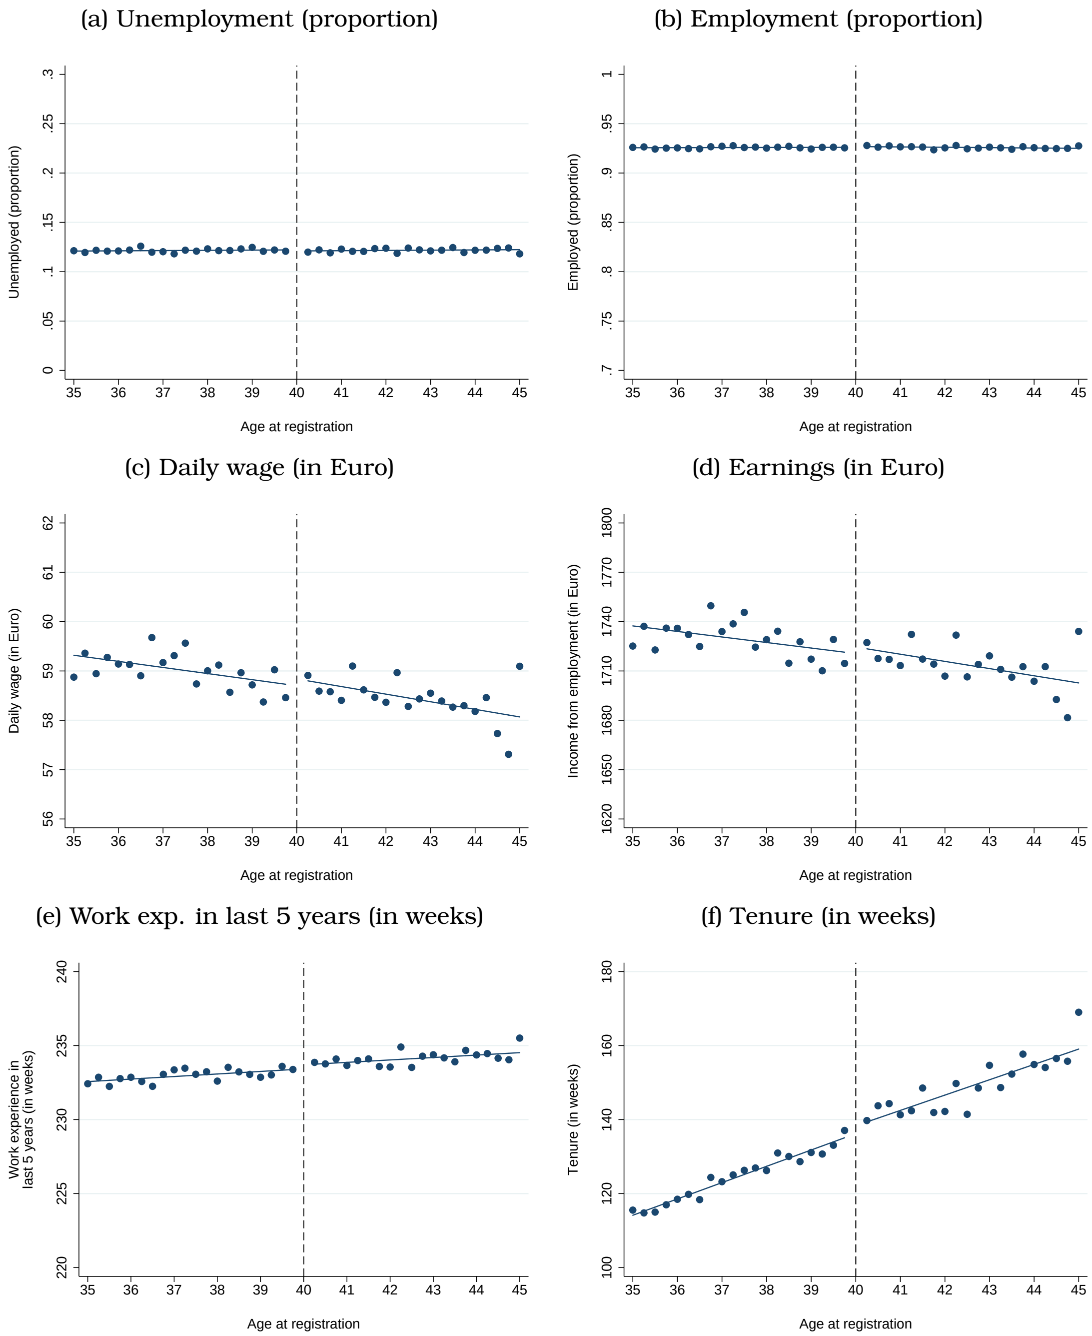

Figure 2.B1: Covariates around the age threshold (Continued)

(g) Elementary school (proportion)

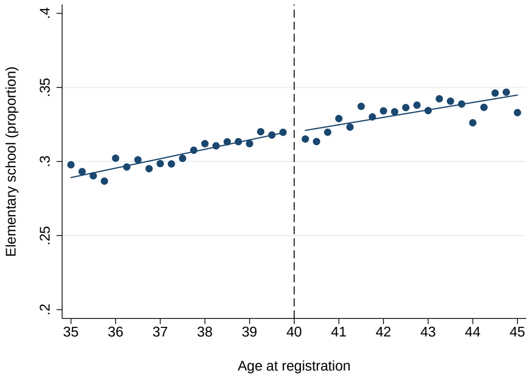

(h) Manufacturing sector (proportion)

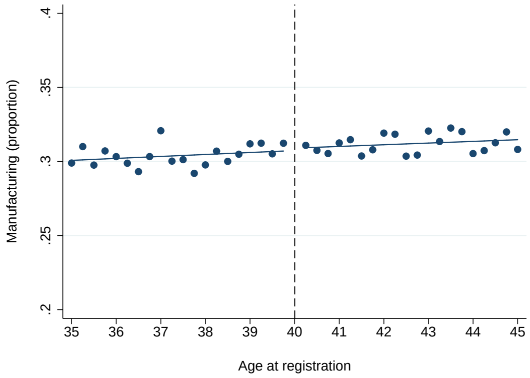

(i) Married (proportion)

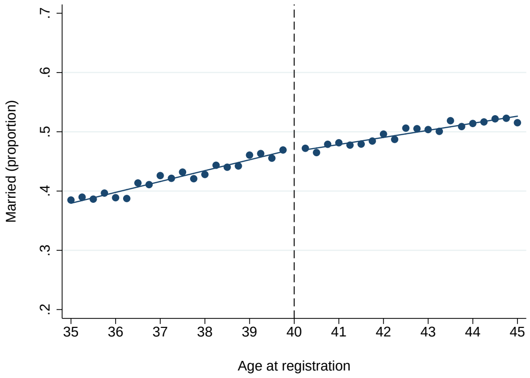

(j) Female (proportion)

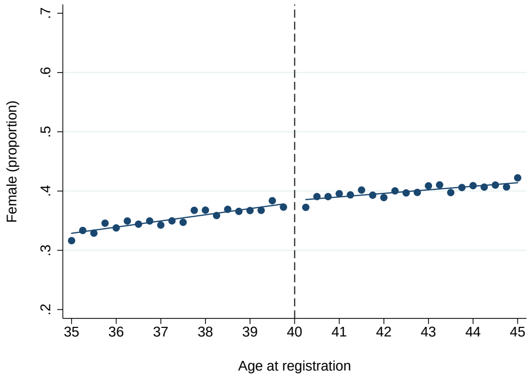

(k) Austrian (proportion)

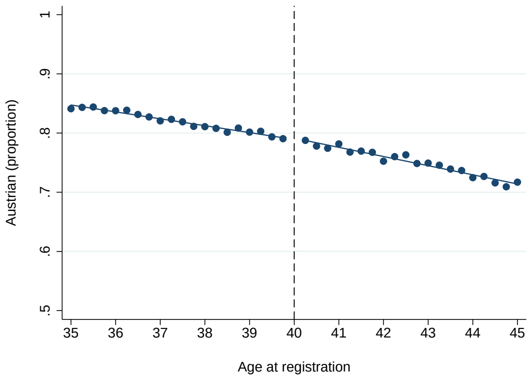

(l) Vienna (proportion)

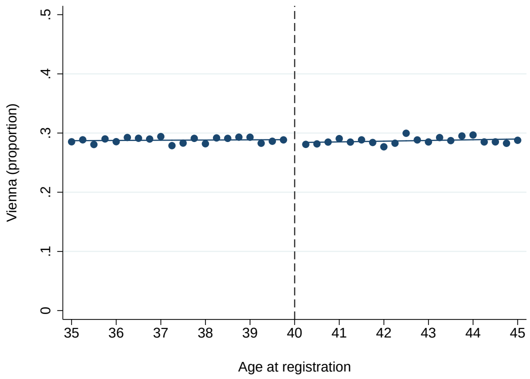

Notes: Figure 2.B1 shows the distribution of covariates as a function of age. The x axis shows age at the start of the unemployment spell. Each dot corresponds to a bin size of one quarter. A window of 5 years around the threshold is shown. The vertical line indicates the threshold value at 40 years old. Job seekers to the left of the threshold are entitled to 30 weeks of unemployment benefits, and job seekers to the right of the threshold are entitled to 39 weeks of benefits. *Source:* Own calculations based on ASSD.

Figure 2.B2: Probability of exit to wage-declining jobs over spell duration

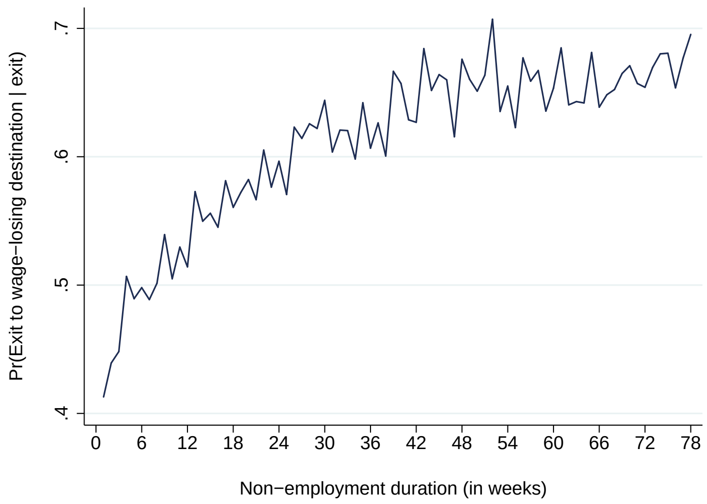

Notes: This figure shows how the proability of exiting to a wage-declining job conditional on exiting evolves over the duration of the spell.
Source: Own calculations based on ASSD.

Figure 2.B3: Average previous wages over the duration of the spell

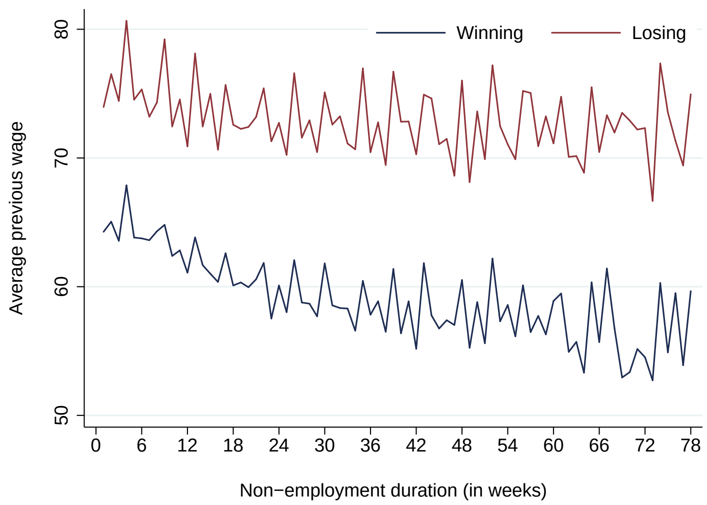

Notes: This figure shows average previous wages as a function of the duration of the spell for exits to wage-improving jobs and exits to wage-declining jobs separately.
Source: Own calculations based on ASSD.

2.C Additional tables

Table 2.C1: RD estimates with controls

D_i	0.660***
	(0.232)
$(X_i - c)$	0.282***
	(0.065)
$D_i(X_i - c)$	0.223**
	(0.111)
Mean past earnings	0.010***
	(0.001)
Mean past wage	-0.322***
	(0.030)
Past unemployment spell	0.050
	(0.227)
Work exp. in past 15 years (in weeks)	-0.008***
	(0.001)
Tenure (in weeks)	0.008***
	(0.001)
Severance pay	2.475***
	(0.251)
Female	2.490***
	(0.175)
Austrian	0.501***
	(0.184)
Married	-3.161***
	(0.142)
Education (reference group is elementary school)	
Less than elementary school	3.033***
	(0.470)
Apprenticeship/High School	-0.769***
	(0.134)
University	2.022***
	(0.538)
Other	-3.515***
	(0.217)
Industry (reference group is manufacturing)	
Wholesale and retail trade	-2.017***
	(0.168)
Financial, insurance activities, extraterritorial bodies	-1.414***
	(0.170)
Transportation	-4.633***
	(0.266)
Health and social activities	-2.126***
	(0.256)
Other	-2.138***
	(0.249)
Industry (reference group is Vienna)	
Lower Austria	-3.928***
	(0.207)
Upper Austria	-6.074***
	(0.166)
Burgenland	-4.681***

Table 2.C1 – continued

	(0.464)
Carinthia	-6.613***
	(0.267)
Salzburg	-8.668***
	(0.233)
Styria	-3.629***
	(0.258)
Tyrol	-10.998***
	(0.257)
Vorarlberg	-3.867***
	(0.393)
Unknown	-8.409***
	(0.409)
Constant	37.477***
	(0.719)
Bandwidth	3.456
Obs.	177,800
Clusters	83

Notes: This table shows the RD estimates for non-employment duration including all control variables. Estimates are calculated using a polynomial p of order 1 and the optimal bandwidth B. Standard errors clustered by age in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on ASSD.

Chapter 3

How Can Consumers Use Electricity More Efficiently? Exploring the Role of Information*

Abstract

This paper describes the results of a randomized field experiment which analyzes the electricity saving potentials of information. We consider three types of information: (i) real-time feedback on one's own electricity consumption using smart metering technology, (ii) personalized electricity savings tips through expert advice, and (iii) social information about one's own and a peer households' electricity consumption. Real-time feedback through smart meters reduces electricity consumption by 3 to 5 % of daily electricity consumption. Effects are detected shortly after installation of the in-home displays and persist over the study period. Moreover, we find that the largest savings are realized in peak hours. At the same time, households substitute part of their electricity consumption towards low-tariff hours. Social information reduces electricity consumption by around 1.5 % of daily consumption as long as feedback is frequent enough. Expert advice improves the perception of how easy it would be to improve energy-efficient behavior, but fails to translate into electricity savings. Utility is unaffected by the treatments indicating that the social benefits of improved information may be offset by negative social pressure effects or the costs of behavioral changes.

JEL Classification: Q48, Q4, C93

Keywords: electricity consumption, field experiment, smart metering, social comparison

---

*This chapter is jointly written with Lorenz Goette and Rafael Lalive. Contribution: Research design (jointly with Lorenz Goette and Rafael Lalive), data preparation, empirical analysis, writing. I would like to thank the steering committee and the project team of *ewz* for the implementation and coordination of the field study and for numerous helpful comments. Financial support from the Swiss Federal Office of Energy (SFOE) project SI/500475-01 is gratefully acknowledged.

3.1 Introduction

Climate change is one of the most important economic policy challenges of the 21^{st} century. Negative environmental externalities of energy consumption such as greenhouse gas emissions and the risks of nuclear energy - which have been forcefully demonstrated at the nuclear incident of Fukushima in 2011 - reinforced the need of policies that foster energy conservation. Exploitation of existing and development of new energy efficiency potentials have recently become top priorities on the agenda of many OECD countries. Major new energy efficiency programs are being implemented or discussed, such as the Energy Efficiency Directive in the European Union, the Clean Energy Future package to exploit the remaining energy efficiency potential in Australia, or tightened appliance standards in the United States (IEA, 2013). In the Swiss context, reducing CO_2 emissions and phase-out of nuclear power in the medium run are identified as main challenges in the Energy Strategy 2050. To ensure a secure energy supply, the Swiss government commits to concentrate on renewable energy sources and hydroelectric power and to foster energy-efficiency.

Researchers and policymakers have mainly focused on relative prices as main driving force of energy use. However, pricing instruments aimed at regulating the energy demand have often been found to be short-lived and difficult to enforce politically (Costa and Kahn, 2013; Cragg et al., 2013). What is more, pecuniary incentives can completely undo the intended effect of an intervention if they distort the intrinsic motivation to save energy.[1] This is why recently the focus of energy conservation programs has shifted towards behavioral approaches - such as the provision of information or social norm feedback. These are potentially effective and relatively cheap tools to increase energy efficiency (Allcott and Mullainathan, 2010; Gardner and Stern, 2008; Griskevicius et al., 2008). The recent focus on information as a policy tool to conserve energy is also reflected in the Energy Efficiency Directive of the European Union: the directive comprises a commitment for easy and free-of-charge access to data on current and past energy consumption through individual metering, empowering consumers to better manage their energy consumption (European Commission, 2012).

This paper describes the results of a randomized controlled field experiment that was designed to evaluate the electricity saving potentials of different types of information. A *first* group of participants - the *smart metering group* - received detailed real-time feedback about their own electricity consumption using smart metering technology. Each household in this group got an in-home display that provided real time feedback about current and past electricity consumption. A *second* group of households - the *expert advice group* - got an invitation to receive personalized electricity savings tips through an energy expert. In an about one-hour consultation at the service center of Zurich's energy supplier *ewz*, households could learn about how to improve energy efficiency in their own households. A *third* group of households - the *social competition group* - received social information about one's own and a peer household's electricity

[1] For evidence from other contexts see for example Frey and Oberholzer-Gee (1997); Mellström and Johannesson (2008); Ariely et al. (2009); Gneezy and Rustichini (2000)

consumption. Household pairs were matched based on pre-treatment electricity consumption. Both members of each matched pair periodically received electricity statements with their own and their peer household's average per-day consumption over the last period and they were encouraged to use less electricity than their peer household over the whole treatment period of one year. A *fourth* group of participants - the *social comparison group* - also received periodical statements about their own and their peer household's electricity consumption with the goal to consume less electricity than their peer household by the end of the treatment period. The only difference between social competition and social comparison is that in the latter peers were drawn from the control group, who did not receive the statements nor did they know that their electricity consumption was disclosed to the other household.

The empirical analysis provides an impact evaluation of these different types of information on short- and medium run electricity consumption. Impacts vary substantially across treatments and over treatment periods. Detailed real-time feedback through in-home displays triggers the largest electricity savings with effects varying from 1.5 to over 4 % of daily consumption over the treatment period. An in-depth analysis of the smart metering treatment shows that electricity conservation starts already shortly after installation of the in-home displays and persists over the whole treatment period. Breaking the effects down to the hourly level highlights that while there is some electricity conservation over the whole day, the bulk of the reductions is realized in peak hours. What is more, households increasingly substitute electricity consumption from high-tariff to low-tariff hours. Social competition generates savings on the order of 1.5 % of daily consumption. Effects are however only significant as long as feedback is frequent enough. The social *comparison* treatment generates effects which are roughly half of the social competition effects and not significant. However, we find highly asymmetric effects when splitting the pairs into households who used more than their peers and households who used less than their peers. Households who cumulatively used more electricity than their peers in the previous periods reduce their electricity consumption by up to 4 % in the medium-run, whereas the "better" household stops saving electricity after the first quarter. Expert advice does not affect electricity savings, but improves the perception of how easy it would be to improve energy-efficient behavior.

This paper is related to an early literature in environmental psychology which focuses on behavioral interventions aimed at household energy conservation. Various interventions such as commitment, goal setting, information or consumption feedback are evaluated and findings are mixed. Information tends to have positive effect on knowledge levels but does not necessarily change behavior, whereas rewards have effectively encouraged energy savings but only in the short run. Feedback triggers energy conserving behavior especially when given frequently. Many of these studies are however based on small samples and solely focus on short-run effects. Moreover it remains often unclear how the interventions impact onto actual energy use. Nevertheless they do show proof of concept that feedback about one's own or other people's behavior could be a potentially efficient and cheap way to foster energy-efficiency.[2]

A strand of related behavioral economic literature is interested in the power of social norm information on residential energy conservation. The power of social norm information on one's own behavior has been documented in various contexts, such as charitable giving (Frey and Meier, 2004), voting (Gerber and Rogers, 2009), retirement savings (Duflo and Saez, 2003), curbside recycling (Schultz, 1999), the reuse of towels in hotels (Goldstein et al., 2008) or the diffusion of solar panels (Bollinger and Gillingham, 2012). Ferraro and Price (2013) examined a large-scale program aimed at reducing residential water demand. They show that social comparison messages had a greater influence on behavior than simple pro-social messages or technical information. However, the effectiveness of the messages waned over time.

A number of recent studies evaluate large scale randomized policy programs in the United States aimed at residential energy conservation. At the core of these policy programs is the distribution of social norm information through Home Energy Report letters that compare a household's energy use with that of similar neighbors. Typically, these letters were augmented with personalized energy savings tips. Allcott (2011) evaluates a series of these programs and finds that the average program reduces residential energy consumption by 2 %. To combat "boomerang" effects, i.e. to avoid households who have been overestimating their energy consumption to increase the unwanted behavior, the Home Energy Reports not only included descriptive norms about a similar neighbor's energy use, but were augmented with injunctive norms, such as smiley emoticons, which convey that energy conservation is pro-social. Despite the use of injunctive norms in the Home Energy Reports, Ayres et al. (2013) find some evidence for the "boomerang" effect in their analysis of two large-scale randomized field experiments.

Allcott and Rogers (forthcoming) evaluate the short-run and the long-run effects of behavioral interventions. They document almost immediate responses after sending out the initial few Home Energy Reports, which however decay rapidly in the months between the reports. This cyclical pattern of action and backsliding however attenuates after the first few reports and what is left is a durable treatment effect. Repeated interventions help individuals to build up a new capital stock, such as exchanging traditional light bulbs for energy-efficient ones, or to change their energy use habits, which makes the effect persistent.

Dolan and Metcalfe (2013) examine a natural field experiments in the United Kingdom aimed at separating the impact of social norms from an information component. Complementing social norms with information about energy savings possibilities has large positive effects in the short-run. Over the long run however, the social norms only treatment was equally effective than the social norms with information treatment. In a second field experiment Dolan and Metcalfe (2013) investigate the use of online versus offline methods as a way of conveying information and find that sending social norm information through email is less effective than through letters. This reflects the importance of how salient a message is to the consumer, which has been shown in other contexts such as taxation (Chetty et al., 2009), financial markets

[2] See Abrahamse et al. (2005) for a survey of this early literature.

(Barber et al., 2005), or up-front appliance costs vs. subsequent electricity costs (Hausman and Joskow, 1982). See DellaVigna (2009) for a detailed review of that literature.

This field experiment complements the existing literature on the role of feedback for electricity conservation in several ways: *First*, our study design distinguishes between different types of information. Most existing programs study interventions which encompass a mixture of different types of feedback - typically coupling some form of social norm information with personalized energy savings tips. Understanding the role of information on electricity consumption however requires a clear distinction between different types of information. *Second*, exploiting high-frequency smart meter data allows us to evaluate for the smart metering group how intensively the in-home display is used, at what time of the day the reductions are typically realized, and how these patterns of reductions evolve over time. *Third*, we can link the data on electricity consumption to very detailed information on participant's demographics, personality traits and attitudes which we surveyed before, during and after the treatment period. *Fourth*, we provide a measure for consumer welfare using different dimensions of satisfaction as proxies for the utility of a consumer. Impact evaluations which are solely based on the administrative cost effectiveness of interventions are incomplete if they fail to account for the social welfare effects. *Finally*, observing treatments at different feedback frequencies allows us to learn about to what extent the feedback frequency plays a role.

The remainder of the paper is structured as follows: Section 3.2 gives an overview of the experiment. It discusses the possible mechanisms through which the interventions could act, presents the experimental design, and describes the data. Section 3.3 provides selective descriptive statistics and section 3.4 presents an empirical analysis of the four treatments based on periodical electricity readings. Section 3.5 discusses the smart metering treatment more in detail using high frequency smart meter data. Section 3.6 discusses cost effectiveness and social welfare implications of the interventions and section 3.7 provides summary and implications of our findings and discusses avenues for future research.

3.2 How does information affect electricity usage?

Subsection 3.2.1 illustrates in a simple conceptual framework the potential mechanisms through which the information could affect electricity consumption. Subsection 3.2.2 presents the experimental design and subsection 3.2.3 discusses the data.

3.2.1 Conceptual framework

To understand through which mechanisms the treatments could act, we consider a simple conceptual framework in the spirit of Levitt and List (2007). Consider an agent with a utility function that is additively separable over two terms: a positive consumption utility from electricity services c_i and a moral cost component M_i, which formalizes the non-pecuniary costs of electricity use. Electricity services can represent for example warmth or the electric power needed to cook or to watch television. The moral costs reflect the negative environmental externalities of electricity consumption, such as greenhouse gas emissions for example.

The consumption utility from electricity services is an increasing and concave function of an agent's electricity input e_i and a set of personal characteristics φ_i, that is $c_i = c_i(e_i;\varphi_i)$. The moral cost component of an agent is given by $M_i = M_i(\hat{e}_i, \tilde{e}; \varphi_i)$ and is a function of his perceived electricity input, $\hat{e}_i$, a norm electricity use, $\tilde{e}$, and a set of personal characteristics φ_i. Due to the public goods nature of electricity use, the moral costs component is assumed to be increasing and convex in the perceived electricity input $\hat{e}_i$ and decreasing in $\tilde{e}$. The utility function of an agent i is specified as

$$U_i = c_i(e_i;\varphi_i) - M_i(\hat{e}_i, \tilde{e}; \varphi_i)$$

The moral cost component $M(.)$ captures the non-pecuniary impact associated with an agent's electricity choice and can either reflect the intrinsic motives of an agent to conserve electricity or the social-reputational motives such as self-image concerns or the perceived social pressure to adhere to social norms.

What is more, due to limited attention agents fail to observe the true relationship between electricity input e_i and the consumption level of electricity services c_i.[3] Instead, agents only imperfectly observe the relationship between electricity input and consumption of electricity services: As electricity usage is not fully transparent, observing the true electricity input requires attention and knowledge which is limited and the electricity use of the chosen level of electricity services consumption is thus perceived as $\hat{e}_i = c_i^{-1}(\theta c_i; \varphi_i)$, with $\theta \in [0, 1]$. Whereas part of residential electricity consumption is immediately visible or perceivable, for example warmth, electric light or the electric power for the cooking stove, other components of electricity usage are less visible and more costly to observe, such as hidden stand-by electricity use. θ reflects the degree of knowledge and attention which is directed towards observing the true electricity usage. The larger θ, the closer the perceived electricity usage to the true electricity usage. However, if $\theta < 1$, then agents underestimate the electricity used to produce the electricity services, which leads to an over-consumption of electricity services.

In this simple framework we can identify different margins through which information could affect behavior. *First*, information raises attention towards the usage of electricity resources and improves knowledge about the less visible components of electricity usage. Feedback makes the electricity usage more *salient* and shifts the weight in the utility function more towards the moral cost component, making previous levels of electricity consumption no longer desirable. *Second*, the provision of information could affect electricity usage through *social pressure*. Social pressure increases the disutility of any prior consumption level and household reduce electricity consumption to offset the increased moral costs of electricity usage. *Finally*,

[3] Sallee (2013) presents a model of rational inattention for energy efficiency decisions. He argues that energy efficiency might rationally be ignored by agents for some categories of durable goods, such as automobiles or home appliances, because the costs of being fully informed may be substantial.

information could be a helpful tool to illustrate how energy efficiency could be improved at low cost and without reducing the privately optimal electricity consumption level.

The different mechanisms can have different implications for consumer welfare. On the one hand, interventions can increase individuals' utility if they positively resonate with an intrinsic motivation to save electricity.[4] Making one's own electricity usage salient may facilitate a utility-improving readjustment of one's privately optimal electricity consumption level. On the other hand, welfare gains may be completely undone if the behavioral changes have been coerced by social pressure. An intervention can then decrease utility and its desirability can be limited. In the context of charitable giving DellaVigna et al. (2012) showed that the welfare implication can be negative if an intervention affects behavior mainly through social pressure.

3.2.2 Experimental design

Study population and roll-out. The field experiment took place in Switzerland's largest city, Zurich, and was implemented and coordinated by Zurich's publicly owned energy supplier *ewz*. In total, 85,955 households were contacted and invited to participate in the study. Interested households were asked to register at the study web page and to give their consent to a data protection statement. After two weeks, households were reminded once about the study, which doubled the initial response rate. The overall response rate was 8.73 %. At the beginning of the treatment phase, the study population comprised 5,919 private households living in rented apartments or owner-occupied dwellings.

Over a total period of 30 months, the study was rolled-out weekly in 26 cohorts of 250 households each. Some meters were not accessible to meter readers and those participants had to be excluded already at the onset of the study. In the first 12 cohorts we lose a relatively large number of participants because of that: we observe only between 183 and 201 participants per cohort at the first reading. For the remaining 14 cohorts, households were already excluded from the potential study population before they were selected for participation if their meter was not freely accessible. For these cohorts, we observe between 234 and 247 participants per cohort at the first reading. The staggered roll-out of the interventions has several practical advantages. First, a sequential roll-out is less prone to implementation and coordination problems. Second, it avoids congestion and interference with *ewz*'s day-to-day business. Finally, the sequential roll-out also smooths out seasonal effects, such as vacation periods or Christmas holidays for example. Each household remained in the study phase for 16 months, and the additional 14 months are due to the sequential roll-out. The implementation of the study was tested with a pilot cohort which started in June 2010. The recruitment process for the first of the 26 actual study cohorts started in January 2011 and the last cohort terminated the study in December 2012.

[4] Note that we do not necessarily assume that households are intrinsically motivated to save electricity. If they aren't, their utility gain from saving electricity would simply be zero. In the context of our specific sample of participating households – a selection of households which voluntarily opted into the study – assuming intrinsic motivation to save electricity is however not unrealistic.

Treatment groups. The experiment consisted of one control group and four treatment groups. The control group (G0) serves to identify time effects over the study period. Except for the periodical surveys which have been sent to all study participants, there was no intervention in this group. The four treatment groups are the following:

- $G1$ *Smart metering:* This group of households received detailed and continuous feedback on their electricity consumption using smart metering technology. Households received in-home displays which allowed them to observe real-time and historical electricity usage, to measure electricity consumption in a specific measurement period and to set weekly consumption goals. Figure 3.C2 illustrates the different features of the in-home display.

- $G2$ *Expert advice:* This group of households received a one-time invitation for expert advice. Households were invited to *ewz*'s service center where they were provided with targeted and personalized recommendations on how to conserve electricity in the household.

- $G3$ *Social competition:* Each member in this group was matched to another member based on similar pre-treatment electricity consumption. The two members of each pair received feedback on the other's electricity consumption during the past month or quarter. Importantly, the feedback was bi-directional: both sides knew that the other would also see one's own consumption. Households were encouraged to compare and compete with the peer household and to use less electricity over the whole treatment period. Figure 3.C3 shows an example of the information given to a $G3$ household.

- $G4$ *Social comparison:* Each member was matched with a household from G0 and could observe one's own and the peer household's electricity consumption during the last month or quarter. Like in the social competition treatment, households were encouraged to compare and compete with the peer household. However, the feedback was one-directional: the households in G0 did not know that their consumption was shown to someone else, and did not see the other's consumption. Figure 3.C4 shows an example of the information given to a $G4$ household.

Timing of the study. Figure 3.1 shows the timing of the implementation for one cohort. The timing is defined relative to the last meter reading before the intervention start (T). The main intervention steps were the following: Upon invitation around 17 weeks prior to intervention start ($T-17$), interested participants could register on the study homepage. From the pool of registered households, we randomly selected each week around 250 households for participation in the study. Around 15 weeks before treatment start ($T-15$), households were randomly assigned to the smart metering group. The traditional meters of these participants were exchanged with smart meters that immediately started to record electricity consumption at a 15-minutes frequency. The remaining households were only informed about their participation in the study at this stage. One week before the start of the intervention (T), the remaining participants were assigned to one of the remaining treatment groups or the control group.

The treatment period started in week $T+1$. For the smart metering ($G1$) treatment, in-home displays were distributed which enabled households to continuously monitor their electricity use and to set weekly goals. The $G2$ group received a one-time invitation for an expert advice at the service center of *ewz* aimed at providing them with targeted recommendations on how to cut down on electricity usage in their households. The social treatment groups $G3$ and $G4$ received a first email informing them on their own and their peer households electricity consumption in the pre-treatment period. Over the course of the treatment phase, households in $G3$ and $G4$ received updated information about their own and their peer household's electricity consumption monthly during the first quarter, and then quarterly for the rest of the treatment period.

During the treatment period, *ewz* read electricity meters first at monthly frequency in the first quarter after intervention start, then at quarterly frequency for the rest of the treatment period, allowing us to observe the evolution of treatment effects over the one-year treatment period. What is more, we sent out comprehensive web surveys to all treatment groups and the control group prior to, during, and after the treatment period in order to assess whether treatments affected behavior, attitudes or knowledge of energy-relevant issues. Additionally, these surveys were used to collect a number of household, socio-demographic and personality characteristics of the participating individuals.

Figure 3.1: Timing of interventions

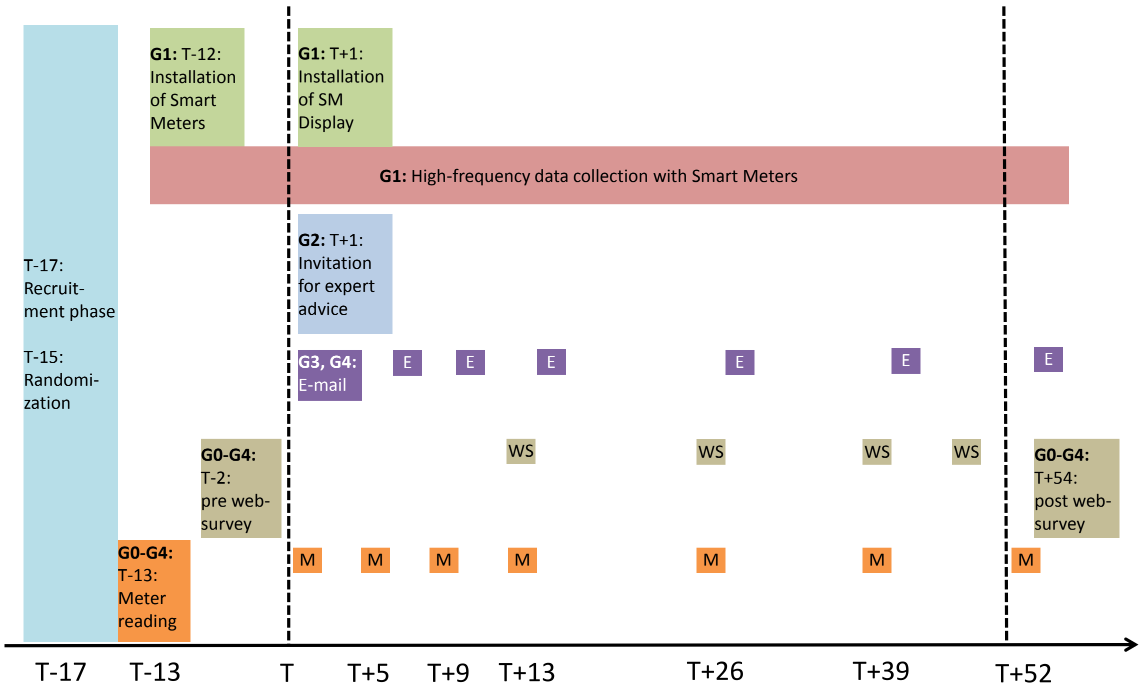

Notes: This figure shows the timing of the most important intervention steps of the study over the implementation phase.

Randomization and matching procedure. The randomized treatment assignment is based on a two-step procedure. In the first step, around 15 weeks prior to treatment start, we assigned households to the smart metering group. Specifically, within each cohort, participants were sorted according to their electricity consumption in 2009 and assigned to groups of five participants each. Within each of these groups of five, one household was randomly attributed to the smart metering group. This mechanism makes sure that the variance in electricity consumption between the groups is minimized but assignment to treatment is random within the groups of five. The remaining households were only informed about their participation in the study at this stage. In the second step of the randomization procedure - one week prior to treatment start - the remaining around 200 participants per cohort were assigned to the remaining treatment groups and the control group. Treatment assignment in this second stage is based on a minimization of the between-group variance of baseline electricity consumption in the quarter prior to treatment start. The remaining around 200 households per cohort were sorted in ascending order according to their average pre-treatment electricity consumption and grouped into brackets of four households each. Within each bracket, households were randomly selected into one of the remaining treatment groups G2, G3 or G4 or the control group G0.

Finally, based on a similar pre-treatment consumption in the quarter prior to the intervention start, households in group G3 were matched with a peer household from G3 and households in group G4 with a peer household from the control group G0. The rationale for using pre-treatment electricity consumption to match pairs is the following: Electricity usage depends on a number of observable household characteristics, such as apartment surface, heating system, number of household members, household income, or the type and number of appliances in the household, and a number of unobservable characteristics, such as the household-specific taste for electricity consumption. Matching pairs with respect to baseline electricity consumption is a simple way of summarizing observed and unobserved characteristics into one single dimension. The drawback of this type of matching is that two households with similar baseline electricity usage could nevertheless differ substantially in terms of other observables. If for example a large but extremely energy-conscious household is matched with a small but energy-wasting household, electricity usage could be comparable, but a comparison between the two members of the pair could be perceived as "unfair", because the former has exhausted all savings potentials already while the latter could easily conserve electricity. However, because participation was voluntarily, participating households are expected to be relatively homogeneous with respect to their energy-conscientiousness so that extreme examples as the above should be relatively rare. Figure 3.2 graphically shows match quality of the pairs in terms of baseline electricity consumption. On average, match quality is very good. The mean difference between a household's and its peer household's baseline electricity consumption amounts to 0.124 standard deviations. Because pairs had to be matched in cohorts of relatively small size, there is however also quite some variability in match quality.[5]

Figure 3.2: Match quality of household pairs

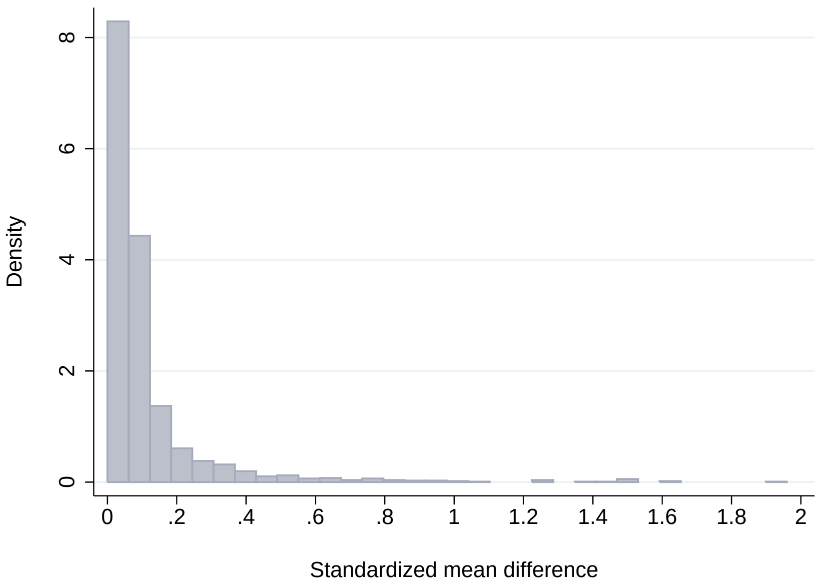

Notes: This figure shows the mean of the difference between a household and its peer household's baseline electricity consumption normalized by the standard deviation of baseline consumption.

3.2.3 Data

Periodical electricity readings. In addition to the yearly readings, electricity meters of the participants were read eight times during the 16 months lasting study period. The additional meter readings were integrated in the usual readings schedule of the meter readers. The meters were read off for the first time around 13 weeks prior to and for the second time around one week before the intervention start. Based on these two readings, we calculated average baseline consumption prior to treatment. After intervention start, meters were read off at a monthly rhythm - i.e. around 5, 9 and 13 weeks after treatment start - during the first quarter of the treatment period, and at a quarterly rhythm - i.e. around 26, 39 and 52 weeks after treatment start - for the remaining study period.

Due to congestion with the daily business of the meter readers, some readings - in particular for cohorts 8, 12, 16 and 20 - had to be advanced by up to three weeks. In order to guarantee comparability of the meter readings across households, we calculated average electricity

[5]In terms of other observable household characteristics, the match quality is ambiguous. The mean difference between household and peer household normalized by the standard deviation is 0.95 for apartment surface, 0.82 for the number of adult household members and 0.68 for the number of children in the household - differences that appear rather high. Because these statistics are based on categorical variables, we might classify households as relatively unequal in terms of the categorical variable, which might not necessarily be true for the underlying variable. A household with an apartment size of 65 m^2 would classify in category 3. If it is matched to a household with apartment size of 66 m^2, then the difference between the two households is relatively large in terms of the categorical variable, even though differences in the underlying variable are negligible.

consumption per day (kWh per day) between two different readings t and $t-1$ as follows:

$$\text{Electricity consumption (kWh) per day} = \frac{\text{reading (t) - reading (t-1)}}{\text{number of hours btw. t and t-1}} * 24$$

Besides the cumulative electricity level measure in kWh, the raw data contains time stamps with the exact date and time of the reading, tariff type, tariff class and the anonymous household identification number, which enables us to link the electricity data with the rich survey data. The tariff type defines a participant's power mix, which can contain different shares of clean energy sources. Households can either opt for one or for a mixture of different tariff types. Tariff class refers to the off-peak and the peak-hour tariff schemes. Peak-hours are from Monday through Saturday from 6 am to 10 pm. The remaining hours are off-peak hours with a more favorable pricing scheme. In order to calculate total electricity use, electricity consumption is aggregated across tariff types and tariff classes.

The quality of the raw data is very good. Less than 4 % of observations are non-standard and could potentially lead to measurement error. These non-standard readings mostly occurred when a meter was either defective, had to be changed or was untraceable. Those readings usually displayed an aggregate electricity level of 0 and we replaced them with missing values. In some cases per day electricity consumption was zero or even negative for other reasons such as meter changes or misreadings.[6] We replaced negative values with missing values. Zero consumption levels were checked case-by-case and replaced by missing values if the zero consumption was most likely coming from defective meters.[7] Finally, from the total number of readings around 3.4 % had to be discarded because the anonymous household identifier was missing.

High frequency smart meter records. In addition to the periodical electricity readings, we observe high-frequency electricity records for households in the smart metering treatment. Due to delivery delays, most in-home displays could not be installed at the scheduled time. Figure 3.3 illustrates the cumulative number of installations of the in-home displays over the study period. Only around 8 % of household received the in-home display in the first month after treatment start. Two months after scheduled intervention start, around one third of households had an installed smart meter device. After one quarter, around 80 % of households were equipped with a display. The remaining 20 % of in-home displays were mostly installed by the end of the second quarter. This shifting of the actual treatment start has important implications for the short-run analysis of the smart meter treatment using the periodical electricity readings.

For each household with an installed in-home display, we observe meter readings at a 15-

[6] In general, the level of the old meters were read off before the change and the level of the new meter was read off just after the installation. For some records, the old reading was however missing which resulted in a negative electricity consumption per day.

[7] Some of the misreadings could be corrected manually, if a confusion of digits was probable. This was the case for the digits *3* and *8*, or *2* and *7* for example. Moreover, some negative consumption levels could be corrected manually, if the time stamp of the meter reading was not in the correct order.

Figure 3.3: Installation of in-home displays

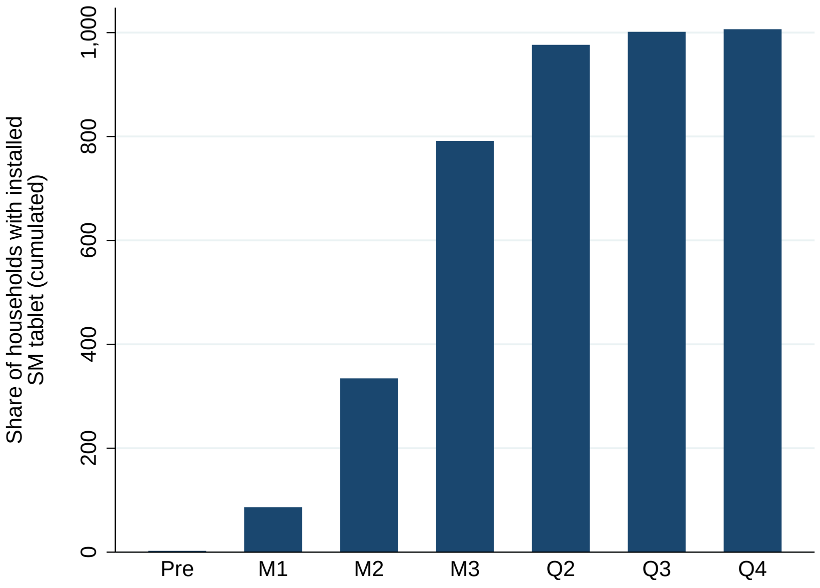

Notes: This figure shows the cumulative number of installations of the in-home displays over the study period.

minute frequency for both high- and low-tariff hours. Because the smart meters already started to record electricity data around 60 days before the installation of the in-home display, before-after comparisons within the smart metering group allow us to estimate short-run effects of smart meters and to analyze where savings come from over the day load profile. 3.5 % of the observations are recorded prior to the exchange of the traditional meter with the smart meter and were discarded. For 18 households, the installation date of the in-home displays was not recorded, but could be reconstructed from the log data of the display. Moreover, the smart meters failed to record the electricity data properly for nine households and had to be discarded.

In addition to the electricity records, we also recorded each activity with the in-home display, so that frequency and use of the display can be linked to the electricity records. We also observe whether G1 households have set weekly target levels (in kWh per week) and whether these targets were met.

Survey data. The survey data is based upon six web based questionnaires. The first web survey was sent to all participants two weeks before treatment start. It contains a broad range of socio-demographic characteristics, such as age, sex, nationality, education, income, the number of adults and children in the same household or whether participants are property owners or renting. Furthermore, a number of household characteristics such as apartment surface, type of heating and water heating system, as well as existence and use of electricity intensive appliances are assessed. Participants are also asked to what extent they already implemented electricity conserving measures in their household, such as the share of energy efficient light bulbs and the use of stand-by modes in appliances.

An important part of the surveys captures a wide range of personality measures, attitudes towards environmental issues, and energy-relevant knowledge.[8] Participants were also asked to give their best estimate for the electricity use of different appliances and the savings potential of different measures such as drying clothes on a clothes line instead of using the dryer for one load of laundry for example. A range of questions related to attitudes and energy-relevant knowledge are borrowed from the Swiss Environmental Survey 2007 (Diekmann et al., 2009) and a U.S. study that assesses public perceptions of energy consumption and savings (Attari et al., 2010). The personality measures are based on the 60-item version of the HEXACO (Ashton and Lee, 2009) and the short version of the big five inventory (Gosling et al., 2003). In a number of questions we assessed different measures of satisfaction in general, satisfaction with *ewz*, and satisfaction with the study.[9]

Participants were asked to repeatedly answer the same questions about attitudes, energy-relevant knowledge, use of appliances, and satisfaction over the course of the study period, which allows us to identify behavioral changes due to the treatments. The response rate of the pre-treatment survey was around 85 % and slightly decreased to roughly 75 % in the second and third survey, and to around 68 to 70 % in the fourth to sixth survey.

3.3 Descriptive statistics

Summary statistics. Table 3.1 summarizes baseline electricity per-day electricity consumption, a number of selected socio-demographic characteristics, and pre-treatment knowledge about energy-relevant topics.[10] Columns 1 to 5 show averages for each treatment group, column 6 depicts the overall average and column 7 contains p-values of a test of equal means across the groups for each characteristic.

Panel A of the table shows that baseline electricity consumption is well balanced between control and treatment groups and varies between 5.9 and 6.2 kWh per day across groups. For context, one kilowatt-hour is enough to run a standard 60-watt light bulb for 17 hours or to watch television for around seven hours for example. Most cohorts stated the treatment phase between May and September and therefore experience the high-electricity season in their second or third quarter of the treatment. Table 3.B1 in the appendix shows that average electricity consumption peaked in the second quarter. Average annual electricity consumption is around 2,300 kWh and lies around 300 kWh below average electricity consumption in Zurich (see Appendix B in Degen et al. (2013)) and much below the US average of 11,280 kWh (Allcott

[8]Three surveys also contained choice experiments, which we can use to elicit social preferences, participant's tendency to overestimate themselves and time preferences. See subsection 3.A for details about the choice experiments.

[9]The full list of the questions of all six web surveys (in German) can be found in Degen et al. (2013).

[10]Tables 3.B4 to 3.B6 present descriptive statistics for a more exhaustive number of household and individual characteristics. These tables confirm that the randomization of treatment assignment has worked very well with respect to a large number of characteristics.

and Rogers, forthcoming). A possible explanation for this could be that participation in the study was voluntary and might predominantly have attracted energy-conscious households.

Panel B summarizes a selection of socio-demographic variables: Participants are on average in their late 30ties and around 37 % of participants are female. The study participants are on average well educated and high-earners: Around 44 % of participants have a university degree and net household income is on average between 7,000 and 9,000 Swiss francs per month. The average household has around 2 members with an apartment size of around 80 m^2. Because single family house owners were excluded from the study, it is not surprising that almost 90 % of participants are renting an apartment and only 10 % are living in their own dwelling. The sociodemographic items are well balanced across treatment groups.

Panel C shows average attitudes and knowledge about energy-relevant topics. Mean attitudes is an average over ten items that cover participants' attitudes about environmental issues. Higher values indicate more energy-oriented values and attitudes. The average of all ten attitude dimensions is almost 3.8, indicating that participants have on average an environmental friendly attitude. Mean risk perception is the average perception of six dimensions of risk, such as the risks of nuclear power or genetic engineering for example. High values are associated with large perceived risks. With an overall average of around 3.5, participants perceive the six dimensions as relatively large risks. Mean energy-efficiency potential measures the average potential to change energy-efficient behavior in the household, such as drying clothes on a clothes line instead of using the dryer or increasing the temperature of the fridge by 1° C. Low values indicate that households can easily change their energy efficient behavior or already did it. High values indicate that changes are difficult or not possible. The average potential to change energy efficient behavior is 1.8. Participants thus perceive it as relatively easy to change their energy efficient behavior in the household already before the interventions. Moreover, participants answered on average 3.4 out of 7 questions related to electricity consumption of different appliances, and 1.1 out of 4 questions related to electricity conservation correctly, which is above what we would expect had they randomly ticked one of the answer options. The test of equality of means is rejected at the 1 % level for the mean energy-efficiency potential. Households in the expert advice group perceive it slightly more difficult and households in the social competition group slightly less difficult to change energy-efficient behavior. Furthermore, households in the smart metering group are slightly more environmentally friendly than their counterparts. Absolute differences are however small.

Finally, panel D summarizes a selection of satisfaction dimensions. High values stand for high satisfaction levels. Participants are very satisfied with their life and their quality of life and think that the study addresses an interesting topic. Satisfaction with the environment and the electricity prices of *ewz* is a bit lower but still at a relatively high level. Treatment groups are well balanced with respect to their satisfaction levels.

Table 3.1: Selected descriptive statistics

	G0	G1	G2	G3	G4	Total	p-value
			A. Baseline electricity usage				
Baseline electricity usage (kWh per day)	6.107	5.990	6.220	5.879	5.953	6.031	0.239
			B. Individual and household characteristics				
Age[a]	3.840	3.918	3.935	3.867	3.848	3.882	0.306
Female	0.361	0.365	0.386	0.368	0.383	0.372	0.726
Tertiary education	0.442	0.459	0.425	0.441	0.421	0.437	0.369
Household income[b]	6.574	6.389	6.311	6.179	6.374	6.366	0.098
Household size	2.128	2.064	2.073	2.082	2.080	2.086	0.732
Apartment surface[c]	4.332	4.323	4.289	4.272	4.248	4.293	0.512
Renting	0.888	0.890	0.900	0.905	0.906	0.898	0.581
			C. Attitudes and knowledge				
Mean attitudes[d]	3.802	3.862	3.777	3.792	3.815	3.81	0.012
Mean risk perception[e]	3.474	3.508	3.471	3.492	3.496	3.488	0.614
Mean energy-efficiency potential[f]	1.806	1.803	1.852	1.773	1.801	1.807	0.005
Knowledge (consumption)[g]	3.424	3.445	3.344	3.401	3.377	3.398	0.479
Knowledge (conservation)[h]	1.069	1.114	1.115	1.194	1.074	1.113	0.091
			D. Satisfaction[i]				
. . . with life	4.204	4.194	4.206	4.179	4.208	4.198	0.869
. . . with environment	3.502	3.484	3.519	3.523	3.485	3.503	0.741
. . . with life quality	4.230	4.232	4.214	4.204	4.211	4.219	0.893
. . . with *ewz*'s electricity prices	3.936	3.952	3.981	3.963	3.924	3.951	0.617
. . . with study: interesting topic	4.619	4.661	4.625	4.580	4.613	4.620	0.122

Notes: Columns 1 to 5 show averages of the items for groups $G0$ to $G4$. Column 6 shows the average across all groups and column 7 displays the p-values of an F-test of inequality of means across groups.
[a] 1 = 0-19 years; 2 = 20-29 years; 3 = 30-39 years; 4 = 40-49 years; 5 = 50-64 years; 6 = 65+ years
[b] 1 = $<$ CHF 3,000; 2 = CHF 3,000-3,999; 3 = CHF 4,000-4,999; 4 = CHF 5,000-5,999; 5 = CHF 6,000-6,999; 6 = CHF 7,000-7,999; 7 = CHF 8,000-8,999; 8 = CHF 9,000-9,999; 9 = CHF 10,000-11,999; 10 = CHF 12,000-14,999; 11 = $\geq$ CHF 15,000
[c] 1 = $< 15m^2$; 2 = 15-40m^2; 3 = 41-65m^2; 4 = 66-80m^2; 5 = 81-115m^2; 6 = $> 115m^2$
[d] 1 = "Do not agree" and 5 = "Agree"
[e] 1 = "No risk" and 5 = "Large risk"
[f] 1 = "Do it already"; 2 = "Easy"; 3 = "Rather difficult"; 4 = "Very difficult"; 5 = "Not possible"
[g] Maximum score is 7
[h] Maximum score is 4
[i] 1 = "Not satisfied" and 5 = "Very satisfied"
Source: Own calculations based on survey data.

Attrition and representativeness of the sample. We observe electricity consumption for 5,919 households at the beginning of the treatment phase. Table 3.B2 in the appendix shows the number of participating households across treatment groups and periods. In the pre-treatment period (Pre), we observe between 1,166 and 1,207 households per group. The number of observations drops to between 1,009 and 1,082 observations per group by the end of the intervention period (Q4). Attrition is slightly higher in $G1$ (14 %), $G3$ (13 %) and $G4$ (14 %) than in $G0$ (11 %) or $G2$ (10 %). Attrition can have different sources: The three most prominent sources of attrition are moving, technical problems with the meter readings and opting-out due to lost interest. Naturally, the two former sources of attrition are not caused by the treatment, whereas the latter possibly is. Unfortunately, the source of the attrition was recorded only for around a quarter of all drop-outs. Among these, around two thirds of drop outs are due to relocation, around 18 % were due to lost interest and 15 % had do be dropped because of technical problems with the meters, including households which had to be dropped because their meter was not freely accessible.

Table 3.2 compares the means of a number of selected characteristics for households who stayed in the study until the end of the treatment phase and for households who dropped out of the study at some point earlier. Panel A highlights that stayers had on average a higher baseline consumption than drop-outs. Clearly, households with higher baseline usage might also have larger savings potentials. If participants with lower per-day consumption dropped from the study because they believed their electricity conservation potential already to be exhausted and thus were no longer interested in participating in the study. Selection of this kind could positively bias our estimates.

Panel B shows how stayers and drop-outs differ with respect to a number of individual and household characteristics. Stayers are on average older than drop-outs, are more less to have a university degree, live in larger apartments, and are more likely to own their dwelling instead of renting it. These are all attributes which decrease mobility. In line with this pattern, Allcott (2011) found that moving households had smaller electricity consumption, were younger, lived in smaller homes, were more likely to rent and had lower incomes, whereas households who opted out had on average higher baseline electricity consumptions and were older. The observed difference in baseline electricity consumption between stayers and drop-outs is presumably driven to a large extend by differential characteristics of movers versus stayers. Selective drop-out of households who lost their interest in participating should therefore play only a second order role.

Panel C shows that stayers and drop-outs do not differ with respect to their attitudes and their knowledge. Drop-outs however perceive it on average slightly easier to change their energy-efficient behavior. This could indicate that on average, drop-outs are already more energy-efficient and lose the interest to participate in the study because they do not feel that they would acquire new information which could help them to conserve electricity. Absolute differences are however too small to bias the estimates.

Finally, panel D shows differences in satisfaction. Stayers and drop-outs do not differ with respect to their general satisfaction with life, the environment and the quality of life. Also, they do not differ with respect to their interest in the study. They do however differ slightly with respect to their satisfaction with *ewz*'s electricity prices. Drop-outs are less satisfied with the electricity prices than stayers, but again, absolute differences are small.

Table 3.B3 provides further evidence that selective drop-outs are not a first order issue. The table shows mean baseline electricity consumption for stayers and for drop-outs for all groups. If drop-out was selective due to the treatment, then we would not expect differences in baseline electricity consumption in the control group. We however observe differences in the same order of magnitude for the control group as for the treatment groups, which speaks against the presence of important selection effects.

Table 3.2: Attrition bias

Characteristics	Drop-outs	Stayers	Total	*p*-value
A. Baseline electricity usage				
Baseline electricity usage (kWh per day)	5.678	6.082	6.031	0.008
B. Individual and household characteristics				
Age	3.4	3.947	3.882	0.000
Female	0.388	0.37	0.372	0.409
Tertiary education	0.467	0.433	0.437	0.080
Household income	6.335	6.37	6.366	0.796
Household size	2.035	2.093	2.086	0.234
Apartment surface	4.012	4.331	4.293	0.000
Renting	0.952	0.89	0.898	0.000
C. Attitudes and knowledge				
Mean attitudes	3.816	3.809	3.81	0.772
Mean risk perception	3.455	3.493	3.488	0.153
Mean energy-efficiency potential	1.762	1.813	1.807	0.011
Knowledge (consumption)	3.387	3.4	3.398	0.825
Knowledge (conservation)	1.101	1.115	1.113	0.758
D. Satisfaction				
... with life	4.216	4.196	4.198	0.496
... with environment	3.46	3.508	3.503	0.169
... with life quality	4.193	4.222	4.219	0.386
... with *ewz*'s electricity prices	3.885	3.96	3.951	0.069
... with study: interesting topic	4.592	4.624	4.62	0.309

Notes: This table shows means of a number of individual and household characteristics across groups for stayers, drop-outs and overall. Stayers are participants who stayed in the study until the end. Drop-outs are participants who dropped out of the study at some point prior to study end.
Source: Own calculations based on electricity meter readings.

Table 3.3 analyzes the representativeness of the study population with respect to the population in the city of Zurich, the average of the six largest Swiss cities[11] and Switzerland as a

[11] The largest Swiss cities are Basel, Bern, Geneva, Lausanne, Winterthur and Zurich.

whole. Females are clearly underrepresented in the study population. Only around 37 % of the study participants are female compared to 51 % in Zurich, the largest Swiss cities and Switzerland. The study population is not representative of Zurich with respect to the age structure neither. Young persons from 0 to 19 years and the 65 years and older are underrepresented. The relative lack of young persons in the study might be due to individuals below 19 mostly still living at their parent's home. The relative lack of elderly participants might be explained by the fact that we required participants to have access to the Internet regularly, which might be less true for elderly persons. Consequently, the 30 to 39 years old, the 40 to 49 years old and the 50 to 64 years old are overrepresented among the study population. Because German was the only language of correspondence, language barriers might explain why Swiss and German citizens were overrepresented, whereas other nations were underrepresented.

Table 3.3: Representativeness of the study population

	Study population	Zurich	Largest Swiss cities	Switzerland
Gender (in %)				
Female	37.44	50. 71	51.61	50. 73
Male	62.56	49. 29	48.39	49. 27
Age (in %)				
0-19	0.18	16.20	17.19	20.87
20-29	13.78	14.54	14.80	12.79
30-39	29.89	20.60	18.38	13.90
40-49	23.08	15.72	15.43	16.33
50-64	21.19	16.18	16.99	19.22
65+	11.88	16.76	17.20	16.90
Nationality (in %)				
Swiss	81.65	69.50	67.17	77.6
German	10.94	7.78	5.06	3.35
Italian	1.45	3.49	4.10	3.65
Portuguese	0.15	2.16	3.27	2.7
Serbian	0.15	1.60	1.52	1.55
Other	5.66	15.47	18.87	11.16

Notes: This table shows averages of a number of socio-demographic characteristics for the study population, the population of Zurich, the average population of Switzerland's largest cities and Switzerland as a whole
Source: Own calculations based on survey data.

3.4 Analysis of the electricity readings

In this section we discuss the empirical results of the different treatments. Subsection 3.4.1 discusses identification and the empirical specification. Subsection 3.4.2 presents the short-run causal effects of the interventions and subsection 3.4.3 the medium-run effects. Subsection 3.4.4 discusses the issue of heterogeneous treatment effects with respect to a number of observed characteristics.

3.4.1 Empirical strategy

The empirical analysis is based on a randomized difference-in-differences design with four treatment groups and one control group. In order to discuss estimation and identification, let Y_{1i} be the electricity consumption of a treated individual after the start of the intervention, and Y_{0i} the electricity consumption of an untreated individual after the intervention. D_i is a dummy which equals one if an household was treated and zero otherwise. The observed outcome after treatment can be written as $Y_i = D_i Y_{1i} + (1 - D_i)Y_{0i}$. The average treatment effect on the treated is given by

$$E(Y_{1i} - Y_{0i} \mid D_i = 1) = E(Y_{1i} \mid D_i = 1) - E(Y_{0i} \mid D_i = 1)$$

Under the assumption of parallel time trends of treatment and control groups in absence of the treatment, the average treatment effect on the treated can be estimated in a difference-in-differences framework. Because selection into the treatment groups was random, there is no reason why this identification assumption should be violated. In table 3.1 and tables 3.B4 to 3.B6 we provided evidence that the randomization procedure worked well and that treatment and control groups are very similar with respect to a large number of characteristics and traits. From the above considerations, we develop the following primary empirical specification:

$$\Delta e_{it} = \beta_0 + \gamma' \mathbb{T}_{\mathbb{i}} + \beta_1' (e_{i0} \times \mathbb{Q}_{\mathbb{i}}) + \omega_{w(t)} + \varepsilon_{it} \tag{3.1}$$

where $\Delta e_{it} = e_{it} - e_{i0}$ denotes the change in electricity consumption of household i at measurement t relative to the baseline measurement prior to treatment. By considering changes in electricity consumption rather than levels we eliminate household-specific differences in the electricity level which improves estimation precision a lot. $\mathbb{T}_{\mathbb{i}}$ is a vector of the four treatment dummies, which takes the value of one for the treatment groups only in the after treatment period. e_{i0} is the baseline electricity consumption of individual i, $\mathbb{Q}_{\mathbb{i}}$ is a set of dummies for each quarter after treatment start, and $\omega_{w(t)}$ capture calendar week fixed effects. These week fixed effects take the value one if a week falls within a given measurement period t and zero otherwise.The inclusion of week fixed effects is not necessary for identification, because the staggered roll-out of the cohorts over time should already eliminate seasonal and other time effects which affect all treatment groups and the control group equally. However, including them can increase precision of the estimation results. β_0 measures the average change in electricity consumption in the control group, and β_1' estimates how electricity consumption grows as a function of baseline consumption e_{i0}. By interacting baseline consumption with a set of quarter dummies, we allow baseline consumption to affect electricity consumption growth differently over the different treatment periods. Again, the inclusion of baseline electricity consumption is not needed for identification, but can improve estimation precision somewhat. Finally, γ'

measures the average treatment effect on the treated. The treatment effects are estimated by ordinary least squares (OLS) and we take into account potential heteroskedasticity and clustering on the household level by using cluster-robust standard errors.

3.4.2 Short-run analysis

We begin by estimating the average treatment effects on the treated in the first quarter of the treatment period. Table 3.4 shows the estimation results. Column 1 and 2 show average electricity consumption in kWh per day in the baseline period and the first quarter respectively for all treatment groups and the control group. Columns 3 and 4 show *across* and *within* differences in electricity consumption, that is the difference between the consumption level in the first quarter ($Q1$) and the baseline consumption (Pre). *Across* differences are calculated as the simple difference of the two averages in columns 1 and 2, whereas *within* differences are calculated as the averaged differences between $Q1$ and baseline electricity consumption for each household. By considering *within* instead of *across* differences, the precision of the estimates is improved by a factor of around 3.5. Column 5 shows the difference-in-differences estimate based on differencing the *within* differences of treatment and control group. Column 6 presents the difference-in-differences estimates based on the estimation which additionally controls for weekly effects and column 7 controls for baseline electricity consumption in addition to the weekly fixed effects and represents the specification of equation 3.1.[12]

The smart metering group (G1) reduced electricity consumption on average by 0.101 kWh per day in the first quarter after the scheduled treatment start. This effect is at the margin of being statistically significant with a t-statistic of 1.58. The non-significance of the short-run effect for the smart metering group is however not surprising: As discussed in subsection 3.2.3 many in-home displays were installed with a delay of three months or more due to logistic problems. Many households thus have not been actually treated in the first two to three months after the scheduled treatment start, which makes a short-run analysis of the smart metering group based on the periodical electricity records a challenge. The short run effects of the smart meters can however be identified using the high-frequency smart meter data. Estimates based on smart meter data will be discussed in section 3.5.

Expert advice (G2) did not change electricity savings in the short run. The point estimate is even slightly positive but far from being statistically significant. Rather than interpreting the effect as average treatment effects on treated, the estimates could be taken as intent-to-treat-effects, because around two thirds of participants did not make use of the invitation. Figure 3.C1 shows the cumulative take-up rate of expert advice which tops out at around one third.

Sending out electricity statements with one's own and a peer household's electricity consumption triggered electricity savings on the order of 0.105 kWh per day for the social competi-

[12]Controlling for seasonal effects changes point estimates quite substantially for the smart metering group. Because in-home displays were mostly distributed with a delay of three months or more, the actual treatment did not start in the first month after the scheduled treatment start for most households in the smart metering group, which makes a seasonal correction important.

tion group (G3). A reduction of this magnitude corresponds to roughly 1.7 % of daily consumption or is equivalent to turning off a 60-watt light bulb for 1 hour and 45 minutes per day. Social comparison (G4) generated small but statistically non-significant savings. The social competition and social comparison treatment only differ from each other in how information is disclosed to the peer households: whereas in the social competition treatment households knew that their electricity consumption was disclosed to their peer households and vice versa, households in the social comparison treatment knew that their electricity consumption was *not* shown to their peer households. Knowing that one's own electricity consumption is shown to their peers could have provoked a feeling of competition among the two households or increased social pressure.

Table 3.4: Short-run difference-in-differences estimates

	Pre	Q1	Δ_{across}	Δ_{within}	$\Delta\Delta_{within}$	$\Delta\Delta_{within}$	$\Delta\Delta_{within}$
G1: Smart metering	5.99 (0.117) [N=1176]	6.227 (0.117) [N=3387]	0.237 (0.166)	0.26 (0.05)	-0.059 (0.068)	-0.101 (0.064)	-0.101 (0.064)
G2: Expert advice	6.22 (0.126) [N=1207]	6.572 (0.137) [N=3600]	0.352 (0.186)	0.346 (0.044)	0.027 (0.064)	0.017 (0.061)	0.019 (0.061)
G3: Social competition	5.879 (0.1) [N=1166]	6.108 (0.106) [N=3430]	0.229 (0.145)	0.222 (0.042)	-0.097 (0.062)	-0.104 (0.059)	-0.105 (0.059)
G4: Social comparison	5.953 (0.1) [N=1173]	6.261 (0.112) [N=3444]	0.308 (0.15)	0.29 (0.042)	-0.029 (0.062)	-0.026 (0.059)	-0.027 (0.059)
G0: Control group	6.107 (0.112) [N=1197]	6.426 (0.123) [N=3539]	0.319 (0.167)	0.319 (0.046)			
Weekly FE	-	-	-	-	No	Yes	Yes
Baseline consumption	-	-	-	-	No	No	Yes

Notes: This table shows average electricity consumption (kWh per day) for the pre-treatment period (column 1) and the first quarter after treatment start (column 2) for all treatment groups. Columns 3 and 4 show the *across* and *within* differences for each treatment group. Columns 5 to 7 displays diff-in-diff estimates with and without weekly fixed effects and baseline consumption as additional controls. Standard errors clustered by household in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on electricity readings.

3.4.3 Medium-run analysis

We now turn to a discussion of the medium-run effects of the treatments. Figures 3.4 to 3.6 show the evolution of the treatment effect over the whole treatment period. The point estimates at the different measurement periods are based on the econometric specification discussed in equation 3.1 and are shown together with the 90 and 95 % confidence intervals.

Figure 3.4 confirms, that electricity savings are not significant in the first quarter for the households in the smart metering treatment. However, as soon as almost all in-home displays were rolled out, we start to detect effects on the order of -0.25 kWh per day. A reduction of this order corresponds to a reduction of around 4 % of average daily consumption (2,351 kWh per year divided by 365.25). An effect of this magnitude compares to each household turning off four light bulbs for an hour or turning off television for 1 hour and 45 minutes per day. The effect size is very comparable to those found in existing studies. Based on an evaluation of existing pilot studies with comparable study populations, Baeriswyl et al. (2012) evaluate the savings potential of smart metering between 1.2 and 3.7 % depending on the type of feedback. Moreover, observing the treatment effects over a longer period highlights that the savings effects persist over time and no sign of fading out is observed.

Figure 3.4: Main effects for smart metering group

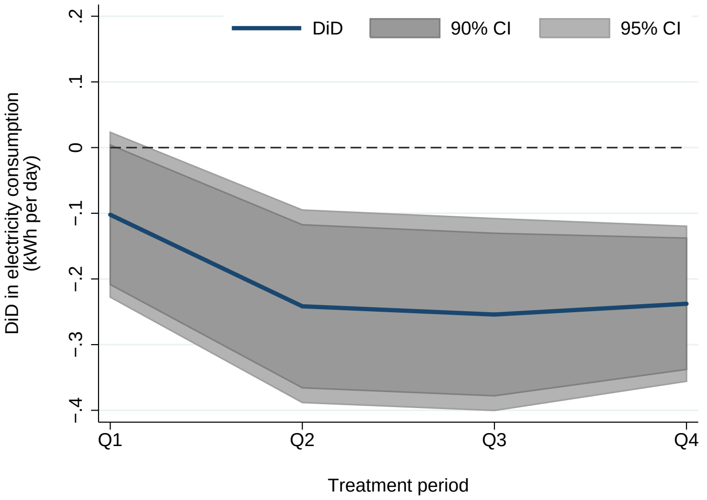

Notes: The figure show the average treatment effect on the treated together with the 90 % and 95 % confidence intervals.
Source: Own calculations based on electricity readings.

Figure 3.5 shows the medium-run effects of inviting participants for an expert advice session at *ewz*'s service center. Expert advice did not induce households to save electricity neither in the short-run nor in the medium-run. Treatment effects are close to zero or even positive and not statistically significant in all periods. Finding no effects for this treatment can have different reasons: First, take-up was relatively low. Only one third of participants chose to follow the invitation. Take-up could possibly have been increased if electricity advice would have been offered at the homes of participants instead asking participants to come to the service center. Second, expert advice might have increased knowledge without necessarily translating into electricity savings. Table 3.B8 in the appendix shows that expert advice improved households perception of how easy they could change energy-efficient behavior in the household. Households with the expert advice treatment - unlike the other treatment groups - perceive it as easier to exchange most of the traditional light bulbs for energy-saving light bulbs or to dry clothes on a clothes line instead of using a dryer. However, households failed to translate the improved knowledge into concrete actions. An effective intervention should help households to bridge the gap from knowledge to behavior, possibly by combining expert advice with other types of feedback.

Figure 3.5: Main effects for expert advice group

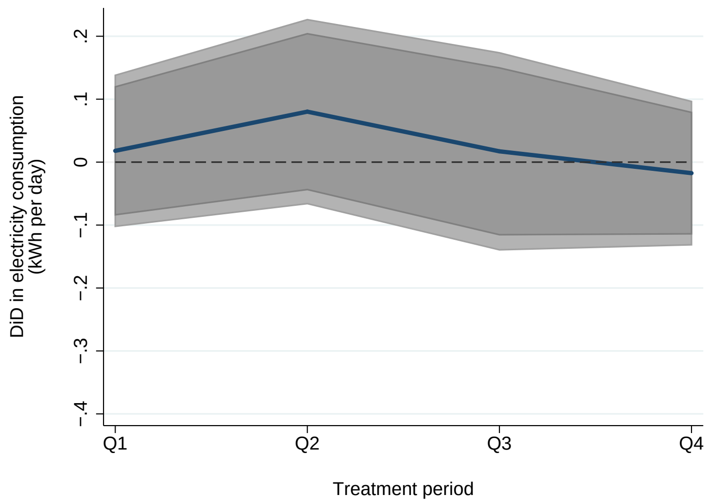

Notes: The figure show the average treatment effect on the treated together with the 90 % and 95 % confidence intervals.
Source: Own calculations based on electricity readings.

The medium-run effects of the social information treatments are assessed in figure 3.6. Subfigure 3.6a shows the average effects for social competition. The marginally significant short-run electricity savings of the social competition group level off in the medium-run with a point estimate virtually reaching zero after three quarters. In the fourth quarter, however, we observe again savings on the order of 1.5 % of daily consumption significant at the 10 % level. This backsliding of the savings effects until the third quarter and the reinforcement in the last quarter can have different explanations. *First*, it underlines the importance of feedback frequency. Electricity savings are significantly negative in the social competition group as long as feedback is given at a monthly frequency, but leveled out when the feedback rhythm changed to a quarterly frequency. *Second*, this cyclical pattern of action and backsliding might be explained by the competitive element of this intervention: households might have reinforced their savings effort towards the end of the treatment period in a last attempt to "win" the competition.

Subfigure 3.6b illustrate average treatment effects for the social comparison group. Point estimates are negative but small and never statistically significant over the whole treatment period. Finding no average savings effects for the social comparison group can have different reasons. On the one hand, the treatment was maybe not strong enough, that is receiving social norm information might just not have induced households enough to achieve significant electricity savings. On the other hand, averaging effects over the whole group might mask asymmetric effects within pairs. Clearly, the response to the treatments can be asymmetric depending on whether one's own consumption was above or below the peer's electricity consumption in the previous periods. Households who used cumulatively less electricity in the previous periods than their peers do not have strong incentives to reduce their electricity consumption even more.[13] Households, however, who used cumulatively more electricity in the previous periods might react more to the treatments due to social pressure or because they want to turn around the result and ”win” the competition.

Figure 3.6: Main effects of social information groups

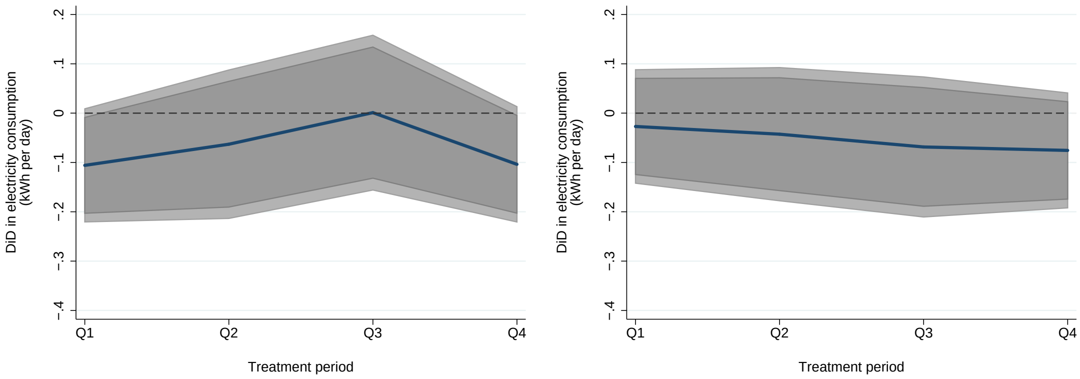

(a) Social competition (b) Social comparison

Notes: The figure show the average treatment effect on the treated together with the 90 % and 95 % confidence intervals.
Source: Own calculations based on electricity readings.

Figure 3.7 shows the medium-run treatment effects split by the cumulative previous electricity consumption. Cumulative previous electricity consumption is the sum of the average per day consumptions until the period prior to the one under study. The blue solid line shows treatment effects for households with higher cumulative electricity together with the 90 % and 95 % confidence intervals. The dashed yellow line depicts treatment effects for households who cumulatively used less electricity consumption than their peer household.[14] In the social competition group (subfigure 3.7a), both households with higher and households with lower cumulative electricity consumption react similarly to the social information letters. The reaction of the high users is however more pronounced than the low users' reaction in the first quarter and the fourth quarter. However, in the social comparison treatment (subfigure 3.7b), clear differential medium-run trends can be identified: while households with lower cumulative consumption save more electricity in the first quarter, households with higher cumulative consumption start to react to the treatment predominantly in the third and fourth quarter. Savings in the third and fourth quarter are significant and on the order of -0.25 kWh per day or 4 % of daily usage for households with a higher cumulative electricity consumption up until the period prior to the one under study.

[13]Apart from the color in which the cumulative percentage difference between one's own and the peer's electricity consumption was displayed– switching from green if one household was using less electricity than the other to red if one household was using more electricity than the other – the social information letters did not contain injunctive norms appealing to the pro-sociality of electricity conservation.

[14]Averaging over the higher and lower group does not exactly yield the average effects we showed in figure 3.6 because the latter is based on the totality of observations whereas here we discarded 3.5 % of observations where partner consumption was missing. Partner consumption might be missing if one of the two households in a pair dropped from the study.

These asymmetric effects are interesting in two aspects: First, we do find tendencies of a ”boomerang” effect especially in the social comparison group: the ”better” households stop saving electricity in the medium-run or even consume more. To some degree but less pronounced this behavior is also observed in the social competition group. Second, the differential effects across households with higher versus households with lower cumulative electricity consumption suggests that social pressure and self-reputational motives could have played a vital role and that the treatment increased the perceived pressure to reduce electricity among households who used more than their peers and coerced them into behavioral changes.

Figure 3.7: Effects by relative position within matched pair: higher versus lower

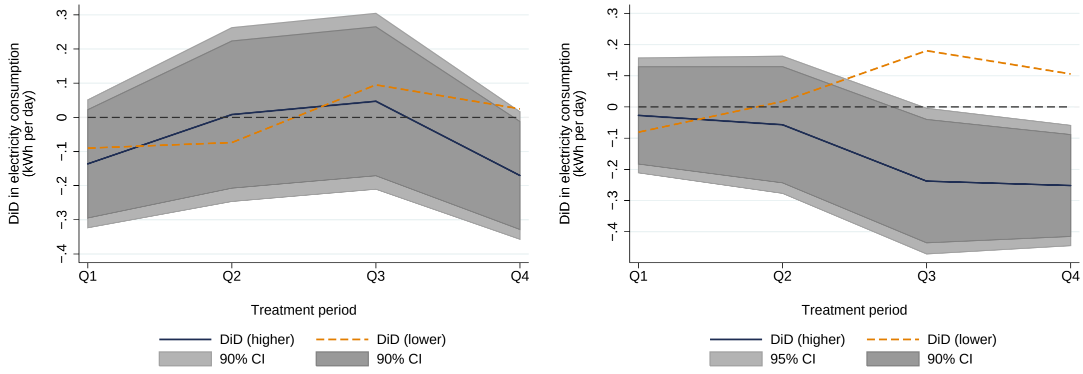

(a) Social competition

(b) Social comparison

Notes: The figure show the average treatment effect on the treated together with the 90 % and 95 % confidence intervals. The treatment effects are split by cumulative previous electricity consumption
Source: Own calculations based on electricity readings.

3.4.4 Heterogeneity of treatment effects

While the empirical focus has been on average treatment effects on the treated so far, the conceptual framework suggests that treatment effects could vary across households as a function of personal and household characteristics. Table 3.5 shows how treatment effects vary by a number of observed characteristics. In contrast to the analysis in subsection 3.4.3 the medium-run treatment effects are averaged over the second to fourth quarter of the treatment phase. The heterogeneity estimates are based on the following augmented specification:

$$\Delta e_{it} = \beta_0 + \gamma_0' \mathbb{T}_{\mathbb{i}} + \gamma_1' (\mathbb{T}_{\mathbb{i}} \times (x_i - \bar{x})) + \beta_1 (x_i - \bar{x}) + \beta_2' (e_{i0} \times \mathbb{Q}_{\mathbb{i}}) + \omega_{w(t)} + \varepsilon_{it} \ , \tag{3.2}$$

where $(x_i - \bar{x})$ represents a given observed trait which is centered to the mean. γ_0' now measures the average response of the treatment and γ_1' measures how the response of the treatment varies with trait x_i. Finally, β_1 measures how electricity consumption changes as a function of trait x_i

Column 1 in table 3.5 presents the main estimates. In accordance with the graphical representation of the medium-run estimates above, estimates are significant only for the social competition treatment in the first quarter. In the second to fourth quarter, the average effect of the in-home displays amounts to -0.244 kWh per day, whereas the effects of the other treatments are not significant.

A first source of potential heterogeneity is baseline electricity consumption. Households with high baseline electricity consumption could have larger savings potentials and reduce their consumption at lower cost. Also, social information could have differential effects for households in different parts of the pre-treatment electricity usage distribution. Social pressure could affect households in the upper part of the distribution more than households in the lower part of the distribution. Estimates show that households with larger baseline electricity consumption react more to the smart metering treatment only in the first quarter. For the other treatments no clear relationship is discernible. While higher baseline consumption tends to strengthen electricity conservation in the social competition group, effects are positive but negligible for the social comparison and the expert advice groups.

Treatment response could also vary by age. The theoretical predictions are however ambiguous: On the one side, one could expect older participants to react less to treatments, because they are less used to modern technologies and therefore might use the in-home displays less frequently or check their email less frequently and are more likely to miss the social information email. On the other side, older participants - in particular participants above retirement age - might have more free time to devote to the interventions which could reinforce electricity conservation. The empirical findings are mixed. While older participants tend to react more to the in-home display both in the short- and the medium-run, age reduces electricity conservation for the social competition treatment in the first quarter. Treatment effects do not differ with respect to gender, education and household income.

In summary, treatment effects are surprisingly homogeneous over the population. Especially in the medium-run, we do not find much evidence that treatment effects vary substantially with respect to different observables. If anything, higher baseline consumption and age tend to reinforce electricity conservation especially early in the treatment phase. Interestingly, the main effects become larger once we include interactions with age, female, education and household income. This points towards some selection effects with respect to households who responded to the questionnaires and those who did not. Households who responded to the surveys are potentially more interested in the study and thus might have reacted more to the interventions.

Table 3.5: Difference-in-differences estimates for electricity consumption (kWh per day)

	Main	Baseline cons.	Age	Female	Tertiary education	Household income
Quarter 1						
Smart metering (G1)	-0.101	-0.106*	-0.100	-0.105	-0.099	-0.073
	(0.064)	(0.064)	(0.069)	(0.070)	(0.070)	(0.071)
Expert advice (G2)	0.019	0.014	0.025	0.029	0.026	-0.017
	(0.061)	(0.059)	(0.064)	(0.065)	(0.064)	(0.066)
Social competition (G3)	-0.105*	-0.108*	-0.099	-0.099	-0.099	-0.094
	(0.059)	(0.060)	(0.063)	(0.064)	(0.063)	(0.067)
Social comparison (G4)	-0.027	-0.022	-0.036	-0.035	-0.038	-0.076
	(0.059)	(0.059)	(0.063)	(0.063)	(0.063)	(0.066)
G1 $\times$ Trait		-0.113*	-0.108**	-0.023	0.201	0.008
		(0.068)	(0.044)	(0.145)	(0.125)	(0.023)
G2 $\times$ Trait		0.008	0.035	-0.071	-0.007	0.018
		(0.039)	(0.042)	(0.110)	(0.111)	(0.022)
G3 $\times$ Trait		-0.050	0.071*	0.043	0.093	-0.009
		(0.048)	(0.042)	(0.113)	(0.109)	(0.023)
G4 $\times$ Trait		0.016	0.047	-0.051	0.044	-0.013
		(0.040)	(0.042)	(0.110)	(0.108)	(0.022)
Quarters 2 to 4						
Smart metering (G1)	-0.244***	-0.245***	-0.305***	-0.311***	-0.307***	-0.272***
	(0.059)	(0.058)	(0.061)	(0.062)	(0.062)	(0.061)
Expert advice (G2)	0.028	0.026	0.004	0.010	0.007	-0.037
	(0.060)	(0.057)	(0.062)	(0.063)	(0.062)	(0.063)
Social competition (G3)	-0.056	-0.054	-0.074	-0.076	-0.073	-0.046
	(0.061)	(0.064)	(0.064)	(0.064)	(0.064)	(0.069)
Social comparison (G4)	-0.062	-0.059	-0.112*	-0.111*	-0.112*	-0.139**
	(0.055)	(0.054)	(0.057)	(0.058)	(0.057)	(0.060)
G1 $\times$ Trait		-0.069	-0.091**	0.037	0.134	0.015
		(0.047)	(0.042)	(0.122)	(0.115)	(0.020)
G2 $\times$ Trait		0.000	0.006	-0.101	0.109	0.028
		(0.037)	(0.043)	(0.121)	(0.114)	(0.021)
G3 $\times$ Trait		-0.012	-0.001	-0.080	0.080	-0.001
		(0.046)	(0.047)	(0.120)	(0.118)	(0.024)
G4 $\times$ Trait		0.018	-0.019	-0.113	0.046	0.016
		(0.028)	(0.038)	(0.109)	(0.105)	(0.019)
Observations	33,398	33,398	28,308	28,072	28,290	25,378
Clusters	5,876	5,876	4,962	4,920	4,959	4,446
R^2	0.105	0.112	0.114	0.112	0.113	0.117

Notes: Standard errors clustered by household in parentheses. *** P<0.01 ** P<0.05 * P<0.1. This table shows average difference-in-differences estimates in the first quarter and the second to fourth quarter respectively. Column 1 presents main estimates and columns 2 to 6 shows estimates which are allowed to vary with respect to a number of observed characteristics. All estimates include weekly FE and baseline consumption interacted with quarter dummies. Standard errors clustered by household in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on electricity readings.

3.5 Analysis of the smart meter data

This section discusses the effects of real-time feedback using only the high frequency smart metering data to identify treatment effects. Subsection 3.5.1 discusses the empirical specification, subsection 3.5.2 presents a descriptive analysis of the usage of the in-home displays and subsection 3.5.3 discusses the aggregate effects of the smart metering treatment. Subsection 3.5.4 the overall effects are broken down to the hourly level.

3.5.1 Empirical strategy

The fact that the smart meters started to record electricity consumption at a high frequency around 60 days before treatment start, i.e. before the in-home displays were installed, helps us to identify the effects of the smart meters using only the smart metering data. A pure *before-after* comparison of households would mix treatment effects with time effects. It is thus critical that we can control for time effects using households as a control group whose in-home displays were not installed yet, but whose smart meter was already recording electricity consumption. This is possible because households were randomly attributed to the cohorts so that the timing of the installation is random. Clearly we can identify treatment effects only in periods where there is a fraction of households whose displays are not installed yet. As soon as all displays are installed, it is no longer possible to disentangle treatment effects from pure time effects. Figure 3.8 shows the share of households without installed display over time. Because the households of the pilot cohort are included in the sample, the share proportion of households without in-home display initially drops from 1 to almost zero and raises steadily due to the weekly roll-out of the 26 remaining cohorts until mid July 2011. The installation of the in-home displays was rolled-out mainly over the period from July 2011 to February 2012, so that the proportion of households without in-home display steadily decreases after July 2011. The share of households without an installed display is above around 20 % in the period from June 2011 to January 2012. During this period, we can identify the treatment effect of providing continuous and real-time feedback via smart meters using only the high-frequency records of the smart meters. For analyses beyond this time period, treatment effects have to be identified using the periodical meter readings discussed above.

We estimate the causal effects of the in-home displays on electricity consumption per day and on average hourly consumption over the daily profile. For the analysis of the effects on daily electricity consumption, we employ the following empirical specification:

$$e_{it} = \alpha_i + \omega_{w(t)} + \beta D_{it} + \varepsilon_{it} \ , \tag{3.3}$$

where e_{it} is electricity consumption in kWh per day of household i at time t. Household-specific differences in electricity consumption are absorbed in a household-specific constant α_i. Furthermore, we control for time effects using a weekly fixed effect $\omega_{w(t)}$. D_{it} a dummy which equals 1 if the in-home display of household i is installed at time t. The coefficient β thus reflects the causal effect of installing the in-home display on daily electricity consumption.

Figure 3.8: Proportion of households without in-home display

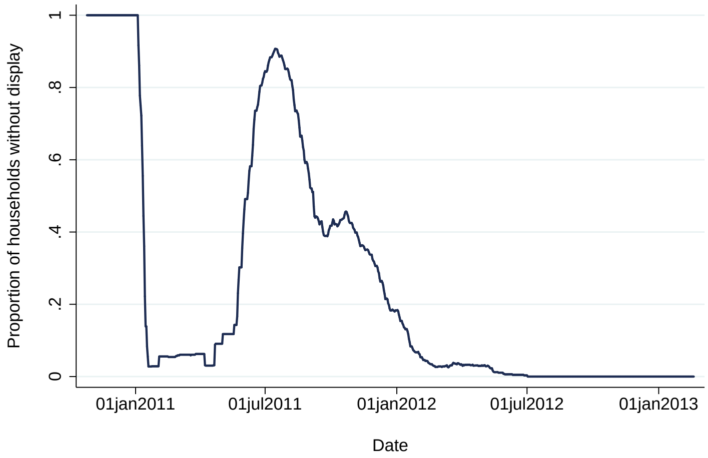

Notes: This figure shows the proportion of housholds without installed in-home display across time.
Source: Own calculations based on smart meter data.

The high frequency records of the smart metering data allow us to investigate how the in-home displays affected electricity consumption over the daily profile. To estimate the hourly effects of the in-home display over the day, we estimate the following estimation equation:

$$e_{ith} = \alpha_i + \delta_h + \omega_{w(t)} + \sum_{h=1}^{24} \beta_h D_{ith} + \varepsilon_{ith} \ . \tag{3.4}$$

The notation is the same as in equation 3.3 with the only difference that the unit of measurement is hourly instead of daily. δ_h are hourly fixed effects. β_1 to β_{24} are the causal effects of the in-home display on electricity consumption (in kWh per hour) in a given hour of the day. We account for potential heteroskedasticity and clustering on the household level by using cluster-robust standard errors.

3.5.2 Usage of in-home displays

Households use the in-home display frequently, but not all features were used at the same frequency. Figure 3.C2 shows the different features of the device. The main screen shows real-time consumption and achievement of the weekly goals. The second screen allows to compare electricity consumption across different measurement periods. On the third screen households could measure current electricity consumption in the household over a specific period. In the last screen households could fix the weekly goals. Figure 3.9 shows how usage of the different features evolved over time. The figure shows the proportion of households who have used the display at least once in a given week after installation. Clearly, this proportion is one in the beginning of the treatment period, because the display has been turned on at least once upon installation. Usage of the device diminishes relatively steeply to around 60 % after 4 weeks. Afterwards, usage decreases only slowly. Three months after installation of the device, 50 % of the households still use the device at least once in a given week, and after one year, that proportion is still around 30 %.[15] The real-time screen was used most often. This is not surprising, because it was the starting screen after turning on the device. The frequency of usage diminishes with the order of the features: the second screen was touched second most often, and the last screen was used least often.

Figure 3.9: Proportion of households using the display

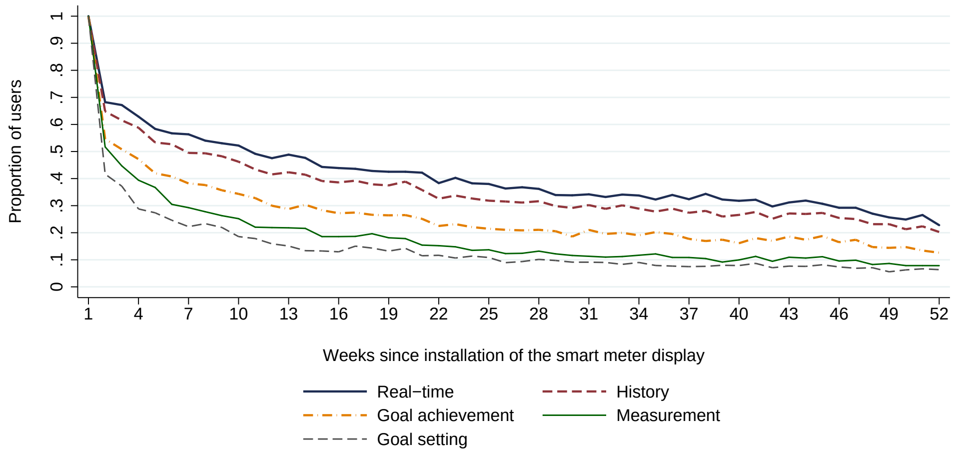

Notes: This figure shows the average proportion using the different functionalities of the display at least once per week across time since installation of the display.
Source: Own calculations based on smart meter data.

Despite the diminishing frequency of usage, time spent with the display is relatively stable conditional on using it. Figure 3.10 shows average duration of usage of the displays in minutes per day for frequent and infrequent users. Frequency of use is defined as the total number of uses over the whole treatment period. Frequent users are households in the upper half of the frequency distribution and infrequent users are households in the lower half of the distribution. Frequent users do not only use the device more frequently than infrequent users, but they also spend more time on the display. Conditional on using it, frequent users spend around 9 minutes per day with the display and infrequent users around 6 minutes.

[15]Part of this surprisingly high frequency of usage might be explained by the fact that households were asked to switch on the device once every three weeks in order to ensure proper storage of the recorded data.

Figure 3.10: Duration of usage per day

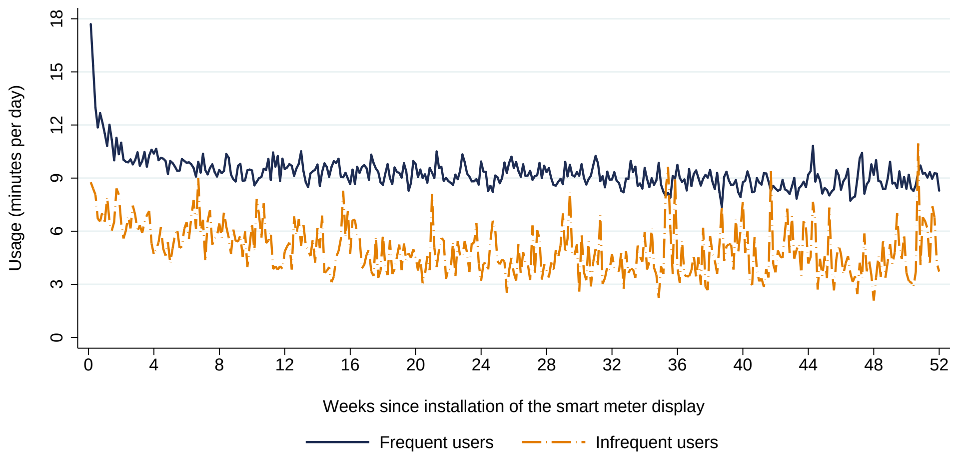

Notes: This figure shows the average duration of usage of the in-home displays in minutes per day.
Source: Own calculations based on smart meter data.

3.5.3 Aggregate effects

Figure 3.11 provides graphical evidence how the displays affected normalized electricity consumption over time. In order to control for household specific heterogeneity, we first normalize electricity consumption with respect to the mean electricity consumption of each household. To eliminate time effects, we then normalize electricity consumption with respect to the mean electricity consumption of the control group in a specific time period. That is, we only take into account observations *before* installation of the in-home displays in the second normalization step. The figure shows the deviation from mean daily consumption (kWh per day) in the 7 weeks before and in the 26 weeks after installation of the in-home displays. A value of zero thus means that on average a household with display uses as much electricity as a household without display and negative values imply that households with the display consume less electricity than their counterparts without display in a given week.

Before installation of the in-home displays the deviation from mean daily consumption fluctuates around zero. Within a few weeks after installation of the device we observe a pronounced drop in electricity consumption relative to households without a display. Four weeks after installation of the display, households reduced electricity consumption by around 0.4 kWh. Over the remaining period, electricity reductions fluctuate between 0.2 and 0.4 kWh per day or 3 to 6 % of daily consumption on average. The installation of the in-home displays thus triggered fast and persistent reductions in electricity consumption.

Table 3.6 presents the supporting regression estimates for the average treatment effect on electricity consumption (columns 1 and 2) and expenditures (columns 3 and 4). Treatment effects in columns 1 and 3 are based on hourly effects which are aggregated up to the daily level and include hourly and weekly fixed effects. Columns 2 and 4 are based on daily data and

Figure 3.11: Deviation from daily mean

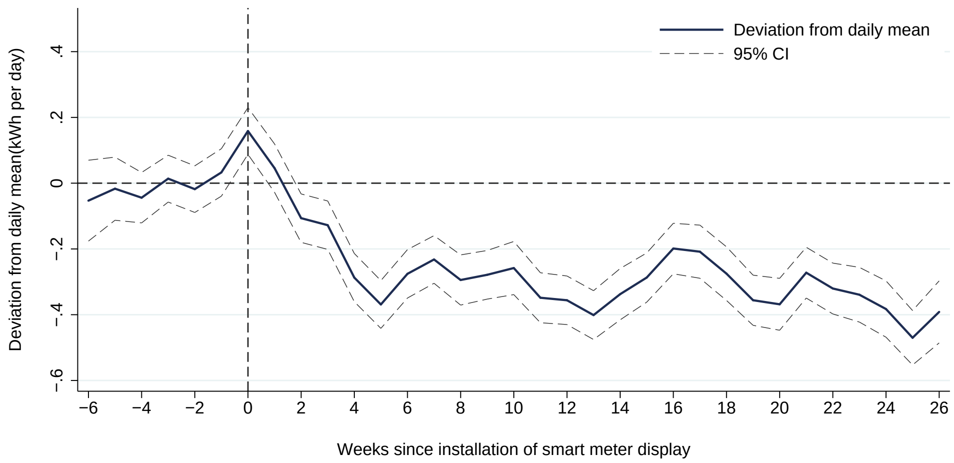

Notes: This figure shows the deviation from mean daily consumption (kWh per day) in the 7 weeks before and 26 weeks after installation of the in-home display. Mean daily consumption is based on observation *before* installation only. Week 0 denotes the week prior to installation.
Source: Own calculations based on smart meter data.

include weekly fixed effects only. All specification account for individual heterogeneity through individual fixed effects.

Panel A presents estimates for the overall sample. Installing the in-home display reduces electricity consumption on average by 0.23 kWh per day or 84 kWh per year. Estimates are robust to the estimation method. Irrespective of whether we estimate hourly effects and aggregate them up to the daily level (column 1) or whether we estimate daily effects (column 2), estimates are on the same order of magnitude. The electricity savings translate into reductions in electricity expenditures on the order of 4 cents per day or 14 Swiss francs per year. Despite different identification strategies, estimates are very similar to the estimates which are based on the periodical electricity readings.

Panel B and C display estimates for frequent and infrequent users of the display respectively. Frequency of usage is defined as above, based on the total number of usages over the whole treatment period. The table shows that frequent users thus not only use the display more frequently and longer, but also conserve more electricity and save more in terms of electricity expenditures. Frequent users conserve around 0.28 kWh per day of electricity, whereas infrequent users save on average around 0.19 kWh - a difference which is however not statistically significant. Nevertheless, observing larger effects for frequent users suggests that behavioral changes are not triggered by the mere existence of the display in the household, but rather by an active usage of the device. Understanding how the intervention affected behavior would however need a more in-depth analysis of the underlying psychological mechanisms and is beyond the scope of this paper.

Table 3.6: Average treatment effects

	Electricity consumption kWh per day		Electricity expenditures CHF per day	
	(1)	(2)	(3)	(4)
	A. Total			
Treatment effect	-0.234***	-0.230***	-0.039***	-0.039***
	(0.060)	(0.060)	(0.011)	(0.011)
N	7,435,617	310,893	7,435,617	310,893
Clusters	1,003	1,003	1,003	1,003
	B. Frequent users			
Treatment effect	-0.284***	-0.267***	-0.052***	-0.049***
	(0.087)	(0.087)	(0.016)	(0.016)
N	4,531,535	189,272	4,531,535	189,272
Clusters	506	506	506	506
	C. Infrequent users			
Treatment effect	-0.194**	-0.200**	-0.027*	-0.029**
	(0.082)	(0.082)	(0.014)	(0.014)
N	2,904,082	121,621	2,904,082	121,621
Clusters	497	497	497	497
Hourly FE	Yes	No	Yes	No
Weekly FE	Yes	Yes	Yes	Yes

Notes: This table shows average treatment effects on treated for electricity consumption (columns 1 and 2) and for electricity expenditures (columns 3 and 4). Panel A shows overall effects, panel B displays effects for frequent users of the smart meter display, and panel C for infrequent users respectively. Columns 1 and 3 are based on an aggregation of hourly effects, and columns 2 and 4 are based on daily data. Standard errors clustered by household in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on smart meter data.

3.5.4 Disaggregation over the daily profile

The aggregate analysis showed that displays showed first effects relatively fast after installation which persisted over time. An interesting question is whether these reductions were achieved through adjustments at the extensive or at the intensive margin. Adjustments at the extensive margin are reductions that are obtained through a readjustment of one's privately optimal electricity consumption level, for example through turning off the television, switching off light bulbs or using the coffee machine less often. Adjustments at the intensive margin are attained through a more efficient use of electricity without reducing the level of consumed electricity services, for example through exchanging traditional light bulbs with energy-efficient ones or consequently turn off standby switches of appliances which are not currently used. Adjustments at the intensive margin should be visible as general reductions over the whole day load profile, whereas adjustments at the extensive margin are expected to be more concentrated in times that households spend at home such as evening hours for example.

Figure 3.12 shows the average day load profile of smart metering households before treatment start. Electricity consumption is relatively low during night hours, grows gradually over the morning hours until it first peaks around midday, and stays relatively flat in the afternoon. The bulk of the daily electricity is consumed in the evening hours between 4 and 11pm.

Figure 3.12: Day load profile before treatment

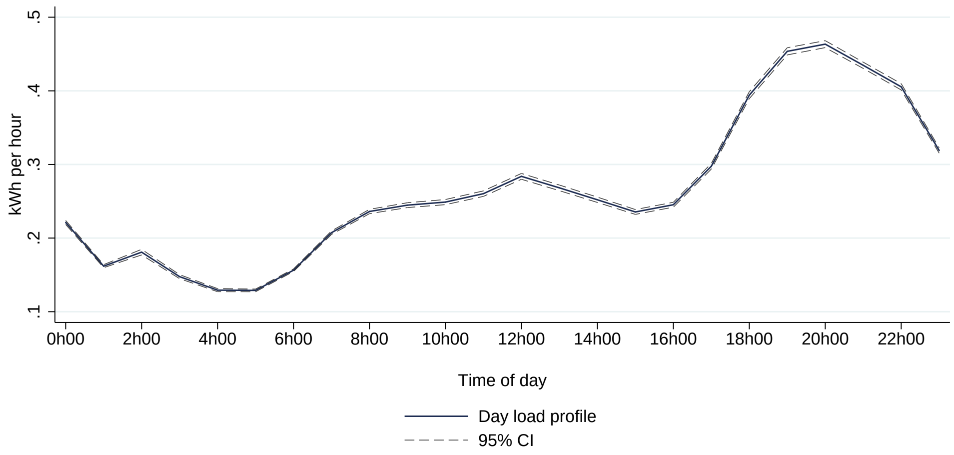

Notes: This figure shows the average day load profile of smart metering households before treatment start.

In figure 3.13 we show deviations from the hourly average of control households in the 50 days before and after installation of the display. Clearly, deviations from hourly means fluctuate around zero before installation of the display, because this is how we defined the normalization. After installation of the smart meters the daily profile of the reductions changes significantly. *First*, we observe a general reduction of electricity consumption over all times of the day starting around 6am. *Second*, the bulk of the reduction is achieved in the hours between 5 and 10pm. In peak hours between 8 and 9pm, the reduction amounts to around 0.035 kWh per hour or almost 8 % of electricity use in that period. *Finally*, we also observe a substitution of electricity consumption away from high-tariff to low-tariff hours. Electricity consumption slightly increases during the night hours between 1am and 5am.

The graphical evidence is confirmed in table 3.B7. The table contains causal estimates for the reduction in electricity consumption (columns 1 to 3) and electricity expenditures (columns 4 to 6) after installation of the displays and is estimated for both the overall sample and for frequent and infrequent users separately. In contrast to the graphical evidence of figure 3.13, which covers the short-run effects in the 50 days after installation, the estimates in table 3.B7 represent average treatment effects over the whole treatment period. The overall picture remains the same: Households reduce their electricity consumption significantly over all times of the day. The largest reductions are observed around 6 and 7pm, and from 11pm to 4am we observe a substitution from high-tariff to low-tariff consumption. Frequent users conserve more electricity in the peak hours, but they also substitute more towards low-tariff hours than

Figure 3.13: Deviation from hourly mean

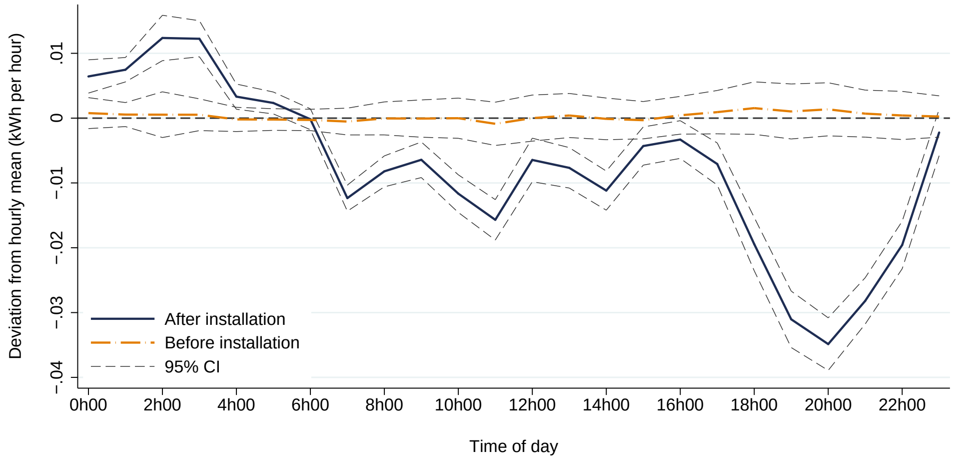

Notes: This figure shows the deviation from mean hourly consumption (kWh per hour) over the day load profile in the 50 days before and after installation of the in-home display.
Source: Own calculations based on smart meter data.

infrequent users.

Figure 3.14 graphically confirms the tendency of household to increasingly substitute consumption from high- to low-tariff hours. For the sake of visibility, confidence intervals are only displayed for the pre-installation effects. The figure splits up the average treatment effects into three periods. Households start to reduce electricity consumption already in the first three weeks after installation of the display. In the period from 4 to 6 weeks after installation, the treatment becomes fully effective and large reductions are observed especially during the evening peak hours. Also households start to substitute consumption towards low-tariff hours. In the period from 7 to 26 weeks after installation, household continue to conserve electricity over the whole time of the day, but at the same time substitute some of their consumption into night hours.

Figure 3.14: Deviation from hourly mean - over time

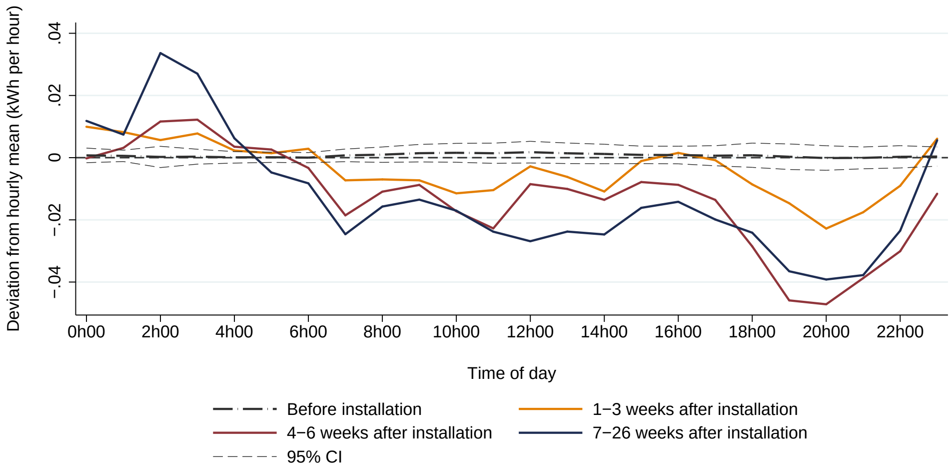

Notes: This figure shows the deviation from mean hourly consumption (kWh per hour) over the day load profile in the 50 days before and after installation of the in-home display in different time periods after installation of the display.
Source: Own calculations based on smart meter data.

In summary, the analysis of the treatment effects over the daily profile highlights some interesting observations: *First*, households react very fast to the treatment. Reductions are already discernible in the first three weeks after treatment start. *Second*, while there is some reduction over the daytime, the largest savings are achieved during the peak hours in the evening. This points towards adjustments at the extensive margin. Households predominantly tend to reduce electricity consumption in the evening hours and reductions over the remaining daytime play only a second order role. *Finally*, the observed substitution electricity consumption from high-tariff to low-tariff hours also supports the view that reductions were predominantly achieved through selected behavioral changes rather than an exchange of traditional for more energy-efficient appliances.

3.6 Welfare implications

Cost effectiveness. Especially when comparing different types of feedback, cost effectiveness in terms of cents of costs per kilowatt-hour of electricity conserved is a measure of great practical interest. While the smart metering treatment is the most effective treatment in terms of kilowatt-hours conserved, it may not be the most cost-effective one, because it requires installation of costly smart meter displays. In this section, we will give a stylized analysis of in-sample cost effectiveness. In-sample means that we only account for the total electricity savings between the beginning and the end of the treatment period and do not extrapolate further in the future. The measure of cost effectiveness we use is the annualized costs of the intervention divided by kilowatt-hours saved per year:

$$\text{Cost effectiveness} = \frac{\text{Cost per treatment per year}}{\hat{\gamma}_j \times 365} \,,$$

where $\hat{\gamma}_j$ denotes the average electricity conservation in kWh per day of treatment j over the whole treatment period.

On average, the smart metering treatment reduced electricity consumption by 0.23 kWh per day or 84 kWh per year (see table 3.6). The main costs associated with the smart metering treatment are the investment and installation costs of the smart meters and the in-home displays. Baeriswyl et al. (2012) estimate that investment costs of smart meters and in-home displays are around 215 Swiss francs per household. The installation of the devices costs another around 200 Swiss francs. We thus assume costs on the order of 415 Swiss francs per household for the smart metering treatment.[16] Assuming a durability of 10 years for smart meters and in-home smart meter displays and no time discounting, the annualized cost of the smart metering treatment is 41.5 Swiss francs per household. Based on these numbers, the cost effectiveness of the smart metering treatment is around 50 cents per kilowatt-hour saved.

Neglecting the additional indirect costs of additional resources needed in the service center of *ewz* due to the treatment, expert advice was virtually costless. At the same time it did not not induce any electricity conservation, which makes a cost effectiveness analysis of that treatment pointless.

Social information letters generated on average electricity conservations of 0.082 kWh per day or 30 kWh per year for the social competition and insignificant 0.044 kWh per day or 16 kWh per year for the social comparison group. Allcott (2011) or Allcott and Rogers (forthcoming) calculate the costs of the energy reports with one dollar per letter. Because in our field experiment the social information letters were sent by email, distribution of the information is virtually costless. However, the pairwise matching of the households and the exporting of the information into the electricity statements incur additional costs. Neglecting fixed indirect costs of adapting the data management system or implementation of a pairing algorithm, we

[16] The effective costs of the smart metering treatment were around 2250 Swiss francs per household. Basing cost effectiveness on this number would highly overstate the cost side because the effective costs reflect in large parts the R&D costs of developing an in-home display prototype.

use a cost of one Swiss franc per social information e-mail. Additionally, the meters have to be read-off more frequently. According to Baeriswyl et al. (2012) meter readings in an urban area incurs costs of around four Swiss francs per reading. During the one-year treatment phase, meters were read-off eight times. Neglecting fixed costs of program implementation, the total costs of the social information treatment add up to around 39 Swiss francs per household per year. Cost effectiveness is thus 1.3 Swiss franc per kilowatt-hour saved for social competition and 2.45 Swiss francs per kilowatt-hour saved for the social comparison group.

Using annualized costs and disregarding time discounting, the smart metering treatment is not only the most effective treatment with respect to electricity savings, but is also most cost effective. Empirical evidence further suggests that savings were persistent only for the smart metering group. Extrapolating into the future and making assumptions about persistence would thus make even more of a difference. Given that the distribution of social information through email is virtually costless, interventions which combine smart meters (instead of traditional meters which have to be read off manually) with social information could be a potentially cost-effective way of reducing household electricity consumption.

Social welfare. Impact evaluations of interventions are highly incomplete if they only focus on administrative cost effectiveness without accounting for social welfare effects. The different mechanisms through which the interventions could act can have different implications for consumer welfare. On the one hand, interventions can increase individual's utility if they positively resonate with their intrinsic motivation to save electricity. Interventions may facilitate a utility-improving readjustment of one's own privately optimal electricity consumption level if they mainly act through an information improvement. On the other hand, welfare gains may be completely undone if the behavioral changes have been coerced by social pressure.

Quantifying the social welfare benefits and costs of the interventions is difficult in practice, because the costs of changing usage behavior or changing the capital stock are usually unobserved. Without appropriate measures for the welfare implications of the interventions, such as lost leisure or cost constraints due to treatments, an analysis of social welfare remains relatively limited and imprecise. Nevertheless – as a first step towards analyzing social welfare – we can make use of the rich information from the surveys to find proxies for the overall utility gains or losses due to the interventions. We use three different dimensions of satisfaction - satisfaction with life, satisfaction with the environment and satisfaction with the quality of life - as proxies for utility. These measures are assumed to incorporate monetary and non-monetary gains and costs of the interventions.

Table 3.7 reports the average effects of the treatments on the three satisfaction measures. None of the treatments has a significant impact on satisfaction with live, environment and quality of life. Point estimates are small and vary in their sign. The interventions therefore neither deteriorated nor improved participants' utility. A zero overall effect could imply that the social benefits of the information improvement are offset by the negative social pressure effects or the costs of changing the capital stock and/or the behavioral habits.

Table 3.7: Difference-in-differences estimates for satisfaction

	Life	Environment	Quality of life
Smart metering	0.019	0.001	0.022
	(0.023)	(0.031)	(0.026)
Expert advice	-0.017	-0.019	0.011
	(0.024)	(0.032)	(0.027)
Social competition	0.040	-0.009	0.021
	(0.026)	(0.032)	(0.028)
Social comparison	-0.005	0.027	0.018
	(0.025)	(0.033)	(0.028)
Observations	14,839	14,775	14,726
Clusters	4,579	4,566	4,557
R^2	0.007	0.005	0.006

Notes: Standard errors clustered by household in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on data from ewz.

3.7 Conclusion

We analyze the electricity savings potentials of information in a randomized controlled field experiment. The three types of information are (i) real-time and detailed feedback using smart metering technology, (ii) personalized electricity savings tips through energy experts and (iii) social information letters with monthly or quarterly electricity statements of one's own and a peer household's electricity consumption.

Real-time feedback reduced electricity consumption on average by 3 to 5 % of daily consumption. Households started to reduce electricity consumption shortly after receiving the in-home display and the savings persisted over the whole treatment period. Breaking the effects down to the hourly level we find that the largest savings were realized during peak hours. At the same time, household increasingly substituted electricity consumption from high-tariff to low-tariff hours. Observing targeted reductions predominantly in evening hours suggests that households react to the treatment with behavioral changes rather than with investing in energy-efficient appliances, which would reduce electricity consumption uniformly over the daily profile.

Personalized savings tips through energy experts improved the perception of how easy a change in energy-efficient behavior would be, but did not translate into concrete actions. Moreover, because the treatment was a relatively weak one, the take-up rate was relatively low. In summary, expert advice can be helpful to improve information, but fails to bridge the gap from knowledge to behavioral changes. In combination with treatments that induce behavioral changes, expert advice could be a helpful complement to boost electricity conservation.

Social competition reduces electricity consumption by around 1.5 % of daily usage. Treatment effects are significant only in the first quarter and the fourth quarter and disappear in the between-period. This pattern underlines the importance of the feedback frequency. Electricity savings leveled out as soon as feedback frequency changed from once-per-month to once-per-quarter. Moreover, this pattern of action and backsliding suggests that the competitive element of this treatment could have played an important role: households might have reinforced savings efforts towards the end of the treatment period in a last attempt to "beat" their peer household.

Surprisingly - given the conceptual similarity of the two social treatment groups - social comparison did not significantly reduce electricity consumption on average. While point estimates are negative, they are too small to statistically distinguish them from zero. However we find that the responses to the social comparison treatment are highly asymmetric. While households with electricity consumption below their peers' consumption tend to save in the beginning of the treatment period but stop saving electricity in the later periods of the treatment, households with higher electricity consumption save up to 0.25 kWh per day or around 4 % of daily usage in the third and fourth quarter of the treatment period. The findings are interesting in two respects: First, we do find tendencies of a "boomerang" effect: the "better" households stop saving electricity in the medium-run or even consume more. This is a deficiency which could be corrected easily in a real application. Social information could be sent out only to higher users or the social information could be supplemented with injunctive norms that appeal to the pro-sociality of electricity consumption. Second, the asymmetric response also shows that being above the other household might have increased the perceived pressure to reduce electricity and coerced households into behavioral changes.

Analyzing the effects of the interventions on utility, we find neither positive nor negative effects on various measures of satisfaction. This can imply that the social benefits of improved information and knowledge are offset by the negative effects of social pressure or the monetary and non-monetary costs of behavioral changes. The identification and separation of the underlying mechanisms is thus crucial for an effective policy intervention, but studies which can separate the underlying mechanisms are still scarce. Separation of positive welfare effects of improved information from the negative social pressure effects is addressed in a follow-up experiment on energy conservation during showering. Tiefenbeck et al. (2014) exploit the abundance of personality measures from the surveys to disentangle the positive utility effects of real-time feedback from the negative social pressure effects and find no evidence for social pressure as underlying mechanism of the energy reductions .

3.A Choice experiments

In addition to the personality measures collected in the surveys, we integrated three choice experiments into the surveys to elicit social preferences, overconfidence and time preferences.

The first choice experiment is a public goods game. The set-up of the experiment closely follows Fischbacher et al. (2001). In each cohort, participants were grouped into groups of ten. Each subject received a budget of 100 units and could either keep these 100 units for herself or invest a fraction g_i into a group project. Each unit invested into the group project was doubled. The total amount from the group project was equally shared among the ten group members irrespective of their contribution. The individual payoff function was the following:

$$\pi_i = 100 - g_i + 0.2 \sum_{i=1}^{10} g_i \ .$$

First, a number of examples were shown to the subjects in order to ensure that subjects understood the mechanism of the game. Then, the subjects were asked to make two different types of contribution decisions. In the first decision situation, subjects had to decide how much to contribute to the group project without knowing how much the other nine participants contributed. In the second decision situation, subjects could make their contribution dependent on how much the others contributed to the group project. Participants had to decide how much they were willing to contribute, knowing that the other nine participants contributed on average between 0 and 20 units, between 21 to 40 units, and so on. In each cohort, one group of ten participants was randomly chosen and paid off according to their choices. One unit was directly translated into 1 Swiss franc. One of the 10 participants was randomly chosen to be paid off according to the second decision situation and the nine other participants were paid off according to their choice in the first decision situation.

Clearly this type of decision situation creates a conflict between the individual interest and the group interest. The payoff could be maximized if all subjects would invest the full 100 units into the group project. In this case, all subjects would receive a payoff of 200 units. The standard prediction, however, is complete free riding by all subjects: Investing one unit into the group project costs one unit and generates two units for the group. The marginal payoff of a contribution to the group project is however only 0.2 units. Thus, the size of the contribution to the group project is a measure for a subjects social preferences. Distinguishing the two decision situations allows us to understand different motives of cooperation. The first decision situation measures unconditional cooperation and serves as a measure of altruism, whereas the second decision situation captures conditional cooperation and allows us to get an estimate for reciprocity.

The second choice experiment measures the tendency of participants to over- or underestimate themselves. We asked the participants the following questions:[17]

[17] See Merkle and Weber (2011) for a similar approach.

If you would be compared a 100 times with two randomly chosen participants of this study:

- *How often would your electricity consumption per month be larger than the electricity consumption of the two other households? In . . . out of 100 times.*
- *How often would your electricity consumption lie between the electricity consumption of the two other households? In . . . out of 100 times.*
- *How often would your electricity consumption per month be smaller than the electricity consumption of the two other households? In . . . out of 100 times.*

The questions were asked not only for electricity consumption, but also for income, knowledge about electricity conservation and water usage. If subjects have realistic beliefs about their performance, the average probability in each of the categories should be one third. Typically, this is not what is observed, because individuals have a tendency to overestimate positive attributes.

In the third choice experiment we measured time preferences. Epper et al. (2011) highlight that the energy-efficiency gap in the market for energy-using durables could be attributed to consumers undervaluing future cost savings due to their pure time preferences. Commonly the literature focuses on two aspects of time preferences: First, how much weight do individuals put on utility in the far future compared to the utility in the near future. Second, how important is the bias towards the presence, i.e. how large is an individuals tendency to prefer immediate utility over future utility. We measure both aspects by considering four different decision situations. In each decision situation, subjects have to decide between a payoff A of 100 units and a payoff of 110 units (or 120 units, depending on the decision situation) later than payoff A. Participants had to answer the following question: How many weeks are you willing to wait for the larger payoff B? Subjects could choose a number between 1 and 46 weeks. For each cohort, 5 subjects were chosen and paid off according to the following incentive-compatible scheme: The lottery numbers of the week after the end of the survey determines the waiting time until payoff. If a participant indicated a waiting period which was larger than the lottery number x, then payoff B was paid off exactly x weeks later. If the indicated waiting period was smaller or equal to the drawn lottery number x, then subjects receive payoff A. We distinguished four different decision situations: In case 1 (2), participants had to decide between 100 units now versus 110 (120) units x weeks later. In case 3 (4), participants had to decide between 100 units in 4 weeks versus 110 (120) units in 4 + x weeks. Cases 3 and 4 measure long-term patience and cases 1 and 2 measure the existence of a bias towards present payoffs. We account for possible order effects by randomizing the ordering of the four different cases.

3.B Additional Tables

Table 3.B1: Average daily consumption across treatment periods and groups

Treatment period	G0	G1	G2	G3	G4	Total
Pre	6.107 (0.112)	5.990 (0.117)	6.220 (0.126)	5.879 (0.100)	5.953 (0.100)	6.031 (0.050)
M1	6.276 (0.129)	6.166 (0.121)	6.459 (0.140)	5.998 (0.108)	6.155 (0.112)	6.213 (0.055)
M2	6.400 (0.127)	6.181 (0.119)	6.538 (0.137)	6.077 (0.111)	6.223 (0.116)	6.287 (0.055)
M3	6.604 (0.126)	6.333 (0.129)	6.721 (0.144)	6.251 (0.112)	6.410 (0.119)	6.467 (0.057)
Q2	6.743 (0.127)	6.338 (0.123)	6.954 (0.145)	6.419 (0.119)	6.574 (0.118)	6.611 (0.057)
Q3	6.421 (0.124)	6.047 (0.119)	6.560 (0.139)	6.127 (0.119)	6.228 (0.115)	6.282 (0.055)
Q4	6.161 (0.121)	5.756 (0.114)	6.232 (0.130)	5.728 (0.103)	5.925 (0.109)	5.966 (0.052)
Avg. annual consumption	2,351	2,225	2,403	2,226	2,282	2,299

Notes: This table shows average electricity consumption (kWh per day) of participants across treatment periods and groups. The last row shows average annual consumption for each treatment group.
Source: Own calculations based on electricity readings.

Table 3.B2: Readings across group and periods

Period	Control group G0	Smart meter G1	Expert advice G2	Social competition G3	Social comparison G4
Pre	1,197	1,176	1,207	1,166	1,173
M1	1,192	1,128	1,206	1,155	1,163
M2	1,180	1,120	1,202	1,143	1,150
M3	1,167	1,139	1,192	1,132	1,131
Q2	1,129	1,097	1,158	1,086	1,073
Q3	1,099	1,052	1,125	1,049	1,041
Q4	1,064	1,009	1,082	1,017	1,013

Notes: This table shows how the number of observations varies across groups and periods.
Source: Own calculations based on electricity readings.

Table 3.B3: Attrition: Baseline electricity consumption (kWh) per day

	G0 (1)	G1 (2)	G2 (3)	G3 (4)	G4 (5)	Total (6)
Drop-outs	5.474	5.828	5.450	6.090	5.489	5.678
	(0.277)	(0.392)	(0.312)	(0.294)	(0.272)	(0.143)
	[N=132]	[N=170]	[N=130]	[N=148]	[N=159]	[N=739]
Stayers	6.185	6.018	6.313	5.849	6.026	6.082
	(0.121	(0.119)	(0.136)	(0.106)	(0.107)	(0.053)
	[N=1,065]	[N=1,006]	[N=1,077]	[N=1,018]	[N=1,014v	[N=5,180]
Total	6.107	5.990	6.220	5.879	5.953	6.031
	(0.112	(0.117)	(0.126)	(0.100)	(0.100)	(0.050)
	[N=1,197]	[N=1,176]	[N=1,207]	[N=1,166]	[N=1,173]	[N=5,919]

Notes: This table shows average baseline consumption across groups for stayers and drop-outs. Columns 1 to 5 show averages for the control group and the treatment groups and column 6 shows the overall average.
Source: Own calculations based on electricity meter readings.

Table 3.B4: Household characteristics

	$G0$	$G1$	$G2$	$G3$	$G4$	Total	p-value
Heating system							
Oil	0.398	0.413	0.414	0.418	0.414	0.412	0.920
Gas	0.207	0.213	0.210	0.229	0.242	0.220	0.286
Electric	0.053	0.029	0.040	0.049	0.046	0.043	0.040
Wood	0.020	0.016	0.014	0.020	0.021	0.018	0.664
Heat pump heating	0.059	0.060	0.057	0.050	0.047	0.055	0.653
Other	0.109	0.094	0.101	0.092	0.090	0.097	0.670
Don't know	0.202	0.217	0.214	0.189	0.197	0.204	0.494
Water heating system							
Electric	0.106	0.104	0.106	0.119	0.103	0.107	0.811
Partially electric	0.023	0.028	0.022	0.028	0.023	0.025	0.811
Heat pump	0.047	0.052	0.052	0.041	0.039	0.046	0.472
Oil	0.207	0.205	0.220	0.214	0.203	0.210	0.871
Gas	0.172	0.177	0.161	0.185	0.190	0.177	0.462
District heat	0.134	0.132	0.120	0.116	0.131	0.126	0.691
Wood	0.007	0.008	0.008	0.009	0.011	0.008	0.861
Other	0.030	0.031	0.038	0.041	0.040	0.036	0.519
Don't know	0.333	0.337	0.331	0.321	0.330	0.330	0.956
Appliances							
Washing machine	0.998	0.996	0.997	1.000	0.998	0.998	0.026
Dryer	0.830	0.837	0.832	0.824	0.844	0.833	0.820
Computer	0.991	0.993	0.990	0.991	0.997	0.992	0.199
Television	0.833	0.821	0.851	0.818	0.832	0.831	0.294
Dishwasher	0.735	0.750	0.733	0.743	0.754	0.743	0.731
Coffee machine	0.656	0.654	0.666	0.683	0.655	0.663	0.565
Microwave	0.367	0.364	0.374	0.356	0.355	0.363	0.884
Water boiler	0.672	0.676	0.693	0.689	0.670	0.680	0.745
Humidifier	0.176	0.157	0.178	0.158	0.168	0.167	0.562
Print/Copying machine	0.833	0.825	0.835	0.824	0.834	0.830	0.941

Table 3.B4 – continued

	$G0$	$G1$	$G2$	$G3$	$G4$	Total	p-value
Internet modem	0.931	0.925	0.929	0.937	0.933	0.931	0.859
Telephone	0.588	0.593	0.581	0.585	0.576	0.585	0.959
Mobile phone	0.944	0.954	0.950	0.947	0.945	0.948	0.839
Radio	0.753	0.757	0.743	0.747	0.736	0.747	0.799
Electronic photo frame	0.039	0.036	0.040	0.048	0.037	0.040	0.722
Stereo equipment	0.761	0.758	0.753	0.754	0.773	0.759	0.843
Video player	0.385	0.422	0.407	0.392	0.397	0.401	0.493
Video game console	0.204	0.179	0.181	0.194	0.195	0.190	0.583
Electric heating	0.046	0.048	0.060	0.060	0.056	0.054	0.556
Aquarium	0.019	0.024	0.026	0.025	0.022	0.023	0.743
Sauna	0.006	0.006	0.007	0.007	0.007	0.007	0.991
Solarium	0.003	0.001	0.001	0.002	0.001	0.002	0.835
Use of appliances							
Washing machine[a]	3.840	3.792	3.829	3.742	3.852	3.811	0.149
Dryer[b]	3.526	3.611	3.505	3.563	3.585	3.558	0.362
Computer[c]	2.761	2.908	2.829	2.902	2.745	2.829	0.028
Television[c]	2.595	2.705	2.700	2.706	2.678	2.677	0.330
Dishwasher[d]	1.938	1.937	1.976	1.896	1.894	1.929	0.194
Standby mode[e]	2.111	2.100	2.065	2.058	2.102	2.087	0.553
Energy-saving bulbs[f]	2.928	2.930	2.953	3.008	2.989	2.961	0.454
Cooking							
Gas	0.067	0.065	0.061	0.073	0.051	0.063	0.315
Electric	0.931	0.933	0.938	0.925	0.948	0.935	0.273
Other	0.002	0.002	0.001	0.002	0.001	0.002	0.937
Use of covers[g]	1.299	1.291	1.306	1.299	1.291	1.297	0.973
Lunch at home[h]	2.265	2.216	2.266	2.214	2.174	2.227	0.327
Dinner at home[h]	3.279	3.285	3.272	3.259	3.241	3.267	0.870

Notes: Columns 1 to 5 show averages of the items for groups $G0$ to $G4$. Column 6 shows the average across all groups and column 7 displays the p-values of an F-test of inequality of means across groups.

[a] 1 = Never; 2 = 1-2 times; 3 = 3-5 times; 4 = 6-9 times; 5 = 10-15 times; 6 = $>$ 15 times (per week)

[b] 1 = Always; 2 = 75 % of times; 3 = 50 % of times; 4 = 25 % of times; 5 = Never

[c] 1 = 0-1 hours; 2 = 1-2 hours; 3 = 2-3 hours; 4 = 3-4 hours; 5 = $>$ 4 hours

[d] 1 = 0-1 times per week; 2 = 2-3 times per week; 3 = $>$ 3 times per week

[e] 1 = None; 2 = 1-3 appliances; 3 = 4-6 appliances; 4 = 7-9 appliances; 5 = 10 or more appliances

[f] 1 = None; 2 = 25 %; 3 = 50 %; 4 = 75 %; 5 = 100 %

[g] 1 = Mostly; 2 = Sometimes; 3 = Rarely; 4 = I don't cook

[h] 1 = Never; 2 = 1-2 days per week; 3 = 3-4 days per week; 4 = 5-6 days per week; 5 = Always

Source: Own calculations based on survey data.

Table 3.B5: Attitudes, knowledge, and energy conserving behavior

	$G0$	$G1$	$G2$	$G3$	$G4$	Total	p-value
Attitudes: 1 = "Do not agree" and 5 = "Agree"							
Customization of environment 3.12	3.154	3.085	3.08	3.081	3.104	0.514	
Environmental abuse	4.205	4.194	4.123	4.147	4.196	4.173	0.302
Co-existence of flora and fauna	4.182	4.240	4.232	4.227	4.256	4.227	0.688
Saturation limits reached	4.135	4.225	4.124	4.185	4.170	4.168	0.185
Ecological awareness	3.379	3.367	3.333	3.328	3.363	3.354	0.717
Ecological awareness of others	3.772	3.811	3.730	3.679	3.732	3.745	0.041
Political engagement	3.720	3.817	3.738	3.717	3.788	3.756	0.140
Environmental pollution	4.037	4.124	4.019	3.980	4.023	4.037	0.082

Table 3.B5 – continued

	$G0$	$G1$	$G2$	$G3$	$G4$	Total	p-value
Climate warming	4.018	4.108	3.975	4.024	4.072	4.039	0.044
Conservation vs. job loss	3.462	3.580	3.410	3.565	3.473	3.498	0.001
Risks: 1 = "No risk" und 5 = "Large risk"							
Genetic engineering	3.354	3.369	3.319	3.386	3.342	3.354	0.651
Nuclear power	4.129	4.209	4.140	4.120	4.118	4.144	0.201
Radio antennas	2.745	2.763	2.786	2.767	2.784	2.769	0.904
High voltage power lines	2.691	2.678	2.687	2.702	2.712	2.694	0.943
Climate warming	4.112	4.180	4.109	4.153	4.158	4.142	0.255
Automobile traffic	3.816	3.846	3.790	3.819	3.839	3.822	0.597
Electricity conserving behavior							
% with past expert advice	0.024	0.024	0.035	0.039	0.037	0.032	0.078
Do you consider energy efficiency when buying ...							
... large appliances	0.975	0.977	0.974	0.970	0.971	0.973	0.898
... small appliances	0.672	0.671	0.653	0.677	0.661	0.667	0.803

Notes: Columns 1 to 5 show averages of the items for groups $G0$ to $G4$. Column 6 shows the average across all groups and column 7 displays the p-values of an F-test of inequality of means across groups.
Source: Own calculations based on survey data.

Table 3.B6: Personality and choice experiments

	$G0$	$G1$	$G2$	$G3$	$G4$	Total	p-value
Personality: (based on the 60-item version of HEXACO)							
Honesty	3.749	3.759	3.769	3.737	3.758	3.754	0.853
Emotionality	2.862	2.861	2.863	2.874	2.875	2.867	0.976
Extraversion	3.561	3.542	3.544	3.541	3.584	3.554	0.417
Agreeableness	3.227	3.219	3.239	3.238	3.248	3.233	0.796
Conscientiousness	3.724	3.699	3.704	3.715	3.671	3.703	0.328
Openness	3.696	3.677	3.652	3.674	3.661	3.672	0.612
Personality: (based on *The Swiss Environmental Survey 2007*)							
Trust in others	3.366	3.362	3.331	3.326	3.314	3.34	0.630
Contact to neighbors	2.992	2.906	2.925	2.93	2.935	2.937	0.432
Watch others	3.607	3.492	3.521	3.533	3.466	3.524	0.037
Compare with others	2.737	2.777	2.738	2.698	2.683	2.727	0.260
Compare notes with others	4.101	4.097	4.095	4.132	4.097	4.104	0.849
Dependence	3.247	3.286	3.252	3.214	3.203	3.241	0.408
Measure with others	2.729	2.784	2.727	2.709	2.711	2.732	0.477
Measure own performance	3.908	3.816	3.858	3.782	3.806	3.834	0.020
Public goods game							
Unconditional contribution	71.403	70.4	69.315	72.381	71.246	70.97	0.502
Conditional contribution: Contribution if others contribute ...							
... 0-20 units	25.232	26.139	26.117	26.323	24.913	25.739	0.873
... 21-40 units	38.965	39.334	39.933	40.227	38.519	39.375	0.762
... 41-60 units	54.32	54.282	54.803	54.888	53.339	54.315	0.794
... 61-80 units	68.957	67.891	68.538	68.691	67.413	68.291	0.839
... 81-100 units	82.916	80.344	81.979	81.901	80.354	81.483	0.487

Table 3.B6 – continued

	$G0$	$G1$	$G2$	$G3$	$G4$	Total	p-value
Overconfidence[18]							
Electricity consumption							
% larger	25.522	25.053	25.15	29.179	29.389	26.797	0.000
% equal	35.317	31.363	32.188	33.526	34.412	33.337	0.000
% smaller	39.442	43.746	42.301	37.792	37.053	40.155	0.000
Knowledge of energy-relevant issues							
% larger	40.577	41.021	42.657	43.455	42.684	42.03	0.035
% equal	37.125	35.412	35.528	35.187	36.066	35.869	0.204
% smaller	23.072	23.965	22.834	22.638	23.258	23.173	0.458
Household income							
% larger	35.52	34.675	34.952	34.47	35.146	34.952	0.893
% equal	33.406	32.769	33.361	33.048	32.489	33.014	0.825
% smaller	30.952	32.655	32.255	33.172	32.648	32.327	0.315
Warm water consumption							
% larger	27.256	25.327	26.661	27.24	26.733	26.612	0.293
% equal	34	34.81	34.014	33.267	34.279	34.094	0.665
% smaller	40.613	41.299	40.609	40.006	40.803	40.683	0.909
Time preferences (waiting period measured in weeks)							
Decision between 100 units today or ...							
... 110 units in X weeks	22.523	21.931	22.352	21.627	21.777	22.043	0.824
... 120 units in X weeks	24.19	23.81	24.385	23.767	23.985	24.02	0.950
Decision between 100 units in 4 weeks or ...							
... 110 units in 4+X weeks	22.677	22.136	22.485	22.396	22.005	22.338	0.939
... 120 units in 4+X weeks	24.658	24.145	24.656	24.238	24.127	24.361	0.939
Time preferences (proportion with maximum waiting period of 46 weeks)							
Decision between 100 units today or ...							
...110 units in X weeks	0.22	0.227	0.214	0.226	0.199	0.217	0.711
...120 units in X week	0.229	0.24	0.237	0.238	0.228	0.235	0.977
Decision between 100 units in 4 weeks or ...							
... 110 units in 4+X weeks	0.215	0.22	0.205	0.228	0.202	0.214	0.759
... 120 units in 4+X weeks	0.227	0.245	0.23	0.247	0.234	0.237	0.873

Notes: Columns 1 to 5 show averages of the items for groups $G0$ to $G4$. Column 6 shows the average across all groups and column 7 displays the p-values of an F-test of inequality of means across groups.
Source: Own calculations based on survey data.

[18] This choice experiment was integrated into the Q2 online survey for the questions regarding electricity consumption, knowledge and household income. The question regarding warm water consumption was asked in the post-treatment survey.

Table 3.B7: Hourly treatment effects

	Electricity consumption kWh per hour			Electricity expenditures CHF per hour		
	Overall	Frequent users	Infrequent users	Overall	Frequent users	Infrequent users
0h00	0.020***	0.022***	0.016***	0.003***	0.004***	0.003***
	(0.004)	(0.007)	(0.006)	(0.001)	(0.001)	(0.001)
1h00	0.012***	0.013***	0.011**	0.002***	0.003***	0.003***
	(0.004)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
2h00	-0.000	-0.003	0.000	0.001	0.000	0.001
	(0.006)	(0.010)	(0.008)	(0.001)	(0.001)	(0.001)
3h00	0.021***	0.027***	0.012	0.004***	0.004***	0.003**
	(0.007)	(0.009)	(0.009)	(0.001)	(0.001)	(0.001)
4h00	0.010**	0.013**	0.007	0.002***	0.002***	0.002**
	(0.005)	(0.005)	(0.007)	(0.001)	(0.001)	(0.001)
5h00	0.003	0.004	0.002	0.001**	0.001**	0.002*
	(0.003)	(0.004)	(0.006)	(0.001)	(0.001)	(0.001)
6h00	-0.003	-0.003	-0.001	-0.000	-0.000	0.000
	(0.003)	(0.004)	(0.005)	(0.001)	(0.001)	(0.001)
7h00	-0.012***	-0.013***	-0.010*	-0.002***	-0.002***	-0.002*
	(0.003)	(0.004)	(0.005)	(0.001)	(0.001)	(0.001)
8h00	-0.012***	-0.013***	-0.010**	-0.002***	-0.002**	-0.002**
	(0.003)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
9h00	-0.011***	-0.010**	-0.011**	-0.002***	-0.002*	-0.002**
	(0.003)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
10h00	-0.011***	-0.010**	-0.011**	-0.002**	-0.002*	-0.002*
	(0.003)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
11h00	-0.015***	-0.014***	-0.015***	-0.003***	-0.003***	-0.003***
	(0.004)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
12h00	-0.012***	-0.011**	-0.013**	-0.002***	-0.002**	-0.002**
	(0.004)	(0.006)	(0.005)	(0.001)	(0.001)	(0.001)
13h00	-0.010***	-0.009*	-0.011**	-0.002**	-0.002	-0.002**
	(0.004)	(0.006)	(0.005)	(0.001)	(0.001)	(0.001)
14h00	-0.010***	-0.010*	-0.010**	-0.002**	-0.002*	-0.002*
	(0.004)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
15h00	-0.005	-0.005	-0.006	-0.001	-0.001	-0.001
	(0.003)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
16h00	-0.012***	-0.014***	-0.010**	-0.002***	-0.003***	-0.002*
	(0.004)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
17h00	-0.031***	-0.038***	-0.027***	-0.006***	-0.007***	-0.005***
	(0.004)	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
18h00	-0.047***	-0.050***	-0.047***	-0.009***	-0.010***	-0.009***
	(0.005)	(0.008)	(0.008)	(0.001)	(0.001)	(0.001)
19h00	-0.041***	-0.050***	-0.032***	-0.008***	-0.010***	-0.006***
	(0.006)	(0.008)	(0.008)	(0.001)	(0.002)	(0.002)
20h00	-0.039***	-0.052***	-0.025***	-0.007***	-0.010***	-0.004***
	(0.005)	(0.008)	(0.007)	(0.001)	(0.002)	(0.001)

Table 3.B7 – continued

	Electricity consumption kWh per hour			Electricity expenditures CHF per hour		
	Overall	Frequent users	Infrequent users	Overall	Frequent users	Infrequent users
21h00	-0.028***	-0.041***	-0.015**	-0.005***	-0.008***	-0.003**
	(0.004)	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
22h00	-0.016***	-0.030***	-0.002	-0.002*	-0.003**	0.001
	(0.006)	(0.010)	(0.006)	(0.001)	(0.001)	(0.001)
23h00	0.016***	0.015*	0.015**	0.003***	0.003***	0.003***
	(0.005)	(0.008)	(0.007)	(0.001)	(0.001)	(0.001)
N	7,435,617	4,531,535	2,904,082	7,435,617	4,531,535	2,904,082
Clusters	1,003	506	497	1,003	506	497
Hourly FE	Yes	Yes	Yes	Yes	Yes	Yes
Weekly FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows deviations from hourly mean for electricity consumption (kWh per hour) and electricity expenditures (CHF per hour). Means are based on observations *before* installation of the smart meter display. Columns 1 and 4 show overall estimates, columns 2 and 5 show estimates for frequent users of the display and columns 3 and 6 for infrequent users respectively. All estimates contain hourly and weekly FE. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on smart meter data.

Table 3.B8: Potential for changing energy-efficient behavior

Treatment group	Light bulbs	Tele-vision	Laundry	Fridge	Turning-off	Power strips	Average
Smart metering	0.029	-0.063	0.049	-0.041	-0.012	-0.001	-0.008
	(0.033)	(0.045)	(0.038)	(0.034)	(0.026)	(0.030)	(0.016)
Expert advice	0.096***	-0.029	0.070*	0.051	0.046	-0.009	0.037**
	(0.036)	(0.043)	(0.038)	(0.036)	(0.028)	(0.032)	(0.016)
Social competition	0.021	-0.052	-0.053	-0.039	-0.046*	-0.020	-0.035**
	(0.034)	(0.046)	(0.038)	(0.035)	(0.026)	(0.031)	(0.016)
Social comparison	-0.001	-0.008	0.026	-0.040	-0.007	0.025	-0.002
	(0.034)	(0.046)	(0.040)	(0.035)	(0.028)	(0.032)	(0.017)
Observations	14,861	14,643	14,878	14,785	14,806	14,835	14,959
R-squared	0.010	0.004	0.005	0.007	0.007	0.006	0.010

Notes: This table presents DiD estimates for the household's potential to change their energy efficient behavior. These changes comprise the ability to change most of the existing light bulbs with energy-efficient light bulbs (column 1), to watch television 2 hours less per week (column 2), to dry clothes on a clothes line instead of using a dryer (column 3), to increase the temperature of the fridge by 1^{circ} Celsius (column 4), to turn off the power of appliances after using them (column 5), and to install power strips for appliances (column 6). Column 7 comprises an average over all six items. Standard errors clustered by household in parentheses. *** P<0.01 ** P<0.05 * P<0.1.
Source: Own calculations based on survey data.

3.C Additional Figures

Figure 3.C1: Cumulative take-up rate of expert advice

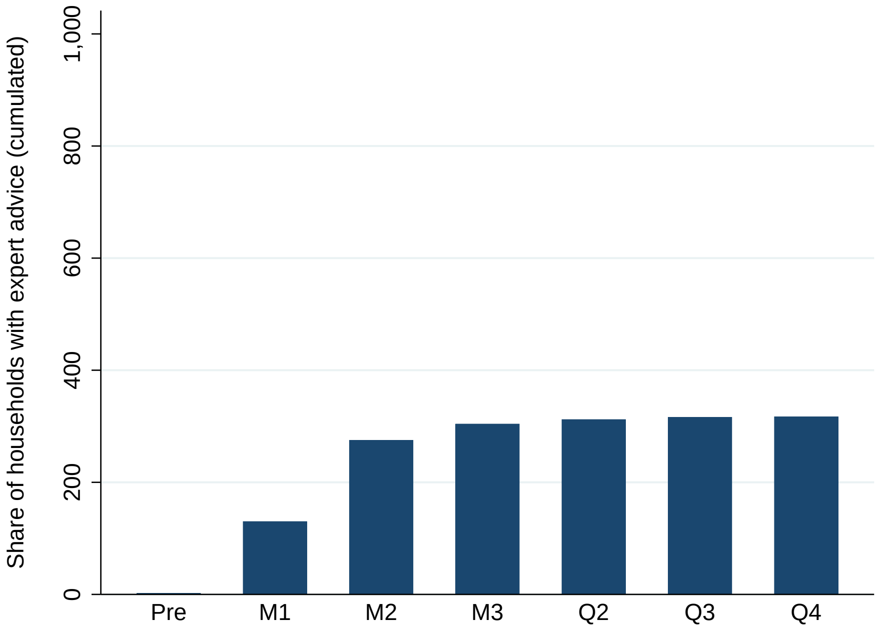

Notes: This figure shows the cumulative take-up rate of expert advice over the study period.

Figure 3.C2: Smart meter device

Notes: The top left screen shows the start screen of the device. It shows a household's real-time electricity use. Also it allows households to monitor their achievement of their weekly goal setting. The top right screen shows a screen which allowed households to observe their historical electricity use at different measurement frequencies. On the bottom left screen households can set weekly goals and monitor the achievement of their weekly goalsetting over time. The bottom right screen allows households to measure electricity consumption in a given period.
Source: *ewz* Smart Metering field experiment

Figure 3.C3: Sample letter for social competition treatment

Ihr Stromverbrauch im Vergleich.

Sehr geehrter Herr Muster

Hiermit senden wir Ihnen den <<ersten/nächsten>> Zwischenstand. Sie erhalten regelmässig Informationen zu Ihrem Stromverbrauch und zum Stromverbrauch ihres Vergleichshaushalts. Dieser erhält im Gegenzug die gleichen anonymisierten Informationen. Die Vergleichsgrafik zeigt Ihnen den prozentualen Unterschied Ihres Gesamtverbrauchs. Ist der Balken rot, dann verbrauchten Sie mehr Strom, ist der Balken grün, dann haben Sie weniger Strom als der Vergleichshaushalt konsumiert.

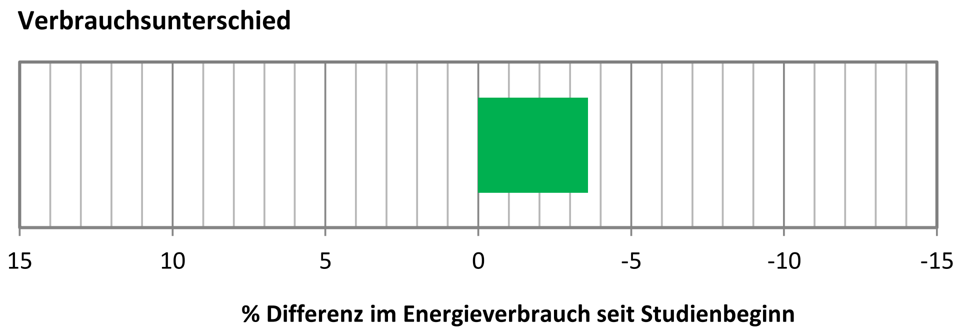

Sie verbrauchen X% «mehr/weniger» Strom als ihr Vergleichshaushalt

In der unten stehenden Grafik sehen Sie die Entwicklung Ihres Stromverbrauchs.

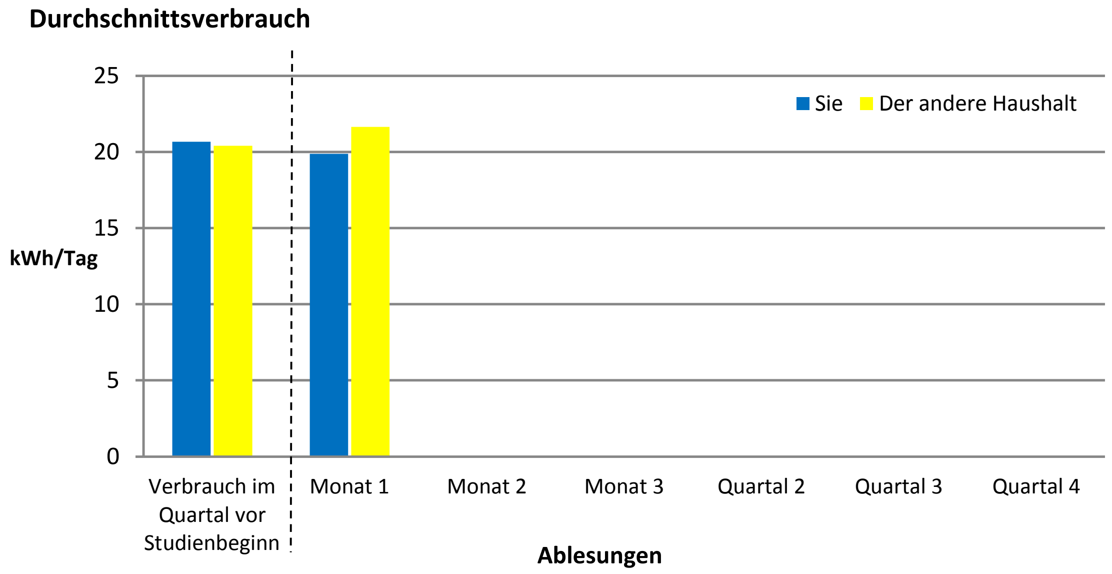

Ihr durchschnittlicher Verbrauch pro Tag im letzten <<Quartal/Monat>> betrug **<<x>> kWh/Tag**. Ihr Vergleichshaushalt hat in derselben Zeitperiode <<**x>> kWh/Tag** konsumiert. Damit haben Sie im vergangenen Quartal **<<x>>kWh/Tag <<mehr/weniger>>** verbraucht.

In rund <<4 Wochen/3 Monaten>> werden wir Ihnen erneut einen Zwischenstand zusenden, der Sie darüber informiert, wo sie gegenüber dem anderen Haushalt stehen.

Figure 3.C4: Sample letter for social comparison treatment

Ihr Stromverbrauch im Vergleich.

Sehr geehrter Herr Muster

Hiermit senden wir Ihnen den <<ersten/nächsten>> Zwischenstand. Sie erhalten regelmässig Informationen zu Ihrem Stromverbrauch und zum Stromverbrauch ihres Vergleichshaushalts. Dieser erhält diese Informationen jedoch nicht.

Die Vergleichsgrafik zeigt Ihnen den prozentualen Unterschied Ihres Gesamtverbrauchs. Ist der Balken rot, dann verbrauchten Sie mehr Strom, ist der Balken grün, dann haben Sie weniger Strom als der Vergleichshaushalt konsumiert.

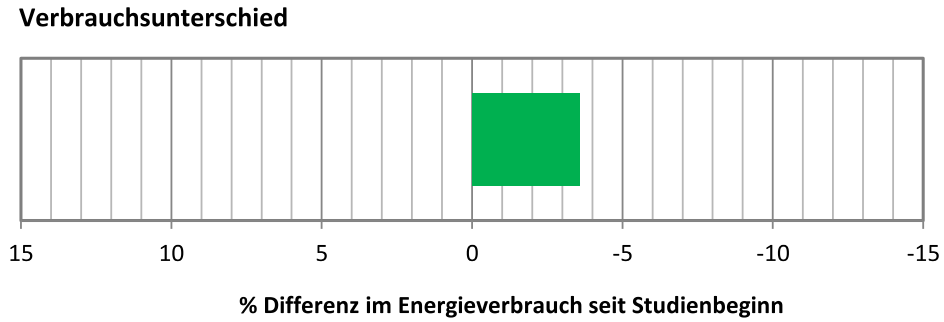

Sie verbrauchen X% «mehr/weniger» Strom als ihr Vergleichshaushalt

In der unten stehenden Grafik sehen Sie die Entwicklung Ihres Stromverbrauchs.

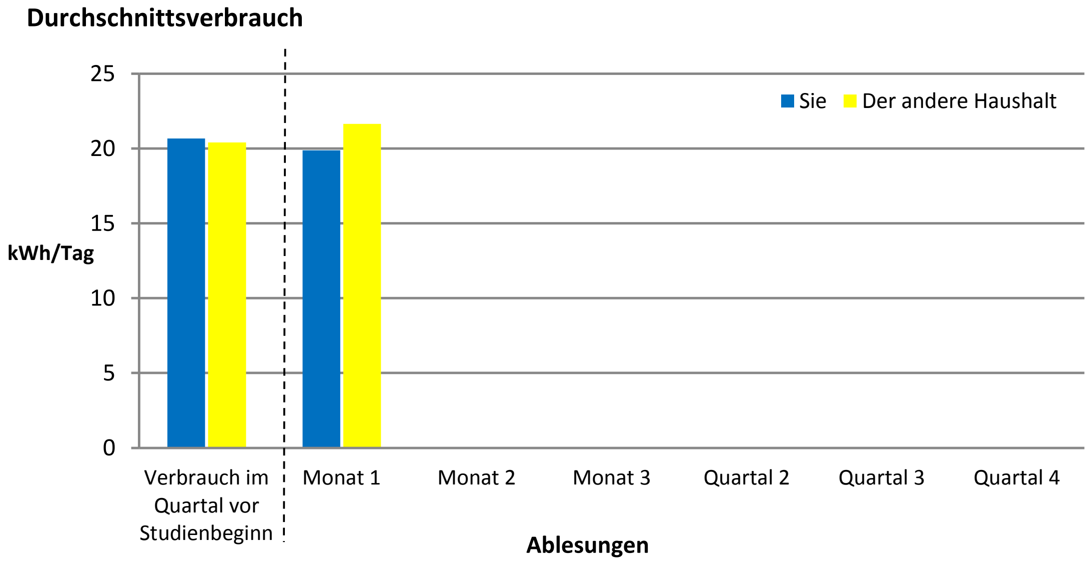

Ihr durchschnittlicher Verbrauch pro Tag im letzten <<Quartal/Monat>> betrug **<<xy>> kWh/Tag**. Ihr Vergleichshaushalt hat in derselben Zeitperiode <<**xy>> kWh/Tag** konsumiert. Damit haben Sie im vergangenen Quartal **<<xy>>kWh/Tag <<mehr/weniger>>** verbraucht.

In rund <<4 Wochen/3 Monaten>> werden wir Ihnen erneut einen Zwischenstand zusenden, der Sie darüber informiert, wo sie gegenüber dem anderen Haushalt stehen.

General Conclusion

Over thirty years ago, Leamer (1983) – among many others – expressed doubts about the quality and usefulness of empirical analyses for the economic profession by stating that "hardly anyone takes data analyses seriously. Or perhaps more accurately, hardly anyone takes anyone else's data analyses seriously" (p.37). Improvements in data quality, more robust estimation methods and the evolution of better research designs seem to make that assertion no longer justifiable (see Angrist and Pischke (2010) for a recent response to Leamer's essay). The economic profession and policy makers alike often rely on empirical evidence as a means to investigate policy relevant questions. The approach of using scientifically rigorous and systematic evidence to identify policies and programs that are capable of improving policy-relevant outcomes is known under the increasingly popular notion of evidence-based policy.

Evidence-based economic policy often relies on randomized or quasi-natural experiments in order to identify causal effects of policies. These can require relatively strong assumptions or raise concerns of external validity. In the context of this thesis, potential concerns are for example endogeneity of policy reforms with respect to the business cycle in the first chapter, the trade-off between precision and bias in the regression-discontinuity setting in chapter 2 or non-representativeness of the sample due to self-selection in chapter 3. While the identification strategies are very useful to gain insights into the causal effects of specific policy questions, transforming the evidence into concrete policy conclusions can be challenging. Policy development should therefore rely on the systematic evidence of a whole body of research on a specific policy question rather than on a single analysis. In this sense, this thesis cannot and should not be viewed as a comprehensive analysis of specific policy issues but rather as a first step towards a better understanding of certain aspects of a policy question.

The thesis applies new and innovative identification strategies to policy-relevant and topical questions in the fields of labor economics and behavioral environmental economics. Each chapter relies on a different identification strategy. In the first chapter, we employ a difference-in-differences approach to exploit the quasi-experimental change in the entitlement of the maximum unemployment benefit duration to identify the medium-run effects of reduced benefit durations on post-unemployment outcomes. Shortening benefit duration carries a double-dividend: It generates fiscal benefits without deteriorating the quality of job-matches. On the contrary, shortened benefit durations improve medium-run earnings and employment possibly through containing the negative effects of skill depreciation or stigmatization.

While the first chapter provides only indirect evidence on the underlying behavioral channels, in the second chapter I develop a novel approach that allows to learn about the relative importance of the two key margins of job search – reservation wage choice and search effort. In the framework of a standard non-stationary job search model, I show how the exit rate from unemployment can be decomposed in a way that is informative on reservation wage movements over the unemployment spell. The empirical analysis relies on a sharp discontinuity in unemployment benefit entitlement, which can be exploited in a regression-discontinuity approach to identify the effects of extended benefit durations on unemployment and survivor functions. I find evidence that calls for an important role of reservation wage choices for job search behavior. This can have direct implications for the optimal design of unemployment insurance policies.

The third chapter – while thematically detached from the other chapters – addresses one of the major policy challenges of the 21^{st} century: climate change and resource consumption. Many governments have recently put energy efficiency on top of their agendas. While pricing instruments aimed at regulating the energy demand have often been found to be short-lived and difficult to enforce politically, the focus of energy conservation programs has shifted towards behavioral approaches – such as provision of information or social norm feedback. The third chapter describes a randomized controlled field experiment in which we discuss the effectiveness of different types of feedback on residential electricity consumption. We find that detailed and real-time feedback caused persistent electricity reductions on the order of 3 to 5 % of daily electricity consumption. Also social norm information can generate substantial electricity savings when designed appropriately. The findings suggest that behavioral approaches constitute effective and relatively cheap way of improving residential energy-efficiency.

Bibliography

Abrahamse, W., Steg, L., Vlek, C., and Rothengatter, T. (2005): "A review of intervention studies aimed at household energy conservation." *Journal of Environmental Psychology*, 25(3): 273–291.

Acemoglu, D. and Autor, D. (2011): "Skills, tasks and technologies: Implications for employment and earnings." In O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 4, chapter 12, pages 1043–1171. Elsevier.

Adamchik, V. (1999): "The effect of unemployment benefits on the probability of re-employment in Poland." *Oxford Bulletin of Economics and Statistics*, 61: 95–108.

Addison, J. T. and Blackburn, M. L. (2000): "The effects of unemployment insurance on postunemployment earnings." *Labour Economics*, 7(1): 21–53.

Addison, J. T. and Portugal, P. (2004): "How does the unemployment insurance system shape the time profile of jobless duration?" *Economics Letters*, 85(2): 229–234.

Allcott, H. (2011): "Social norms and energy conservation." *Journal of Public Economics*, 95(9): 1082–1095.

Allcott, H. and Mullainathan, S. (2010): "Behavior and energy policy." *Science*, 327(5970): 1204–1205.

Allcott, H. and Rogers, T. (forthcoming): "The short-run and long-run effects of behaviroal interventions: Experimental evidence from energy conservation." *American Economic Review*.

Angrist, J. D. and Pischke, J.-S. (2010): "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives*, 24(2): 3–30.

Ariely, D., Bracha, A., and Meier, S. (2009): "Doing good or doing well? Image motivation and monetary incentives in behaving prosocially." *American Economic Review*, 99(1): 544–555.

Arni, P., van Ours, J. C., and Lalive, R. (2013): "How effective are unemployment benefit sanctions? Looking beyond unemployment exit." *Journal of Applied Econometrics*, 28(7): 1153–1178.

Ashton, M. C. and Lee, K. (2009): "The HEXACO-60: A short measure of the major dimensions of personality." *Journal of Personality Assessment*, 91(4): 340–345.

Attari, S. Z., DeKay, M. L., Davidson, C. I., and Bruine de Bruin, W. (2010): "Public perceptions of energy consumption and savings." *Proceedings of the National Academy of Sciences*, 107(37): 16054–16059.

Ayres, I., Raseman, S., and Shih, A. (2013): "Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage." *Journal of Law, Economics, and Organization*, 29(5): 992–1022.

Baeriswyl, M., Müller, A., Rigassi, R., Rissi, C., Solenthaler, S., Staake, T., and Weisskopf, T. (2012): "Folgeabschätzung einer Einführung von *Smart Metering* im Zusammenhang mit *Smart Grids* in der Schweiz." Technical report, Bundesamt für Energie.

Baily, M. N. (1978): "Some aspects of optimal unemployment compensation." *Journal of Public Economics*, 10(3): 379–402.

Barber, B., Odean, T., and Zheng, L. (2005): "Out of sight, out of mind: The effects of expenses on mutual fund flows." *Journal of Business*, 78: 2095–2102.

Blanchard, O. J. and Diamond, P. A. (1994): "Ranking, unemployment duration, and wages." *Review of Economic Studies*, 61(3): 417–34.

Blunkett, D. (2000): "Influence or irrelevance: can social science improve government." *Research Intelligence*, 71: 12–21.

Bollinger, B. and Gillingham, K. (2012): "Peer effects in the diffusion of solar photovoltaic panels." *Marketing Science*, 31(6): 900–912.

Cahuc, P. and Zylberberg, A. (2004): *Labor Economics*. MIT Press, Cambridge, MA.

Caliendo, M., Tatsiramos, K., and Uhlendorff, A. (2013a): "Benefit duration, unemployment duration and job match quality: A regression-discontinuity approach." *Journal of Applied Econometrics*, 28(4): 604–627.

Caliendo, M., Tatsiramos, K., and Uhlendorff, A. (2013b): "Benefit duration, unemployment duration and job match quality: A regression-discontinuity approach." *Journal of Applied Econometrics*, 28(4): 604–627.

Card, D., Chetty, R., and Weber, A. (2007a): "Cash-on-nand and competing models of intertemporal behavior: New evidence from the labor market." *The Quarterly Journal of Economics*, 122(4): 1511–1560.

Card, D., Chetty, R., and Weber, A. (2007b): "The spike at benefit exhaustion: Leaving the unemployment system or starting a new job?" *American Economic Review*, 97(2): 113–118.

Card, D., Chetty, R., and Weber, A. (2007c): "The spike at benefit exhaustion: Leaving the unemployment system or starting a new job?" *American Economic Review*, 97(2): 113–118.

Card, D. E. and Levine, P. B. (2000): "Extended benefits and the duration of UI spells: Evidence from the New Jersey extended benefit program." *Journal of Public Economics*, 78(1): 107–138.

Carling, K., Edin, P.-A., Harkman, A., and Holmlund, B. (1996): "Unemployment duration, unemployment benefits, and labor market programs in Sweden." *Journal of Public Economics*, 59(3): 313–334.

Centeno, M. and Novo, A. (2006): "The impact of unemployment insurance on the job match quality: A quantile regression approach." *Empirical Economics*, 31(4): 905–919.

Centeno, M. and Novo, A. (2009): "Reemployment wages and UI liquidity effect: a regression discontinuity approach." *Portuguese Economic Journal*, 8(1): 45–52.

Chetty, R. (2008): "Moral hazard versus liquidity and optimal unemployment insurance." *Journal of Political Economy*, 116(2): 173–234.

Chetty, R. (2009): "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods." *Annual Review of Economics*, 1(1): 451–488.

Chetty, R., Looney, A., and Kroft, K. (2009): "Salience and taxation: Theory and evidence." *American Economic Review*, 99(4): 1145–77.

Clark, K. B. and Summers, L. H. (1981): "Demographic differences in cyclical employment variation." *Journal of Human Resources*, 16(1): 61–79.

Costa, D. L. and Kahn, M. E. (2013): "Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment." *Journal of the European Economic Association*, 11(3): 680–702.

Cragg, M. I., Zhou, Y., Gurney, K., and Kahn, M. E. (2013): "Carbon geography: The political economy of congressional support for legislation intended to mitigate greenhouse gas production." *Economic Inquiry*, 51(2): 1640–1650.

Degen, K., Efferson, C., Frei, F., Goette, L., and Lalive, R. (2013): "Smart Metering, Beratung oder Sozialer Vergleich. Was beeinflusst den Elektrizitätsverbrauch?" Technical report, Bundesamt für Energie (Publikation 290850).

DellaVigna, S. (2009): "Psychology and economics: Evidence from the field." *Journal of Economic Literature*, 47(2): 315–72.

DellaVigna, S., List, J. A., and Malmendier, U. (2012): "Testing for altruism and social pressure in charitable giving." *The Quarterly Journal of Economics*, 127(1): 1–56.

DellaVigna, S. and Paserman, M. D. (2005): "Job search and impatience." *Journal of Labor Economics*, 23(3): 527–588.

Diekmann, A., Meyer, R., Mühlemann, C., and Diem, A. (2009): "Schweizer Umweltsurvey 2007 - Dokumentation und Codebuch." Technical report, ETH Zürich Professur für Soziologie.

Dolan, P. and Metcalfe, R. (2013): "Neighbors, knowledge, and nuggets: Two natural field experiments on the role of incentives on energy conservation." *CEP Discussion Papers 1222*, Centre for Economic Performance, London School of Economics and Political Science, London, UK.

Duflo, E. and Saez, E. (2003): "The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment." *The Quarterly Journal of Economics*, 118(3): 815–842.

Ehrenberg, R. G. and Oaxaca, R. L. (1976): "Unemployment insurance, duration of unemployment, and subsequent wage gain." *American Economic Review*, 66(5): 754–66.

EIROnline (2013): "Austria: Industrial relations profile." Available at http://www.eurofound.europa.eu/eiro/country/austria_1.htm.

Epper, T., Fehr-Duda, H., and Schubert, R. (2011): "Energy-using durables: The role of time discounting in investment decisions." *IED Working paper 11-16*, IED Institute for Environmental Decisions, ETH Zurich.

European Commission (2012): "Energy Efficiency Directive (Directive 2012/27/EU)." Available at http://ec.europa.eu/energy/efficiency/eed/eed_en.htm.

Feldstein, M. and Poterba, J. (1984): "Unemployment insurance and reservation wages." *Journal of Public Economics*, 23(1-2): 141–167.

Ferraro, P. J. and Price, M. K. (2013): "Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment." *Review of Economics and Statistics*, 95(1): 64–73.

Fischbacher, U., Gächter, S., and Fehr, E. (2001): "Are people conditionally cooperative? Evidence from a public goods experiment." *Economics Letters*, 71(3): 397–404.

Fredriksson, P. and Holmlund, B. (2006): "Improving incentives in unemployment insurance: A review of recent research." *Journal of Economic Surveys*, 20(3): 357–386.

Frey, B. S. and Meier, S. (2004): "Social comparisons and pro-social behavior: Testing "conditional cooperation" in a field experiment." *American Economic Review*, 94(5): 1717–1722.

Frey, B. S. and Oberholzer-Gee, F. (1997): "The cost of price incentives: An empirical analysis of motivation crowding-out." *American Economic Review*, 87(4): 746–755.

Gardner, G. and Stern, P. (2008): "The short list: The most effective actions U.S. households can take to curb climate change." *Environment*, 50: 12–23.

Gerber, A. S. and Rogers, T. (2009): "Descriptive social norms and motivation to vote: everybody's voting and so should you." *Journal of Politics*, 71(01): 178–191.

Gerfin, M. and Lechner, M. (2002): "A Microeconometric Evaluation of the Active Labour Market Policy in Switzerland." *The Economic Journal*, 112(482): 854–893.

Gibbons, R. and Katz, L. (1992): "Does unmeasured ability explain inter-industry wage differences?" *Review of Economic Studies*, 59(3): 515–535.

Gneezy, U. and Rustichini, A. (2000): "Pay enough or don't pay at all." *The Quarterly Journal of Economics*, 115(3): 791–810.

Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. (2008): "A room with a viewpoint: Using social norms to motivate environmental conservation in hotels." *Journal of Consumer Research*, 35(3): 472–482.

Görlich, D. and de Grip, A. (2009): "Human capital depreciation during hometime." *Oxford Economic Papers*, 61: i98–i121.

Gosling, S. D., Rentfrow, P. J., and Swann Jr, W. B. (2003): "A very brief measure of the Big-Five personality domains." *Journal of Research in Personality*, 37(6): 504–528.

Green, D. and Riddell, W. (1993): "Qualifying for unemployment insurance: An empirical analysis." *UBC Departmental Archives 93-33*, UBC Department of Economics.

Green, D. A. and Riddell, W. C. (1997): "Qualifying for unemployment insurance: An empirical analysis." *Economic Journal*, 107(440): 67–84.

Griskevicius, V., Cialdini, R. B., and Goldstein, N. J. (2008): "Social norms: An underestimated and underemployed lever for managing climate change." *International Journal of Sustainability Communication*, 3: 5–13.

Grossman, J. B. (1989): "The work disincentive effect of extended unemployment compensation: Recent evidence." *Review of Economics and Statistics*, 71: 159–164.

Haan, P. and Prowse, V. L. (2010): "The design of unemployment transfers: Evidence from a dynamic structural life-cycle model." *IZA Discussion Papers 4792*, Institute for the Study of Labor (IZA).

Hahn, J., Todd, P., and Klaauw, W. V. d. (2001): "Identification and estimation of treatment effects with a regression-discontinuity design." *Econometrica*, 69(1): 201–209.

Ham, J. C. and LaLonde, R. J. (1996): "The Effect of Sample Selection and Initial Conditions in Duration Models: Evidence from Experimental Data on Training." *Econometrica*, 64(1): 175–205.

Ham, J. C. and Rea, J., Samuel A (1987): "Unemployment insurance and male unemployment duration in Canada." *Journal of Labor Economics*, 5(3): 325–53.

Hausman, J. A. and Joskow, P. L. (1982): "Evaluating the costs and benefits of appliance efficiency standards." *American Economic Review*, 72(2): 220–225.

Hunt, J. (1995): "The effect of unemployment compensation on unemployment duration in Germany." *Journal of Labor Economics*, 13(2): 88–120.

IEA (2013): "World Energy Outlook 2013." doi:10.1787/weo-2013-en.

ILO (2012): "Global employment trends: Preventing a deeper jobs crisis." Report, ILO: Geneva.

Imbens, G. and Kalyanaraman, K. (2012): "Optimal bandwidth choice for the regression discontinuity estimator." *Review of Economic Studies*, 79(3): 933–959.

Inderbitzin, L., Staubli, S., and Zweimuller, J. (2012): "Extended unemployment benefits and early retirement: Program complementarity and program substitution." *mimeo*, University of Zurich.

Katz, L. F. and Meyer, B. D. (1990): "The impact of the potential duration of unemployment benefits on the duration of unemployment." *Journal of Public Economics*, 41(1): 45–72.

Kroft, K., Lange, F., and Notowidigo, M. (2013): "Duration dependence and labor market conditions: Evidence from a field experiment." *The Quarterly Journal of Economics*, 128(3): 1123–1167.

Krueger, A. B. and Mueller (2013): "A contribution to the empirics of reservation wages." *Mimeo*.

Krueger, A. B. and Mueller, A. (2011a): "Job search, emotional well-being and job finding in a period of mass unemployment: Evidence from high-frequency longitudinal data." *Brookings Papers on Economic Activity*, 42(1): 1–81.

Krueger, A. B. and Mueller, A. I. (2011b): "Job search and job finding in a period of mass unemployment: Evidence from figh-frequency longitudinal data." Working Papers 562, Princeton University. Industrial Relations Section.

Kuhn, A., Lalive, R., and Zweimüller, J. (2009): "The public health costs of job loss." *Journal of Health Economics*, 28(6): 1099–1115.

Lalive, R. (2007): "Unemployment benefits, unemployment duration, and post-unemployment jobs: A regression discontinuity approach." *American Economic Review (Papers and Proceedings)*, 97(2): 108–112.

Lalive, R. (2008): "How do extended benefits affect unemployment duration A regression discontinuity approach." *Journal of Econometrics*, 142(2): 785–806.

Lalive, R., Landais, C., and Zweimüller, J. (2013): "Market externalities of large unemployment insurance extension programs." *Mimeo*.

Lalive, R., Ours, J. V., and Zweimüller, J. (2006): "How changes in financial incentives affect the duration of unemployment." *Review of Economic Studies*, 73(4): 1009–1038.

Lalive, R., Van Ours, J. C., and Zweimüller, J. (2008): "The Impact of Active Labour Market Programmes on The Duration of Unemployment in Switzerland*." *The Economic Journal*, 118(525): 235–257.

Lalive, R. and Zweimüller, J. (2004a): "Benefit entitlement and the labor market: Evidence from a large-scale policy change." In J. Agell, M. Keen, and A. Weichenrieder, editors, *Labor Market Institutions and Public Policy*, pages 63–100. Cambridge, Massachusetts: MIT Press.

Lalive, R. and Zweimüller, J. (2004b): "Benefit entitlement and unemployment duration: The role of policy endogeneity." *Journal of Public Economics*, 88(12): 2587–2616.

Le Barbanchon, T. (2012): "The effect of potential unemployment benefits duration on unemployment exits to work and on job quality." *Crest working paper 2012-21*, CREST Centre de Recherche en Economie et Statistique.

Leamer, E. E. (1983): "Let's take the con out of econometrics." *American Economic Review*, 73(1): 31–43.

Leamer, E. E. (2010): "Tantalus on the Road to Asymptopia." *Journal of Economic Perspectives*, 24(2): 31–46.

Lee, D. S. and Lemieux, T. (2010a): "Regression discontinuity designs in economics." *Journal of Economic Literature*, 48(2): 281–355.

Lee, D. S. and Lemieux, T. (2010b): "Regression discontinuity designs in economics." *Journal of Economic Literature*, 48(2): 281–355.

Lee, M.-J. and Kang, C. (2006): "Identification for difference in differences with cross-section and panel data." *Economics Letters*, 92(2): 270–276.

Levitt, S. D. and List, J. A. (2007): "What do laboratory experiments measuring social preferences reveal about the real world?" *Journal of Economic Perspectives*, 21(2): 153–174.

Li, H.-H. (2013): "The effects of human capital depreciation on gender segregation." Discussion paper, University of Wisconsin-Madison.

Ljungqvist, L. and Sargent, T. J. (1998): "The European unemployment dilemma." *Journal of Political Economy*, 106(3): 514–550.

Manski, C. F. (2013): *Public Policy in an Uncertain World: Analysis and Decisions*. Cambridge, MA: Harvard UP.

McCrary, J. (2008): “Manipulation of the running variable in the regression discontinuity design: A density test.” *Journal of Econometrics*, 142(2): 698–714.

Mellström, C. and Johannesson, M. (2008): “Crowding out in blood donation: Was Titmuss right?” *Journal of the European Economic Association*, 6(4): 845–863.

Merkle, C. and Weber, M. (2011): “True overconfidence: The inability of rational information processing to account for apparent overconfidence.” *Organizational Behavior and Human Decision Processes*, 116(2): 262–271.

Meyer, B. D. (1990): “Unemployment insurance and unemployment spells.” *Econometrica*, 58(4): 757–82.

Moffitt, R. A. (1985): “Unemployment insurance and the distribution of unemployment spells.” *Journal of Econometrics*, 28(1): 85–101.

Moffitt, R. A. and Nicholson, W. (1982): “The effect of unemployment insurance on unemployment: The case of federal supplemental benefits.” *Review of Economics and Statistics*, 64(1): 1–11.

Mortensen, D. (1977): “Unemployment insurance and job search decisions.” *Industrial and Labor Relations Review*, 30(4): 505–517.

Nekoei, A. and Weber, A. (2013): “Does extending unemployment benefits improve job quality?” *Mimeo*.

Oberholzer-Gee, F. (2008): “Nonemployment stigma as rational herding: A field experiment.” *Journal of Economic Behavior & Organization*, 65(1): 30–40.

Puhani, P. A. (2000): “Poland on the dole: The effect of reducing the unemployment benefit entitlement period during transition.” *Journal of Population Economics*, 13: 35–44.

Roed, K. and Zhang, T. (2003): “Does unemployment compensation affect unemployment duration?” *Economic Journal*, 113(1): 190–206.

Sallee, J. M. (2013): “Rational inattention and energy efficiency.” *NBER Working Paper 19545*, National Bureau of Economic Research, Inc.

Schmieder, J. F., von Wachter, T., and Bender, S. (2012a): “The effects of extended unemployment insurance over the business cycle: Evidence from regression discontinuity estimates over 20 years.” *The Quarterly Journal of Economics*, 127(2): 701–752.

Schmieder, J. F., von Wachter, T., and Bender, S. (2012b): “The effects of extended unemployment insurance over the business cycle: Evidence from regression discontinuity estimates over twenty years.” *The Quarterly Journal of Economics*, 127(2): 701–752.

Schmieder, J. F., von Wachter, T., and Bender, S. (2012c): "The long-term effects of UI extensions on employment." *American Economic Review*, 102(3): 514–19.

Schmieder, J. F., Wachter, T. v., and Bender, S. (2013): "The causal effect of unemployment duration on wages: Evidence from unemployment insurance extensions." *NBER Working Paper 19772*, National Bureau of Economic Research, Inc.

Schultz, W. P. (1999): "Changing behavior with normative feedback interventions: A field experiment on curbside recycling." *Basic and Applied Social Psychology*, 21(1): 25–36.

Shimer, R. and Werning, I. (2006): "On the optimal timing of benefits with heterogeneous workers and human capital depreciation." *NBER Working Papers 12230*, National Bureau of Economic Research, Inc.

Shimer, R. and Werning, I. (2007): "Reservation wages and unemployment insurance." *The Quarterly Journal of Economics*, 122(3): 1145–1185.

Tatsiramos, K. and Van Ours, J. C. (2014): "Labor market effects of unemployment insurance design." *Journal of Economic Surveys*, 28(2).

Tiefenbeck, V., Goette, L., Degen, K., Lalive, R., Fleisch, E., and Staake, T. (2014): "Salience, not social Pressure: Engaging in resource conservation might be more pleasant than we think." *mimeo*, ETH Zurich.

van den Berg, G. J. (1990a): "Nonstationarity in job search theory." *Review of Economic Studies*, 57(2): 255–77.

van den Berg, G. J. (1990b): "Nonstationarity in job search theory." *Review of Economic Studies*, 57(2): 255–77.

van Ours, J. C. and Vodopivec, M. (2006a): "How shortening the potential duration of unemployment benefits affects the duration of unemployment: Evidence from a natural experiment." *Journal of Labor Economics*, 24(2): 351–350.

van Ours, J. C. and Vodopivec, M. (2006b): "How shortening the potential duration of unemployment benefits affects the duration of unemployment: Evidence from a natural experiment." *Journal of Labor Economics*, 24(2): 351–350.

van Ours, J. C. and Vodopivec, M. (2008): "Does reducing unemployment insurance generosity reduce job match quality?" *Journal of Public Economics*, 92(3-4): 684–695.

Winter-Ebmer, R. (1998): "Potential unemployment benefit duration and spell length: Lessons from a quasi-experiment in Austria." *Oxford Bulletin of Economics and Statistics*, 60(1): 33–45.