Helsinki GSE

HELSINKI GSE DISCUSSION PAPERS 18 · 2023

Changing the unemployment insurance duration: heterogeneous effects and an unbudging exit spike

Heikki Korpela

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

HANKEN

Aalto University

Helsinki GSE Discussion Papers

Helsinki GSE Discussion Papers 18 · 2023

Heikki Korpela:
Changing the unemployment insurance duration: heterogeneous effects and an unbudging exit spike

ISBN 978-952-7543-17-7 (PDF)
ISSN 2954-1492

Helsinki GSE Discussion Papers:
https://www.helsinkigse.fi/discussion-papers

Helsinki Graduate School of Economics
PO BOX 21210
FI-00076 AALTO
FINLAND

Helsinki, October 2023

Changing the unemployment insurance duration: heterogeneous effects and an unbudging exit spike

Heikki Korpela*

October 19, 2023

Abstract

I study the effects of the maximum duration of unemployment insurance benefits on unemployment. I find highly heterogeneous effects across two consecutive reforms in Finland. Both reforms cut the maximum duration by 20 weeks. The first reform targeted those with short work histories and did not affect their mean time in unemployment. The second reform applied to everyone and reduced the mean duration by three weeks. The first reform also did not move the spike in job-finding rates: it stayed at the old maximum duration. The second reform moved the spike to the new exhaustion point, also among the short-history group. This difference can be explained by the first reform's unique implementation, which separated the time when insurance ends from when individuals must switch benefits agencies. This switch causes observable frictions, which appear to be an essential driver for the observed spike. I present how the job-finding rates can be adjusted for observed frictions. Adjusting the job-finding rates for observed frictions substantially flattens the spike.

Key words: unemployment insurance, unemployment duration, hazard rates

JEL classes: J64, J65

*University of Helsinki. I thank Tomi Kyyrä, Hanna Pesola and Roope Uusitalo for their help with access to data and numerous helpful suggestions, and Mika Haapanen for his comments. I also gratefully acknowledge financial support from the Yrjö Jahnsson Foundation. Declarations of interest: none. Email: heikki.korpela@helsinki.fi. The online appendix is available at http://iki.fi/heikki.korpela/research/.

1 Introduction

A key question in unemployment insurance (UI) is its maximum duration. Longer entitlements provide more protection against income loss but weaken the incentives to search for and accept jobs. Reviews of the literature by Tatsiramos and Van Ours (2014) and Schmieder and von Wachter (2016) indicate that a longer maximum duration also translates to longer spells of unemployment. Both these reviews also note that there is often a sharp spike in job-finding when the UI entitlement period is exhausted. Less is still known about which groups are the most likely to respond to changes in the entitlement period. This paper uses high-quality register data from Finland to show that the unemployed with short work histories react weakly to duration cuts.

The paper also documents a case where the exit spike does not move when the entitlement is cut and another where it does. The crucial difference between the cases appears to be whether individuals must switch benefits agencies when UI benefits run out. The switch creates frictions, which may cause observed exits to bunch up tightly at the transition time, even if actual job findings are more dispersed. This bunching can be unpacked when high-frequency data is available, providing a more accurate estimate of job-finding rates over unemployment.

The conclusions are drawn by comparing two consecutive reforms in Finland. The first, *targeted reform* in 2014 reduced new entitlements from 100 weeks to 80. The cut applied only to those with less than three years of work history (corresponding to about one tenth of all UI recipients). This reform did not affect the mean duration of unemployment and did not move the exit spike from its old place. Both results are surprising and generally at odds with prior research. The second, *universal reform* in 2017 cut new entitlements by 20 weeks for almost everyone, providing a benchmark case for comparisons. This later reform both reduced the mean duration by 3 weeks and shifted the exit spike, aligning well with existing literature.

The effects of both reforms are investigated using difference-in-differences setups. As the targeted reform only affected those with short work histories, these constitute a treatment group. The difference in unemployment durations between this group and those with slightly longer histories had been stable for at least 15 years before the reform. Those with longer histories thus offer a natural control group. If the reform had a clear impact on unemployment, the difference between the groups should have changed after the reform; it did not. A different but comparable setup is needed for the universal reform, which cut entitlements for everyone simultaneously. As the universal reform primarily serves as a benchmark and its findings are independently less novel, describing the setup is postponed until Subsection 3.2.

The lessons from the two reforms offer two contributions to the literature. The first contribution is that those with short work histories appear to be less responsive to changes in the entitlement. The second contribution is that the spike in exits from unemployment at the end of the UI entitlement period may be largely driven by administrative frictions rather than increased job search.

Both findings are based on the lack of responses to the targeted reform. However, these two non-responses appear to be caused by at least partially distinct drivers: the unique administrative implementation of the reform explains the unbudging spike, while the group being targeted is probably an important explanation for the unchanged mean duration.

There are two major justifications for making this distinction. First, after the universal reform, the exit spike visibly moves for the short-history group. Thus, the exit spike for this group is not generally immune to changes. Second, only about 20% of the the mean reduction caused the universal cut among other unemployed can be attributed to the movement of the spike. Therefore, the unbudging spike alone cannot explain why the short-history group did not respond to the earlier reform at other times during unemployment. Additionally, after the universal reform, mean durations changed less among the short-history group.[1] Even though the later reform moved the spike because it did not share the implementation of the first reform, this appears to have been insufficient to substantially change the mean duration for the short-history group.

Weak overall responses to entitlement cuts by unemployed with short work histories have been documented at least once before. Van Ours and Vodopivec (2006) documented a similar result for Slovenia. However, in that case, the authors attributed their finding to the short duration of the remaining entitlement, which would have left too little time to find jobs. In the present context, the maximum remained at 80 weeks even after the cut, so insufficient time to react cannot explain the non-response. Instead, those with short work history appear to have a weaker behavioural response to duration cuts. Since those with less experience are usually young, the results are also consistent with work by de Groot and van der Klaauw (2019), who find that young workers are less responsive to entitlement cuts.

In contrast, that the spike remains at the *old* exhaustion time when the entitlement is changed seems to be a unique finding. In prior literature, when a surge in exits at the entitlement is observed, it also appears to move consistently when the maximum duration is changed. Schmieder, von Wachter and Bender (2016), Lalive, Van Ours and Zweimüller (2006), Moffitt (1985) and Van Ours and Vodopivec (2006) document such responses of the spike.

The non-response of the exit spike can be easily visualised with minimal assumptions. Figure 1 shows how observed mean benefits develop during unemployment. Before the reform, they fell sharply at 100 weeks; after the reform, at 80 weeks. Figure 2 shows the

[1]These changes had to be examined with two different approaches from the main estimates, and the results are not fully comparable.

exit rates from unemployment. The clear spike in exits remains at 100 weeks before and after the reform. Following Kyyrä, Pesola, and Verho (2019), unemployment is measured as the time for which the person claims unemployment benefits (either UI or a follow-up flat-rate assistance) from data covering each benefit payment individually.

Card, Chetty, and Weber (2007) note that the spike is usually taken as direct evidence that unemployment insurance distorts incentives to find jobs. However, if the fall in benefits drives the spike, then changing the timing of this drop should also have moved the spike. With the targeted reform, this was not the case.

The key to the puzzle appears to be frictions when individuals have to switch benefits agencies. When insurance is exhausted, the individual must usually apply for further, smaller unemployment assistance from another agency. The targeted reform had a quirk: the insurance agencies continued to pay benefits up to the old maximum duration, but for the last 20 weeks, they only paid the flat-rate assistance. This was the only time in Finland when the time of UI expiration and the time of switching agencies were separated. The exit spike stayed at the time of the switch.

The universal reform did not share this quirk. For new entitlements since 2017, both the time when benefits drop and the time of switching were shifted by 20 weeks. This time, the exit spike moved to the new location.

Section 5 documents observable aspects of the frictions caused by the agency shift. Such frictions plausibly explain why the mean time to the next job after an exit also spikes at the maximum duration. A data-driven adjustment to the job-finding time can be applied to assess the potential magnitude of this effect, accounting for these short delays for each spell individually. Previously, authors such as Card, Chetty, and Weber (2007) and Kyyrä, Pesola, and Verho (2019) have shown that different time measures in unemployment may cause the spike to be severely over- or underestimated. The best measure will depend on the institutional context. The proposed adjustment builds on and expands the insights from this prior work.

Applying the proposed new adjustment to Finnish data substantially flattens the job-finding spike. Finland is not the only country where there is a switch between agencies when UI expires. Using the same adjustment on data from other countries could thus reveal whether these frictions might also drive the observed spike elsewhere.

The rest of the paper is organised as follows. Section 2 describes the institutional context and the reforms. Section 3 covers the definitions used, the data, and the groups for the difference-in-differences designs. Section 4 presents descriptive statistics. Section 5 covers non-parametric findings on the exit rates, the exit spike, and evidence on frictions around the spike. Section 6 presents the primary econometric models and the main estimates for the two reforms, while Section 7 considers sensitivity. Section 8 concludes.

Figure 1: Observed mean benefits as a function of benefit weeks.

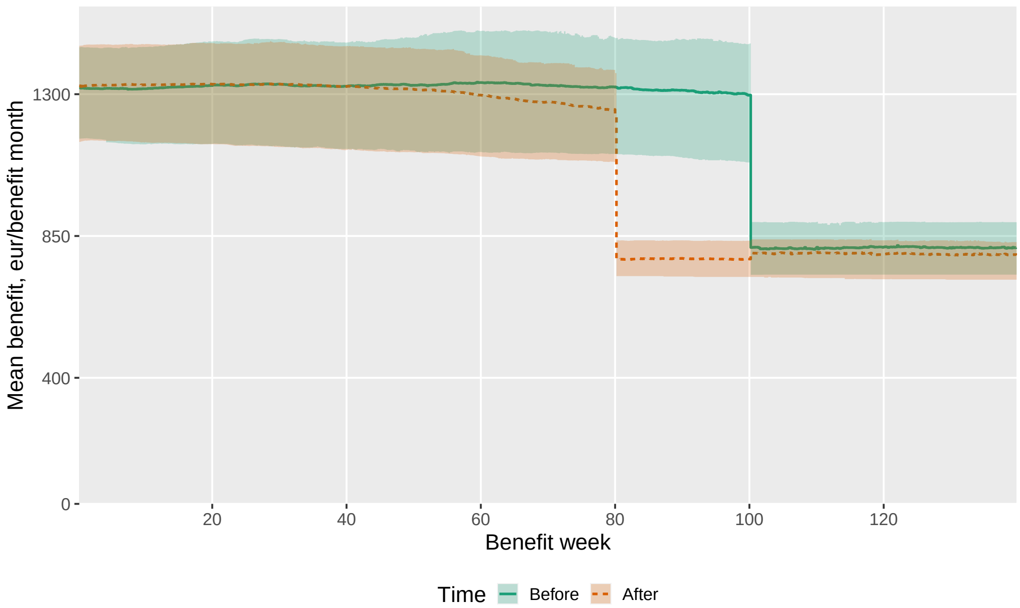

New UI entitlements with less than three years of work history in 2012–2013 (before the reform) and 2014–2015 (after). The shaded areas are the interquartile range.

Figure 2: Exit rate from unemployment as a function of benefit weeks.

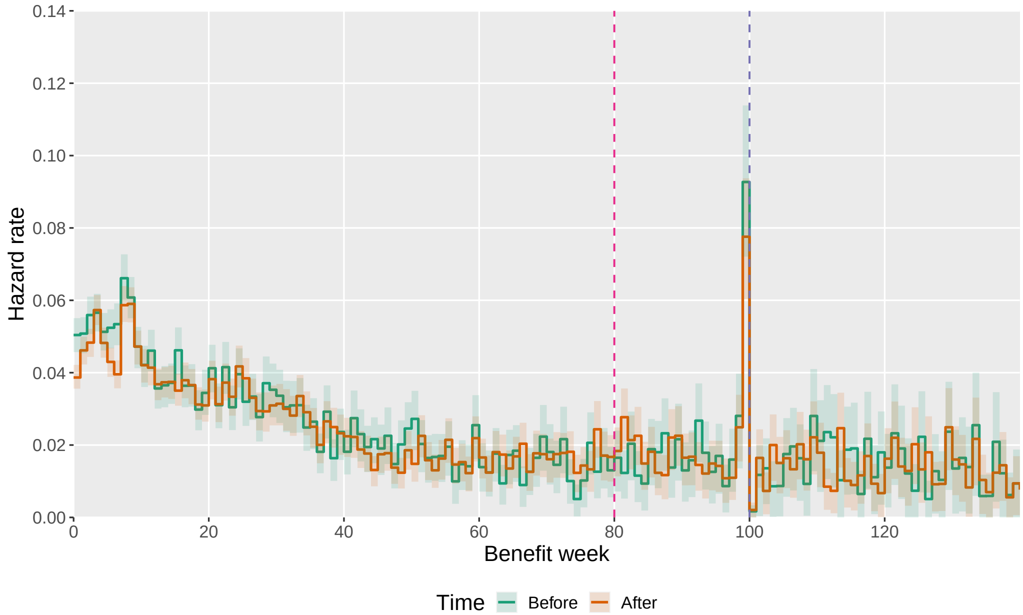

Shaded areas correspond to a bootstrapped 95% confidence interval. The vertical dashed lines correspond to the new and the old insurance entitlement.

2 Institutional setting

Finland has a two-tiered system of unemployment-based benefits. Most newly unemployed job seekers start on unemployment insurance (UI), which is based on prior earnings and has a limited duration (the entitlement). Persons who are not eligible for UI or exhaust the entitlement may apply for a lower, flat-rate unemployment assistance (UA) with no maximum duration. [2] Both benefits are conditional to being a registered unemployed jobseeker and are referred to as unemployment benefits.

Applicants must have sufficient recent employment (employment condition hereafter) and membership in an unemployment fund to claim the insurance. Up to 2013, the insurance had a default maximum entitlement of 100 benefit weeks.[3] Unemployment funds administer the insurance, while the public social security institution, Kela, pays the assistance.

Insurance is based on prior wages. In 2013, the mean observed payment was 361 euros per benefit week for UI and 162 euros for UA. If a person is part-time employed but looking for a full-time job, they may apply for part-time benefits, adjusted for the simultaneous wage. In such cases, benefit weeks are translated to full-time equivalents (FTE) by the benefits agency.[4] Part-time benefits consume the entitlement at the FTE rate. To focus on the connection between exits and the remaining entitlement, this paper uses these FTE weeks as the measure of time in unemployment. The mean ratio of benefit weeks to calendar weeks is around 0.98, and does not change much throughout spells.

The maximum duration is based on the first day for which benefits are claimed. A new entitlement can be earned by satisfying a new employment condition. Changes to the maximum duration only apply to new entitlements.

In force from 2014, the targeted reform cut new effective entitlements of *short-history* entrants (less than three years of work history) from 100 to 80 weeks. The universal reform, from 2017, cut new entitlements for almost everyone, again by 20 weeks. New entitlements dropped from 80 to 60 weeks for short-history individuals and for everyone else aged 57 or lower from 100 to 80 weeks. Work history is checked when the entitlement is determined.[5]

The targeted reform was implemented uniquely. Usually, after the end of entitlement, payments from funds cease, and one must apply for the assistance from the governmental agency. For the targeted reform, short-history individuals were still paid by the funds for

[2]Technically, there are two types of flat-rate transfers, which pay the same rate and only differ slightly in the rules. The details are discussed in Appendix A.

[3]Several special policies apply to those aged 57 by the end of the default entitlement, either extending the entitlement indefinitely or guaranteeing a job placement.

[4]For example, if a person would be entitled to 400 euros per week of full-time unemployment and receives 200 euros adjusted for a part-time wage, FTE weeks is $\frac{200}{400} = .5$.

[5]This measure counts any prior employment regardless of hours and wages, and is distinct from the employment condition, which tracks only recent employment.

weeks 81–100, but at the flat-rate levels of the unemployment assistance. After this, they had to reapply for any follow-up benefits from Kela as usual. The effect on unemployment benefit amounts was the same as a cut to the maximum entitlement period; only the name of the benefit and the payment institution differed. With the universal reform, no such peculiar arrangements were made. For any new entitlements since 2017, after exhausting UI, one must reapply for the assistance from another agency as usual.[6]

The funds inform the individuals about their entitlement with the initial benefit payment. They also tell recipients when the entitlement is exhausted and that they may now be eligible for flat-rate assistance from Kela. For the targeted reform, funds notified beneficiaries twice about expiration: once when their earnings-related UI benefits ran out and again when funds stopped the flat-rate payments.

The insurance and the assistance have shared eligibility criteria: the person must be registered as a jobseeker at public employment service (PES) offices, look for jobs, and accept job offers. The benefits may be increased or decreased for various reasons, such as participation in active labour market programs (ALMPs). Empirical benefit sums presented in the paper include all such adjustments.

A particular feature of the Finnish system is that roughly a third of UI spells are furloughs. An employer in temporary financial trouble may furlough an employee, usually for a fixed period. Furloughs overwhelmingly end in recalls and only very rarely turn into longer-term unemployment.[7] Furlough spell durations also did not empirically change in either reform.

Many unemployment-related policies besides the entitlement also changed in the 2010s. In general, most policy changes are unlikely to have affected either reform's control and treatment groups (to be defined shortly) differently; the major ones are listed in Appendix B. Two changes are more notable regarding the setting here.

First, the employment condition was relaxed in 2014, simultaneously with the targeted reform. Until 2013, applicants needed 34 weeks of qualifying employment over a review period (typically, the preceding 28 months) to be eligible for insurance. From 2014 to 2019, 26 weeks sufficed. The change caused many jobseekers who would not have qualified for insurance in 2013 to qualify for it in 2014. The relative increase was larger for those with shorter overall histories, which caused the composition of the treatment and control groups to change differently. In Subsection 7.1, weights are presented that balance observed recent employment to pre-reform levels; the empirical results are not highly sensitive to these weights.

[6]The targeted reform was also phased out in a peculiar fashion. From 2019, for short-history individuals who were still using their entitlements from 2014–2016, weeks 81–100 were again paid at the normal earnings-related levels. These cases are extremely rare, and unemployment follow-ups for the targeted reform end well before 2019.

[7]79% of UI spells starting on furloughs end within three months, and 92% end within six, and the durations are not sensitive to the business cycle.

Second, in 2014, a 4-week initial benefit increase was phased out. This increase was previously available to those with at least three years of work history (i.e., only the control group in the control period). Subsection 7.5 discusses this change. Various strands of evidence suggest this phaseout probably did not change mean time in unemployment much, as the period on increased benefits was very short.

3 Data and definitions

3.1 Data

The study combines individual-level, high-frequency administrative records on unemployment benefits, registered unemployment, employment, and wages.

Each payment of unemployment benefits over 2010–2021[8] was observed for all those collecting these benefits. These data include exact dates of unemployment which the benefits are claimed for, benefit amounts, any increases and decreases, payment dates, and benefit day counters towards the maximum entitlement.

Records on employment until 2020 include start and end dates of job contracts and annual wages per contract. The data are based on mandatory pension insurance. This mandate has progressively expanded in Finland over the decades to be almost universal and has very high coverage, particularly in recent decades. A companion dataset covers certain non-employment types which extend the review period for the recent employment condition of UI eligibility.

Finally, detailed data were available on registered unemployment, benefit eligibility statements, professions, education, places of residence and birthdates. The benefit data come from the Financial Supervisory Authority (that collects it from the funds) and Kela. Individual background data were obtained from the Ministry of Economic Affairs and Employment, which supervises the PES offices. Data on employment, wages and periods of non-employment were provided by the Finnish Centre for Pensions (ETK).

Some additional datasets were used for the appendices, covered in Appendix Y.

3.2 Definitions and sample restrictions

An unemployment spell is defined as a period for which a person receives either UI or the flat-rate assistance (UA). Its length is measured in FTE benefit weeks, described in Section 2. Since unemployment benefits can be collected during participation in ALMPs or part-time work, the individual may thus alternate between such statuses during a spell as long as they continue to look for a full-time job and collect benefits. A spell ends when there are more than 30 days for which no UI or UA is claimed.

[8] UI was also observed for 1999-2009, but not the flat-rate assistance.

Some restrictions were used across all samples used for estimation and descriptive statistics unless otherwise noted. The largest exclusion comes from dropping spells starting on furloughs: these short spells are unlikely to be affected by the entitlement cut for reasons discussed in Section 2. Persons aged 55 or older when entering unemployment are also dropped, as they have age-specific provisions at the end of their entitlement. Voluntary quits (1–2% of otherwise eligible spells) within six months of a spell's start are excluded.

Finally, to keep the setup tractable, the samples are limited to fresh UI entitlements (15–17% dropped) and those who earn a new entitlement during the spell through part-time employment are dropped (2–4% dropped). Appendices N and O discuss findings when the latter constraints are relaxed.

All unemployment spells are followed for a maximum of two years and ten months from their starting point. This duration captures most exits from unemployment while ensuring that there is no overlap with the COVID-19 pandemic period. A spell continuing past this time is censored; otherwise, an exit occurs at the end of the last observed benefit period.

If an exit is followed by employment, it is classified as job-finding.[9] Additionally, wages and duration in the next job were followed for a maximum of one year after unemployment. [10]

For spells since 2014, entitlement at the start of a spell must be determined from observed prior employment. While accurate register data is available for the entitlement used, it counts weeks elapsed rather than weeks remaining. Thus, actual entitlement can only be directly observed for the highly selective group who survived in unemployment beyond 80 weeks.

Estimating entitlement through employment introduces some measurement error.[11] The observable error is significantly larger for those with less observed histories, and is at its largest at the threshold of three years' history. To reduce the error, the treatment group for the targeted reform is defined as having 1–2.5 years of work history, and the control group as having 3.5–5 years. The weights described in the sensitivity analysis in 7.1 further reduce the observable error to 6.5%. Since the universal reform affected almost everyone, the measurement error for the maximum entitlement can be minimised by limiting the sample to those with 3.5–19 years of history for that reform. Thus, for both reform samples, post-reform treatment spells have an estimated entitlement of 80 weeks, and all others have an entitlement of 100 weeks.

[9] The job must also satisfy a low wage threshold, be ongoing within 30 days after the exit and last at least a month. Appendix D illustrates the effects of these constraints on the job-finding estimates.

[10] For the universal reform, the joint unemployment and employment follow-up would overlap with the COVID-19 pandemic period. For this reason, the empirical results for the later reform are presented separately over two different follow-ups for comparisons.

[11] Appendix W covers the observed error in more detail, using spells continuing past 80 weeks. There appears to be no error related to the entitlement start date, even at the exact turn-of-year thresholds.

For the targeted reform, entitlements starting in 2012–2013 constitute the pre-period, and those beginning in 2014–2015 the reform period. These windows strike a reasonable balance between power, availability of a placebo test given the data, and a sufficiently long follow-up available.

As the universal reform targeted everyone, a different setup must be chosen. Since this reform mainly serves as a benchmark for the targeted reform, a similar difference-in-differences strategy is desirable. This time, the identifying variation comes from the entitlement changing sharply at the year-turn of 2017. Entitlements starting in January 2017 constitute the post-reform treatment group, while those beginning in December 2016 form the control group. Over such a short interval, changes in the macroeconomic environment are unlikely to be driving time in unemployment. However, those entering unemployment in December or January are selected differently and face different immediate job opportunities. Thus, a similarly formed sample from the turn of 2016 is used as a control period to account for differences in the selection by month.

The sample for both reforms is limited to persons aged 18–54 years when entering unemployment. Older cohorts are eligible for age-specific policies that either extend their insurance entitlement indefinitely or guarantee them a job placement at the end of the entitlement.

After the exclusions, 50 330 spells were available for the targeted reform and 18 996 for the universal reform.

There are two remaining significant concerns: leftover error in measuring the actual entitlement (first reform) and specific changes in the inflow to unemployment over time. For additional robustness, the sample was further balanced using entropy balancing weights, which address these potential issues. The main estimates proved not to be highly sensitive to these weights, so further details on the motivation and implementation of the weights are postponed until Subsection 7.1. However, the main estimates in Section 6 are directly presented with and without the weights to ease comparisons.

4 Descriptive statistics

Table 1 describes the sample for the targeted reform. The variables include characteristics fixed at the start of each spell and potential outcomes, such as spell duration and ALMP participation during unemployment. Appendix Z contains definitions and a similar table for the universal reform.

The selection strikes a reasonable compromise between group similarity, classification error and sample size. While the groups are not identical, observed heterogeneity is small compared to the entire unemployed population.

The observed differences remain quite stable over time for most characteristics fixed at the start of a spell. Running a difference-in-differences test using a covariate as a potential outcome shows no significant change for most such variables. The benefit level at the beginning of spells is one major exception, analysed in Subsection 7.5. Most of the other differences over time appear to be related to the loosening of the recent employment condition in 2014, as balancing the observed recent employment also attenuates these differences. This change also helps explain the relative increase in the share of persons who have received UI before.[12]

[12] Among persons who re-enter unemployment each year, more individuals would be eligible for UI in 2014 with less than 1.5 years of experience and then again in 2015, now with between 1.5 and 2.5 years of experience.

Table 1: Means of various outcomes and characteristics for the targeted reform.

Variable	Treatment, before	Treatment, after	Control, before	Control, after
Duration in full-time equivalent benefit weeks	28.0	30.1	28.5	30.2
Duration in calendar days	207	222	211	223
The spell continues past the UI entitlement	6.81%	11.35%	8.01%	8.11%
The spell continues past the agency switch	6.81%	7.52%	8.01%	8.11%
Age	25.4	25.4	29.5	29.3
Woman	49.2%	50.4%	62.5%	59.6%
Received unemployment benefits since 2005	72.0%	75.5%	80.2%	81.1%
Received UI since 2005	12.2%	17.9%	46.1%	47.2%
Received unemployment assistance (UA) since 2005	68.9%	70.8%	67.0%	69.3%
UI spells in the last 1.9 years if any, N	1.5	1.5	1.5	1.6
UA spells in the last 1.9 years if any, N	1.4	1.4	1.3	1.3
Wage basis for benefits (indexed to 2019), euros/mo	2236	2184	2366	2350
Number of children	0.33	0.30	0.65	0.58
Benefit weeks in ALMPs	4.77	5.33	4.97	5.25
Benefit weeks in partial unemployment	2.24	3.33	2.77	3.64
Total unemp. benefit payments, euros	8567	8986	9261	9660
Initial payment, euros/month	1239	1239	1438	1307
Average payment, euros/month	1211	1222	1318	1291
Prior employment, years	1.80	1.79	4.20	4.19
Nationality other than Finnish / residence permit	12.81%	12.40%	9.99%	10.79%
Duration of last job, years	0.75	0.72	1.01	1.04
Time from previous employment to spell, days	33.5	25.7	42.4	33.5
Contribution weeks towards the recent employment condition	58.3	55.1	71.4	70.0
Estimated years of completed education	12.9	13.0	14.2	14.1
Inverse of regional labour market tightness	0.37	0.35	0.47	0.41
Postal code area inv. pop. density (pct of national weighted avg.)	65.2%	68.8%	75.4%	80.7%
Postal code area unemp. rate (pct points over national weighted avg.)	1.18%	1.13%	1.13%	1.10%
Re-enters unemployment after spell	53.7%	53.6%	53.7%	53.7%
Entry in June after fixed-term contact in specific professions	5.92%	5.32%	10.83%	9.08%
Days from exit to next spell (if any)	153	146	163	152
Average payment by Kela, euros/month*	789	770	798	785
Last payment paid by fund, euros/month**		754		767
Last earnings-related payment, euros/month**		1214		1226
First flat-rate payment paid by fund, euros/month**		747		767
Initial payment paid by Kela, euros/month**		768		780
N	8768	12957	12314	16291

* = among those who transfer to Kela after 100 benefit weeks

** = among those directly affected by the reform (flat-rate benefits paid by funds after 80 weeks)

5 Descriptive evidence of the exit spike

Figures 3–4 show the unqualified exit hazards and job-finding rates for the targeted reform.[13] The job-finding hazard is generally parallel to, but lower than, the indeterminate exit hazard. Appendix D shows how relaxing the job-finding criteria changes the rate.

The exit spike is essentially unchanged after the reform that targeted the short-history group. Comparing exits to observed benefits in Figure 5, there is a significant drop in benefits at the 80-week mark after the reform but no corresponding increase in the exit rate. Similarly, there is no clear change in benefits at the 100-week mark after the reform, but a large exit spike remains.

In contrast, Figure 6 shows that after the universal reform, the exit spike shifts. A similar shift, covered in Appendix K, can also be seen among those with short work histories after the universal reform.

Several hypotheses for why one of the reforms did not affect the spike can be rejected with high likelihood. These include the benefits agencies using different eligibility criteria, different rules during the targeted reform, and missing data. These non-explanations are covered in Appendix G.

[13]Spells and their ending were defined in Section 3.

Figure 3: Exit hazard from unemployment benefits, targeted reform.

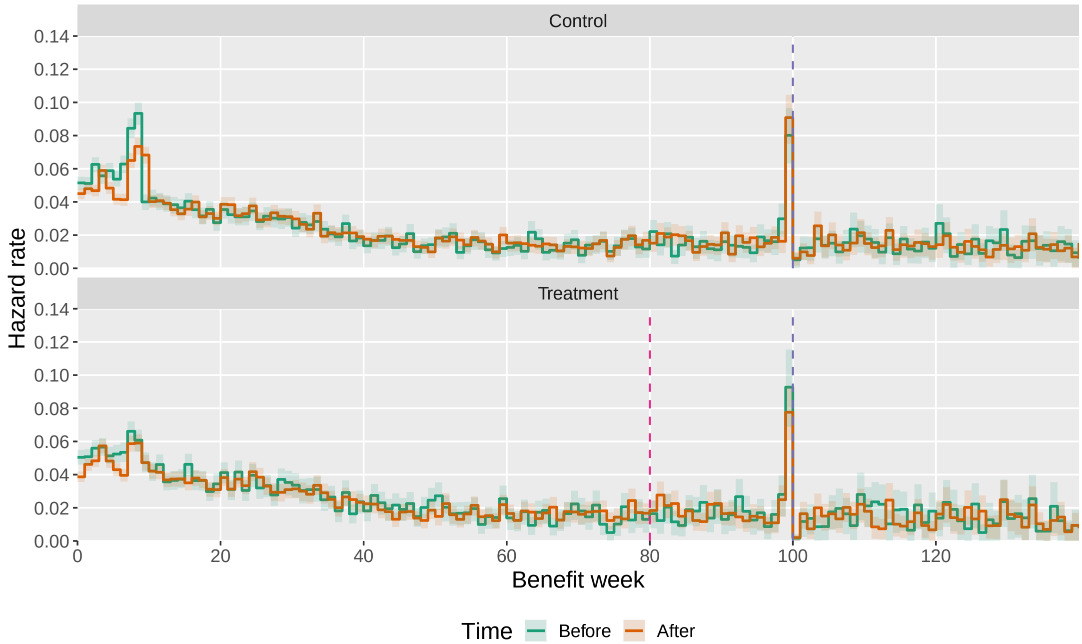

Non-parametric estimate, exits binned to weeks. The shaded areas correspond to a bootstrapped 95% confidence interval. The dashed vertical lines represent the entitlements.

Figure 4: Job-finding rate, targeted reform.

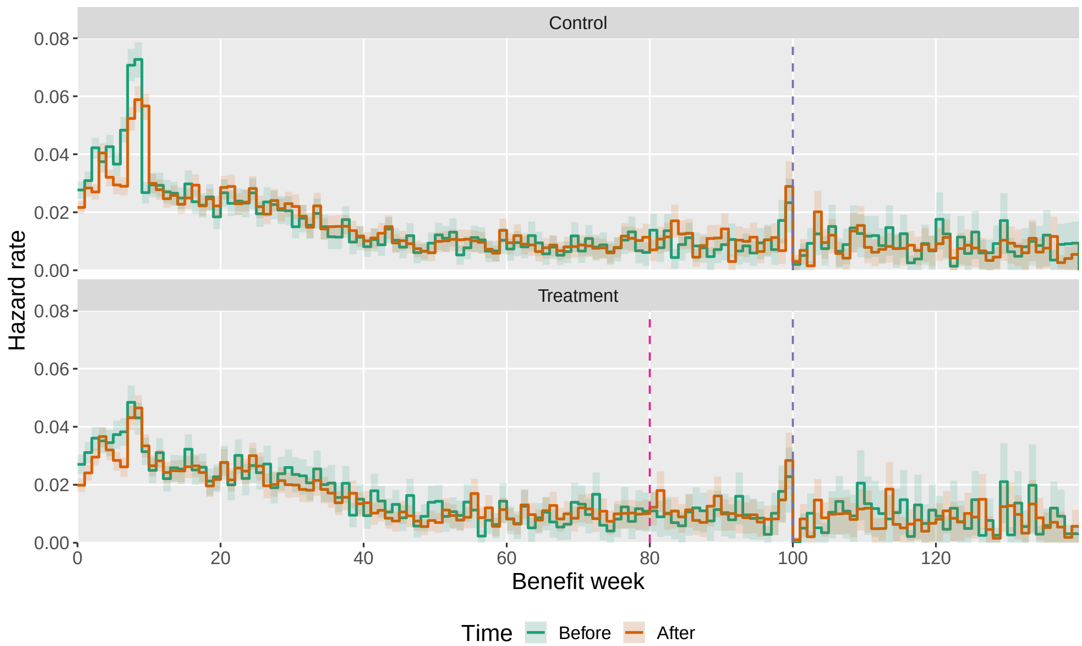

Non-parametric estimate, exits binned to weeks. The shaded areas correspond to a bootstrapped 95% confidence interval. Note that the vertical scale differs from the unrestricted exit hazard.

Figure 5: Mean monthly benefits as a function of benefit weeks, targeted reform.

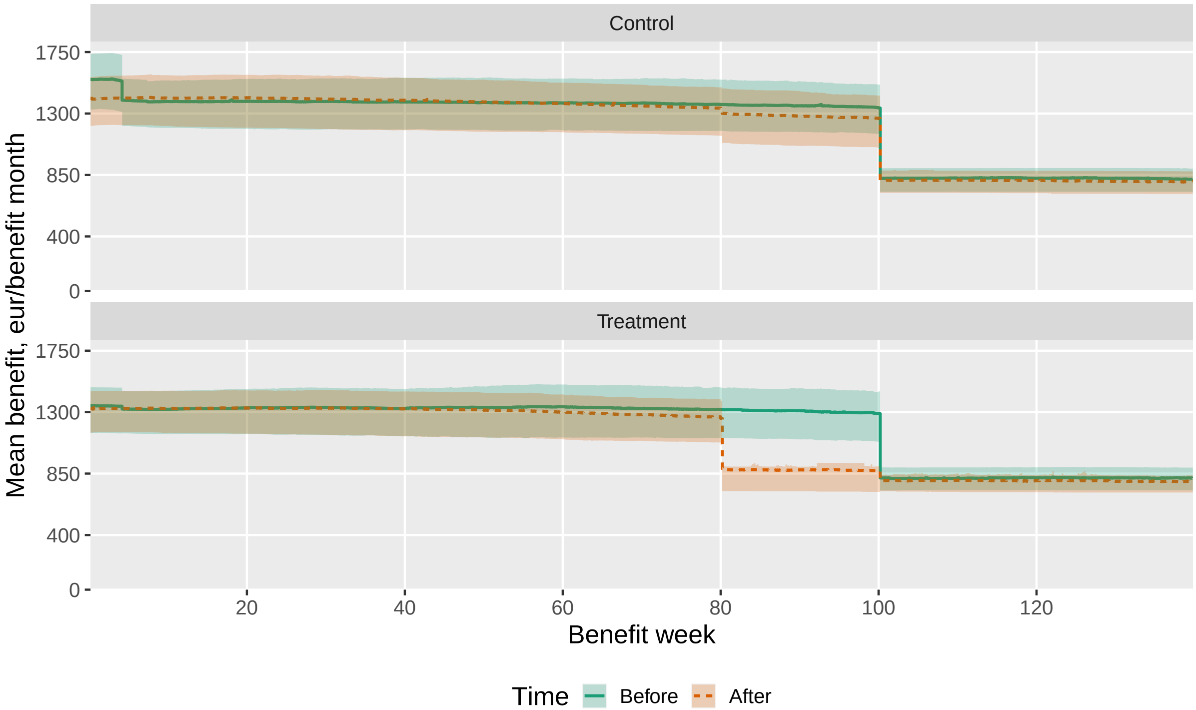

The shaded areas correspond to the interquartile range.

Figure 6: Exit hazard from unemployment benefits, universal reform.

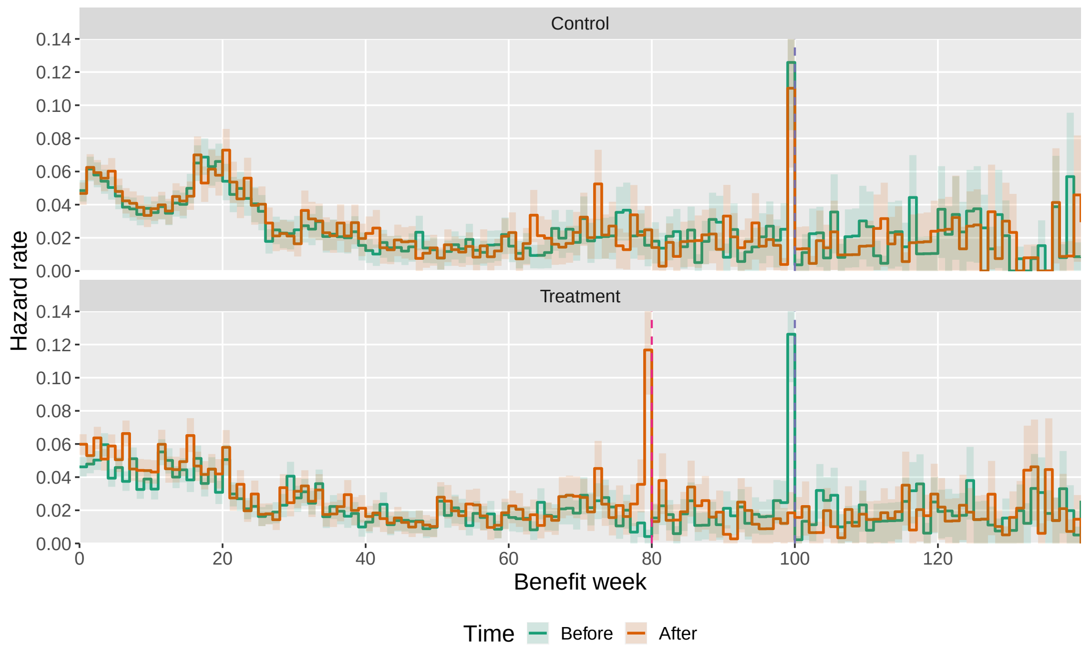

Figure 7: Job-finding rate, universal reform.

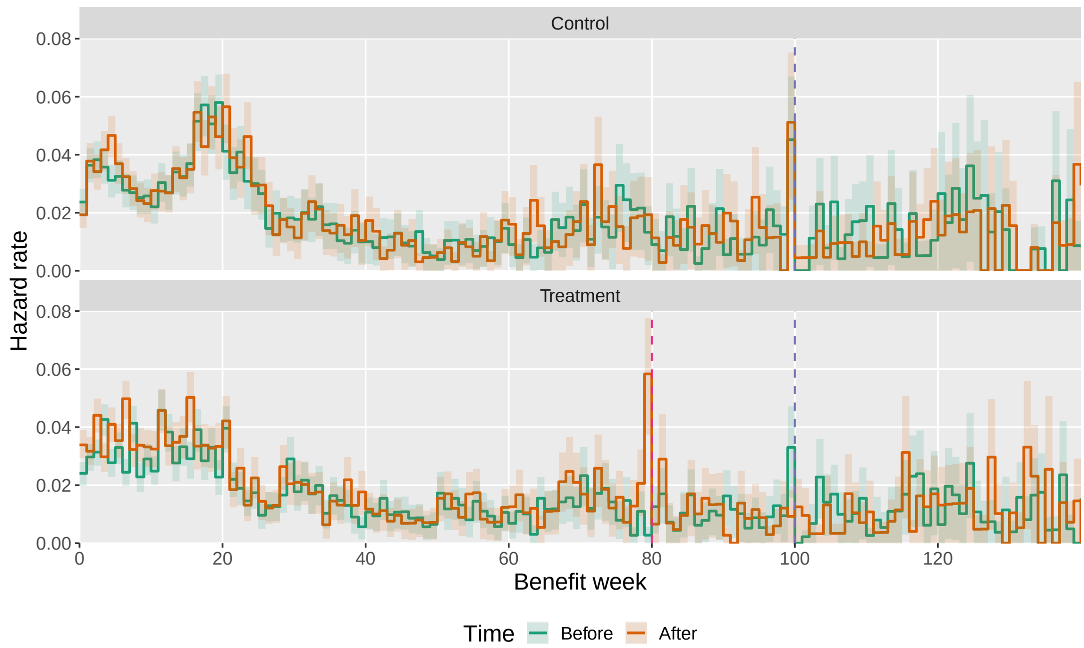

Note that the vertical scale is different from the unrestricted exit hazard.

A more plausible explanation for why the targeted reform did not move the spike is frictions related to a transition between benefits agencies. As explained in Section 2, for the targeted reform, this agency transition remained at the old UI expiration time, and it is also where the spike stayed. At other times, including after the universal reform, the agency transition coincides with UI expiration.

To understand the frictions, consider two unemployed individuals, A and B, who randomly receive a job offer for the same date and accept it. At this point, A only has one week of insurance left, while B has plenty of entitlement remaining. Both individuals now need to decide whether to file for the residual benefits. The two cases can be visualised as follows, with time on the horizontal axis:

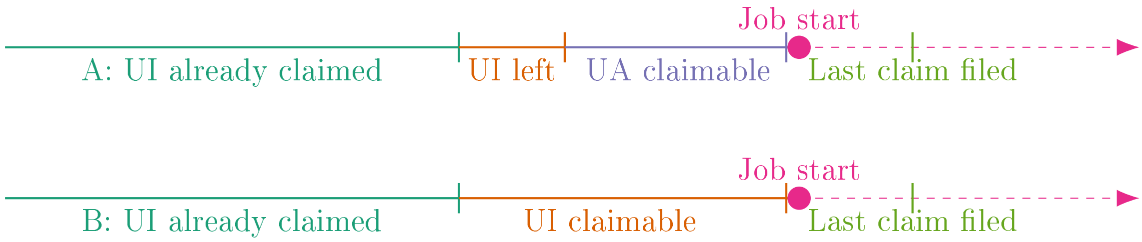

The unemployed usually claim benefits once per month for the time unemployed in the preceding month. If applicants have found a job by the time of filing, they may claim benefits up to the last jobless day. Meanwhile, UI entitlements are defined in weeks, not months. Thus, for the last month of the UI entitlement, most only have a partial month left. Individuals who continue in unemployment after this need to file one last partial-month claim for UI and a new partial-month application for the flat-rate assistance (UA). Figure 8 shows, for each benefit week, the distribution of calendar days *per claim*. For calculation details and additional figures, see Appendix C.[14]

At least three parameters govern each choice to file for benefits: the potential benefit, the probability of getting the benefit, and the disutility of time spent in the application process. For the last UI claim after finding a job, the benefits are high and time costs low: one only needs to fill out a short follow-up form. When switching to UA, one must go through a new application process, which may be perceived as more cumbersome and can, for example, require additional documentation. Additionally, some individuals may mistakenly believe they cannot retroactively apply for a new type of benefits after finding a job.

These factors might drive individuals who exit soon after UI expiration to file for the residual UI but forgo the short amount of the flat-rate assistance. If there are plenty of cases like this, the observed exits bunch up on the last day of insurance, while almost nobody seems to leave unemployment soon after.

Consistent with the example, Figure 9 shows that delays to jobs increase markedly

[14] The pattern shown is driven by the number of *calendar days one can claim at a time*, not partial unemployment. The mean ratio of benefit days to calendar days is much more stable over spells between 0.95 and 0.98.

on the last week of insurance.[15] Fortunately, having this data implies that the empirical magnitude of the potential phenomenon can be estimated.

To take the phenomenon into account, an adjustment to data is proposed. The adjustment is conservative, but still has a significant impact on the empirical job-finding rate. The baseline measure remains the same: full-time equivalent benefit weeks in unemployment to track time until UI exhaustion accurately.[16] Job-finding is still a binary indicator, subject to the same criteria as before. Only if there is a delay from an exit to the next job between 2 and 30 days is this delay converted to benefit days and the exit time adjusted. In terms of the earlier example, A's time in unemployment is extended by the short UA period they could have claimed but did not.

Figure 10 illustrates the job-finding hazard with and without this adjustment. With the adjustment, the spike visibly flattens, and the very low job-finding rate right after the maximum visibly shifts upwards. For similarly adjusted hazards for other groups, see Appendix C.

There are three caveats. First, high-frequency data was available for jobs but not for other destinations such as education or child home care. The generic exit hazard could thus not be adjusted similarly. Therefore, and because the available adjustment would only change the *mean* duration by about 0.3%, the parametric estimates (Section 6) were made using the unadjusted durations. Second, the spike in job-finding is lower to start with for the treatment reform sample. The same adjustment also substantially flattens the spike for other groups but does not remove it, suggesting there is also an actual behavioural response to UI exhaustion.

Third, compared to exits shortly before or after exhaustion, exits are much more likely to happen at the spike for persons with prior part-time or fixed-term disability pensions and those with entrepreneurial or property income. Similarly, exits at exhaustion are more likely to be followed by specific patterns in primary income sources. Overrepresented post-exit income groups include disability pensions, business or property income, and extremely low incomes. These patterns suggest that there are several factors driving exits to bunch at the spike; unfortunately, most of these observed patterns do not suggest that UI expiration is fastening transitions to employment. Appendix E covers these questions in more detail.

The spike has attracted considerable interest in the literature for at least two reasons. First, it has been taken as direct evidence that UI may distort the incentives to search for and accept jobs. Second, the spike is an empirical phenomenon that can be used to select and refine theoretical models, depending on whether the models' predictions are

[15] This figure covers all comparable exits between 2010 and 2016 when work history restrictions are dropped but other sample restrictions are maintained. Appendix C shows figures for the estimation samples used for the two reforms; they are qualitatively similar but noisier.

[16] Appendix I discusses why this measure outperforms time in non-employment or registered unemployment.

Figure 8: Calendar days per claim

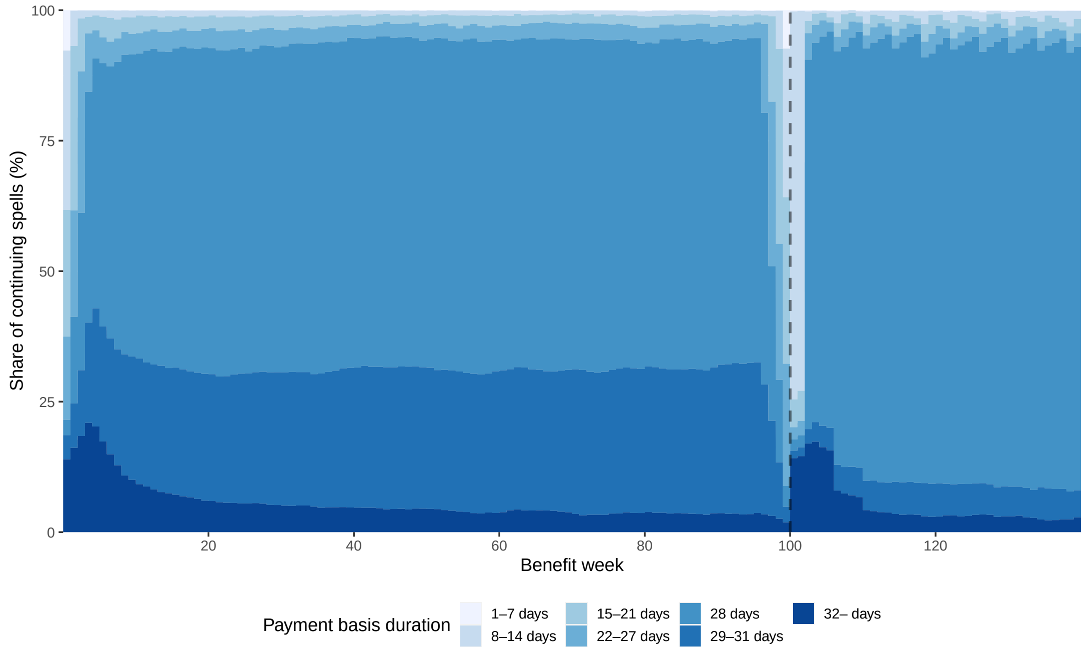

The figure tracks, for each benefit week of ongoing unemployment, how many calendar days of unemployment were being claimed per claim on average. For details and similar figures for other groups, see Appendix C.

Figure 9: Delay to job after exit

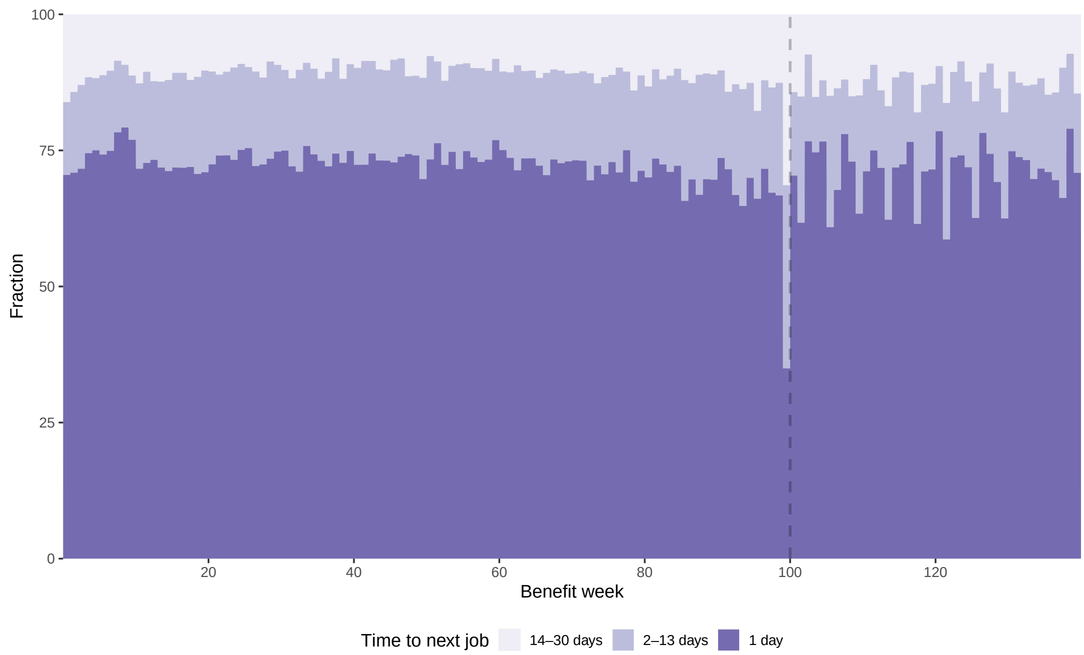

The figure shows the delay to jobs, conditional on the exit being classified as job-finding. For details and accompanying figures, see Appendix C.

Figure 10: Job-finding hazard, adjusted by delay to the next job.

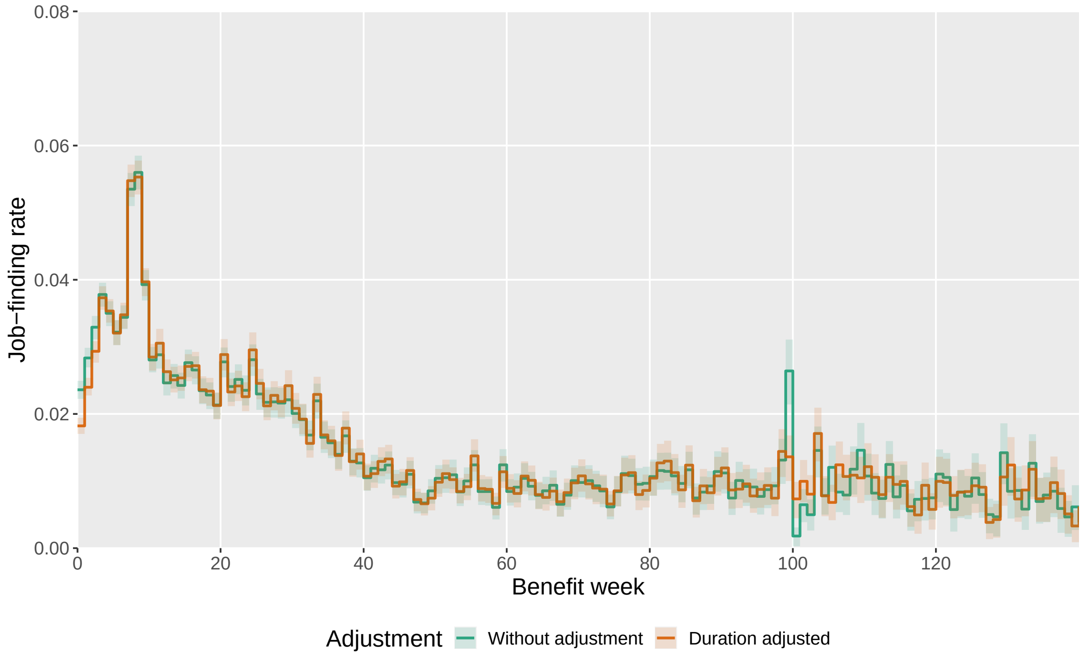

consistent with a spike.

In particular, the drop in exit rates following the spike is inconsistent with standard search models, such as the model by Mortensen (1977). A standard prediction is that exits should increase steadily towards the end of entitlement and then spike and stay elevated when UI expires. The adjustment presented here suggests that the actual job-finding response around exhaustion time in Finland is modest and dispersed over time. Further, while the unadjusted job-finding rates drop to very low levels right after exhaustion, the adjusted rates do not. Thus, the adjustment attenuates the puzzle to be explained.[17]

If the exit spike is much smaller after the adjustment, the spike naturally provides much weaker evidence for the incentive effects of UI. However, the role of the spike has always been that of supportive evidence; an exit spike is neither sufficient nor necessary for UI to have these effects. Empirical evidence that examines the average effects of UI on the duration of unemployment and employment spells addresses the true impact more comprehensively and directly. The following section presents such empirical evidence.

[17]Boone and Van Ours (2012) present a model with storable job offers, and this model is consistent with an observed spike. For this model to be consistent with the commonly observed delays to jobs in the Finnish data, individuals would have to store job offers, wait for UI to expire, forgo the follow-up UA and only cash in the stored offer after a delay.

6 Empirical results

6.1 Effects on average duration of unemployment

For both reforms, the main estimands were the impacts of the entitlement cut on the average duration of unemployment and the probability of re-employment. Among the re-employed, effects on the average duration and wage in the subsequent job were also examined. In each case, the assumption is that the difference between the applicable control and treatment groups would have stayed the same over time had the reform not taken place. When the weights (see Subsection 7.1) are used, the assumption is conditional on potential selection over time by the observables used in weighting. With weights, the estimated effects are also local to the reference group used in weighting.

The effects on mean outcomes are assumed to follow

$$D_i = \omega + \alpha \cdot After + \gamma \cdot Treatment + \delta \cdot (After \times Treatment) + x_i'\beta + \varepsilon_i, \qquad (1)$$

where x_i is a vector of controls fixed at the start of an unemployment spell, i indexes unemployment spells, $After$ indicates the relevant period, and $Treatment$ indicates the reform targeted the individual (based on their work history or time of entry around the turn of the year).

Under the parallel trends assumption, the average effect of the reform is identified by the coefficient δ, the interaction term reported in the result tables. For continuous variables, the effects were estimated with OLS; for re-employment probability, the logit counterpart was used. To obtain standard errors, the entire procedure was bootstrapped with 1 000 repetitions (including reweighing after each resampling). For covariates, mainly the same additional controls were used as for balancing. [18]

For the targeted reform, the overall finding is a relatively precise, near-zero effect. Table 2 collects these results from models with and without additional controls and with and without weights.[19] The results exhibit only minor sensitivity to including weights, additional controls or both.

In contrast, the universal reform reduced time in unemployment, with a point estimate of about −3 weeks, or 10% of the pre-reform mean. The main results are collected in Table

[18] In addition to the weighting variables, extra covariates were added for precision for the targeted reform. These were year/month dummies, whether unemployment was part-time at the start of the spell, number of children, and regional labour market tightness. As their inclusion in weighting would significantly reduce the effective sample size but had a negligible effect on estimates, they were not used for weighting. The point estimates are not overly sensitive to these controls.

The combined use of entropy balancing weights and controls follows Zhao and Percival (2016), who present simulation evidence that this approach may yield added robustness against misspecification in either the weighting or the regression.

[19] In this and other relevant tables, ESS is the effective sample size measuring loss of precision from weights, defined by Greifer and Stuart (2022) as $ESS = \frac{\left(\sum_{i=1}^{n} w_i\right)^2}{\sum_{i=1}^{n} w_i^2}$, where w_i are the weights.

3. The preferred specification with both weights and controls indicates a range in good agreement with previous research regarding Finnish UI duration by Kyyrä and Pesola (2020b) and consensus estimates in the literature cited in the introduction. No impacts on re-employment rates are observed within a month from exits; while exits happened earlier, an exit was, on average, no more or less likely to be towards jobs than earlier.

Because the universal reform happened later, a long follow-up for unemployment followed by another follow-up for the next job would in many cases overlap with the COVID-19 pandemic period. To avoid the overlap, a separate analysis is presented in table 4. In this analysis, unemployment spells were followed for a maximum of only 2 years and 7 weeks (instead of 2 years and 10 months), and the duration and wage in the next job for up to a year after each exit. The results imply that the later reform did not negatively affect the quality of the next job on average.

Table 2: Effects of the targeted reform on unemployment duration, re-employment and post-unemployment outcomes.

Outcome	Unweighted, without additional controls	Unweighted, with controls	Weighted, without controls	Weighted, with controls	Pre-reform mean	N	ESS
FTE weeks of unemployment	0.5329 (0.6576)	0.4000 (0.6491)	-0.2629 (0.6809)	-0.0867 (0.6706)	27	50,330	43,925
Re-employment probability	-0.0102 (0.0089)	-0.0105 (0.0089)	-0.0062 (0.0089)	-0.0089 (0.0090)	0.64	50,330	43,925
Duration in next job, days	3.4160 (2.3645)	0.9655 (2.3749)	4.0299 (2.5518)	2.7104 (2.5808)	193	36,692	32,116
Wage in next job	-8.6270 (13.7444)	-0.6370 (11.9442)	11.5425 (13.7319)	10.3508 (13.5781)	1715	36,692	32,116

Unemployment spells were followed for a maximum of 2 years and 10 months and post-unemployment outcomes for a subsequent maximum of 1 year. Effects for the re-employment probability are bootstrapped marginal effects. Effects on wage and duration of job are for the re-employed. Wages are monthly wages indexed to 2005. Pre-reform means are for the treatment group. ESS = effective sample size. Bootstrapped standard errors in parentheses.

Table 3: Effects of the universal reform on unemployment duration and re-employment.

Outcome	Unweighted, without additional controls	Unweighted, with controls	Weighted, without controls	Weighted, with controls	Pre-reform mean	N	ESS
FTE weeks of unemployment	-3.2900 (1.0359)	-2.9954 (0.9331)	-3.1315 (1.0218)	-3.0524 (1.0809)	31	19,393	17,832
Re-employment probability	0.0121 (0.0136)	0.0100 (0.0138)	0.0034 (0.0143)	0.0039 (0.0144)	0.66	19,393	17,832

Unemployment spells were followed for a maximum of 2 years and 10 months. Post-unemployment outcomes were not followed to avoid overlaps with the COVID-19 pandemic period. Other remarks as in table 2.

Table 4: Effects of the universal reform, using a different follow-up.

Outcome	Unweighted, without controls	Unweighted, with controls	Weighted, without controls	Weighted, with controls	Pre-reform mean	N	ESS
FTE weeks of unemployment	-2.8731 (0.9012)	-2.6169 (0.8764)	-2.8595 (0.9314)	-2.7927 (0.8910)	30	19,393	17,832
Re-employment probability	0.0145 (0.0137)	0.0123 (0.0143)	0.0079 (0.0146)	0.0082 (0.0143)	0.65	19,393	17,832
Duration in next job, days	3.0818 (3.7425)	4.6033 (3.6495)	3.3981 (3.8724)	4.4731 (3.9270)	209	14,981	14,012
Wage in next job	11.4381 (25.8846)	20.4269 (24.0000)	-10.4114 (25.3968)	3.7514 (24.4420)	1957	14,981	14,012

Unemployment spells were followed for a maximum of 2 years and 7 weeks and post-unemployment outcomes for a subsequent maximum of 1 year. Other remarks as in table 2.

6.2 Effects on the time profile of unemployment

To estimate the time profile of unemployment, the empirical hazard function for exit from unemployment is modelled as

$$\theta(i,t) = \lambda(I(t), Treatment, After)\exp(x_i\beta), \tag{2}$$

where t measures elapsed FTE benefit weeks[20] from the start of an unemployment spell, i indexes spells, I maps benefit weeks to interval indices, λ is a function of these that may vary over the spell, and x_i are controls fixed at the start of a spell. As before, $Treatment$ and $After$ indicate the relevant treatment group and period at the start of a spell. As implied by the specification, the assumption is that the additional controls may shift the level but not the shape of the hazard. The same controls and weights were used as for the mean outcomes in the previous subsection.

The baseline hazard λ is allowed to vary by group and time in unemployment, and is modelled as a piecewise constant:

$$\begin{aligned}\lambda(I, Treatment, After) = \exp\big\{&\gamma_{I,0} + \gamma_{I,1}Treatment \\ &+ \gamma_{I,2}After + \gamma_{I,3}(Treatment \times After)\big\}.\end{aligned}$$

In this setting, the coefficients $\gamma_{\cdot,3}$ represent the (relative) causal effect of the reform on the hazards at distinct stages of unemployment.[21]

The model results are illustrated in Figures 11 and 12. In both cases, the hazard estimates for a typical individual in the post-reform treatment group, using medians and modes for covariate values. The treatment effect is calculated as the difference between this hazard and a counterfactual hazard, where the $After \times Treatment$ interaction term is set to 0. The plotted confidence area corresponds to the 95% bootstrapped confidence interval for this difference. Appendix Q lists the hazard estimates by interval.

The targeted reform generally had little effect on exit rates, as seen in figure 11. In particular, the changes around the new and old entitlement are small.

For the universal reform, the estimated shifts in the time profile confirm the overall visual findings from the earlier non-parametric estimates. In figure 12, there are changes near the old and the new entitlement, but also in the initial stages of the time profile, with exit rates increasing by roughly 12%–18% in the first ten weeks in relative terms. A natural explanation would be that some of the unemployed expect that, given some baseline strategy, there is some risk of prolonged unemployment. With a shorter entitlement,

[20] The parametric estimates were calculated using data at the accuracy of .002 FTE benefit weeks.

[21] Models allowing for recurrent spells by an individual did not appear very helpful in this context for two reasons. First, only a small minority in the treatment group for the targeted reform had preceding UI spells. Second, later spells often had different new entitlements, which could itself be a potential outcome. Recurrence is analysed separately in appendices N and O.

there is thus a greater risk of significantly reduced income if unemployment is very long. To avoid this risk from actualising, they increase their search effort or otherwise adjust their choices, leading to earlier exits.

Figure 11: Estimated hazard and the treatment effect, targeted reform.

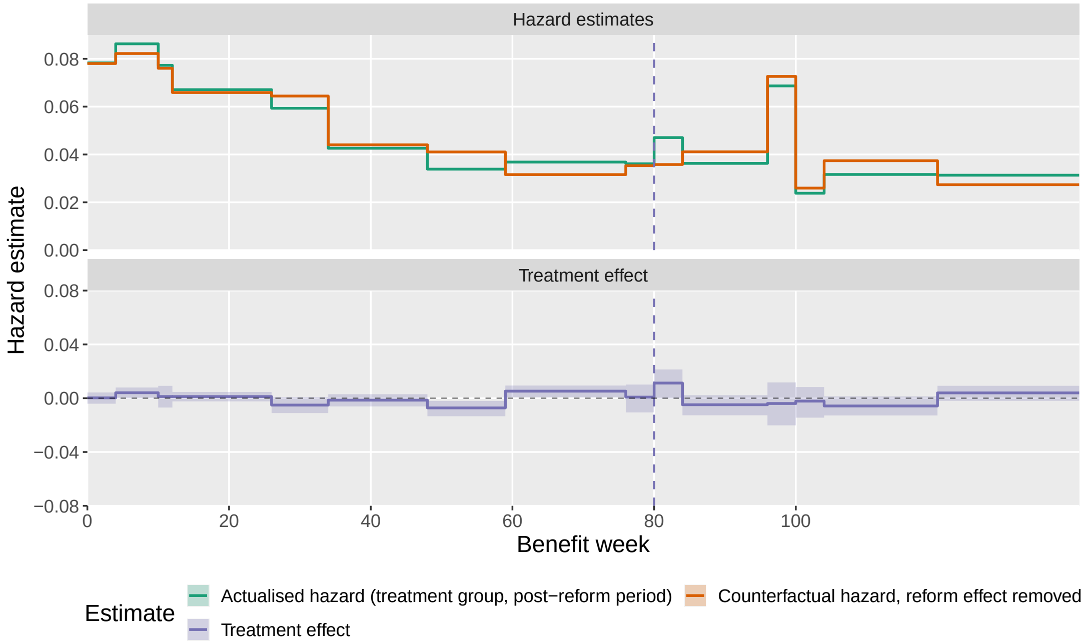

Figure 12: Estimated hazard and the treatment effect, universal reform.

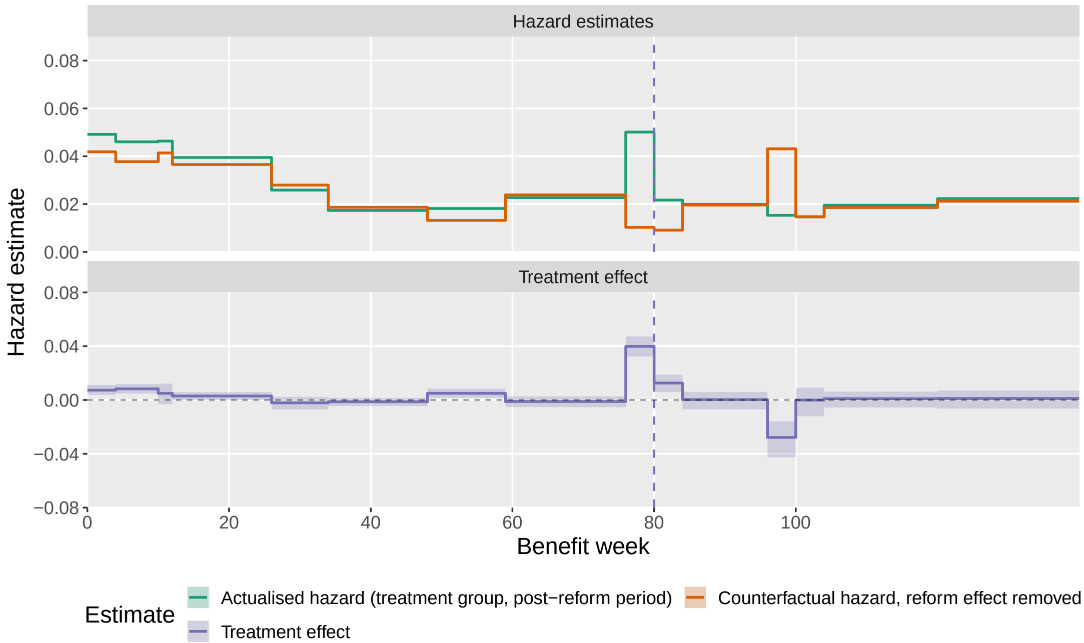

The estimated hazards are for a treated individual in the post-period. The counterfactual corresponds to an estimate where the $Treatment \times After$ parameter is set to zero. The plotted treatment effect is the difference between the two hazards. The shaded areas correspond to bootstrapped 95% confidence intervals for this difference.

7 Sensitivity analysis

7.1 Balancing

To guard against some potential threats to identification detailed below, samples for the targeted and the universal reform were separately balanced using entropy balancing weights. This subsection outlines the different threats in each case and the steps taken to overcome them. The empirical results in Section 6 were presented with and without these weights.

In both cases, the weights were intended to strengthen the primary identification strategy. The difference-in-differences setup is still meant to carry the heavier burden of removing confounding due to unobservable factors, and weighting is only meant to remove potentially important observable changes in the group composition over time. It turned out the main results were not highly sensitive to whether weights were used or not. However, one of the placebo tests used for the universal reform, as shown in Subsection 7.3, fails without the weights but passes with them.

7.1.1 The targeted reform

One concern for the targeted reform was the error in measuring the actual entitlement through observed prior employment, as explained in Subsection 3.2. Fortunately, the data for spells starting *before* the reform includes an indicator confirming whether the funds considered the person to have at least three years of work history.[22] Thus, wrongly classified spells where the estimated employment does not match the fund's assessment can be dropped from the pre-period. The post-reform control and treatment groups are then separately weighted to match their pre-reform counterparts.

The second issue to be tackled was the loosening of the recent employed condition, which was described in Section 2. This change increased the number of new UI entitlements in the post-reform period. Because recent employment and overall work history are correlated, the relative increase was more prominent among the treatment group with shorter histories. Because recent employment is also measured with some error, the observed change is a shift in the entire distribution, which cannot be fully resolved by simply dropping those with less employment than the pre-reform condition. For this reason, the entire distribution of contribution weeks for each group is balanced to match the pre-reform counterpart.

Several other important observables are simultaneously targeted to ensure this weighting does not distort balance in other aspects.[23] The weights are the entropy balancing

[22] This is because all such individuals were granted a short-term increase at the start of entitlement before 2014, and this increase is observed in the data; see Subsection 7.5.

[23] The observables and the balancing method were chosen using a set-aside test sample from another period. The choice criteria were reduction in observable classification error, balance for prognostically

Figure 13: Balance for observed *recent* employment, treatment group.

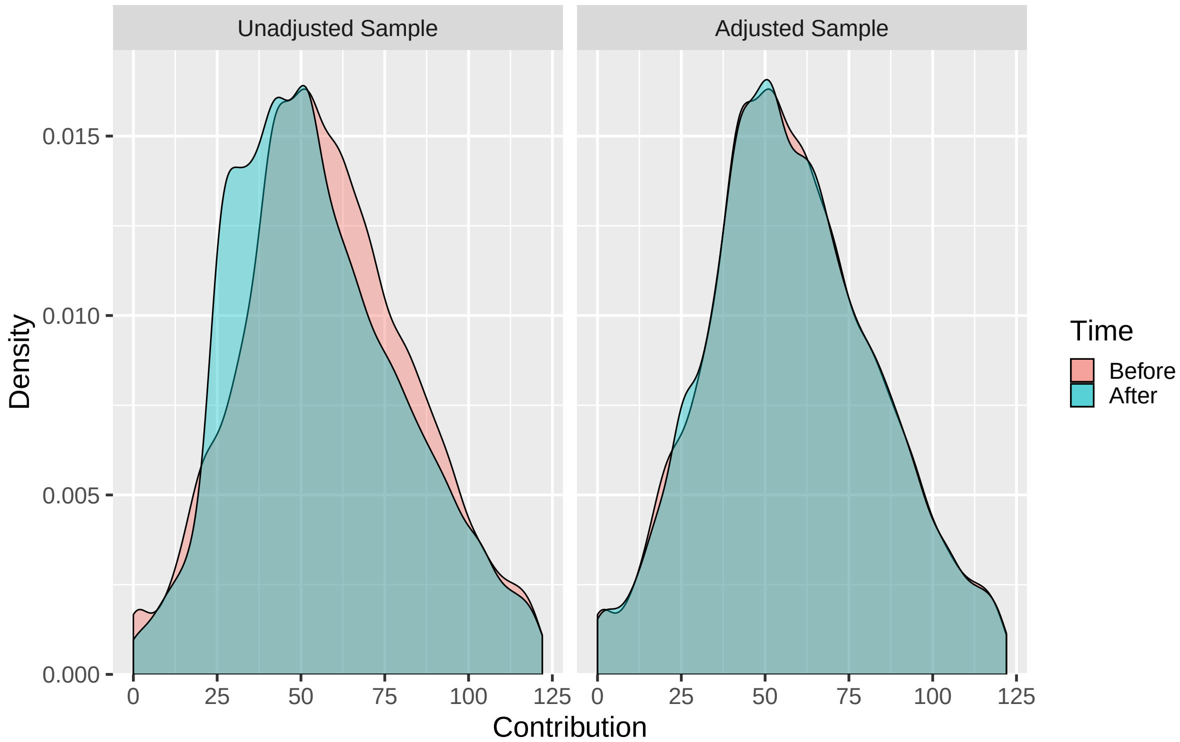

The measure is observed weeks contributed towards the employment condition. The unadjusted distributions to the left are without weights, and those on the right with weights.

weights introduced by Hainmueller (2011), which directly target the balance in covariates. Besides recent employment, balancing was done simultaneously on overall employment history, age, unemployment fund (the roughly 30 different funds serve as a joint proxy for profession and industry), gender and time from previous employment. Extreme weights are trimmed to the 95th percentile.

Observed treatment period measurement error drops to 13.9% for the treated and 7.7% for the controls, or 6.5% for the entire estimation sample. Figure 13 illustrates the balance acquired for recent employment weeks for the treatment group. Balance in some continuous covariates appears in Figure 14. Appendix T has additional figures for balance in discrete covariates, for the control group, and for the universal reform's groups.

7.1.2 The universal reform

For the universal reform, the concern motivating balancing is that the inflows into unemployment in December (controls)/January (treated) may vary differentially across years. For example, during an economic downturn, a larger than usual number of contracts in

significant covariates and effective sample size; see Appendix T.

Figure 14: Balance for selected covariates, targeted reform, treatment group.

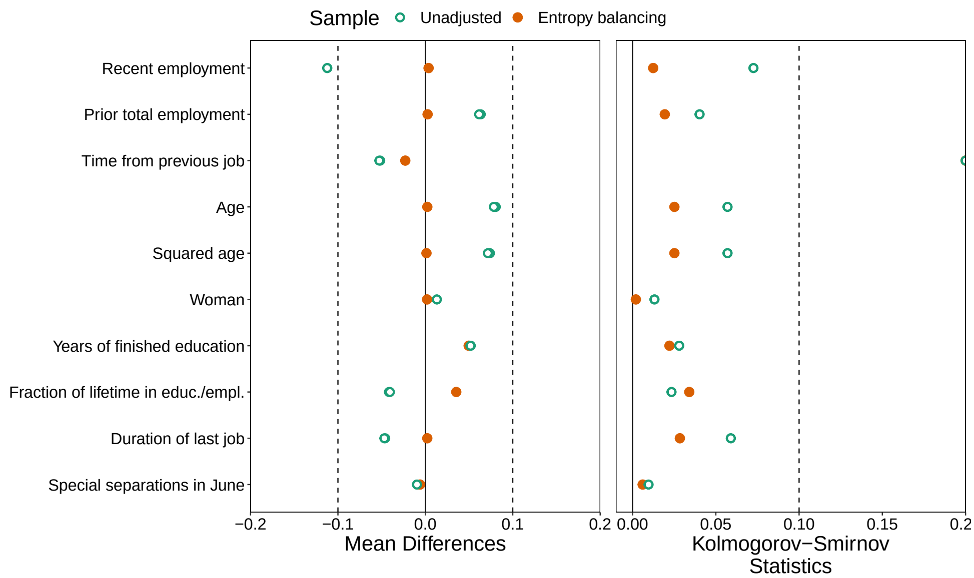

The comparison is for entitlements that started before reform vs. after reform. Several example covariates that were not targeted by weighting are also included to visualise the overall balance.

certain professions or industries may end in December. The weighting method is the same as for the targeted reform, but as the sample and the purpose of balancing are different, the selected covariates are now age, squared age, profession (ISCO level 1), work history, lifetime wages, number of children, residence permit status, local and regional characteristics including lags of labour market tightness and unemployment rates and dummies for the previous unemployment spell's presence and duration.

In this case, all groups are simultaneously weighted to one reference group (the post-period treatment group).

7.2 Prior trends

For the targeted reform, a visual analysis of prior trends supports the parallel trends assumption: the difference in unemployment durations across the two groups is quite stable across almost two decades. Durations for identically defined groups (short histories, slightly longer histories, and other histories) are decomposed in Figure 15 using LOESS.[24] As can be seen, the durations for the control and treatment groups tend to be similar but often diverge from the durations of the other unemployed. Regarding the universal reform, most other year pairs involve other reforms that affected spells starting in December or January very differently. This makes visual trends less informative for the later reform than the direct placebo check discussed next.

[24] The trends cover durations until UI exhaustion, as data on UA is only available from 2010. Trends from 2010– using spells with transitions to UA included are qualitatively similar.

Figure 15: Prior trends for the targeted reform.

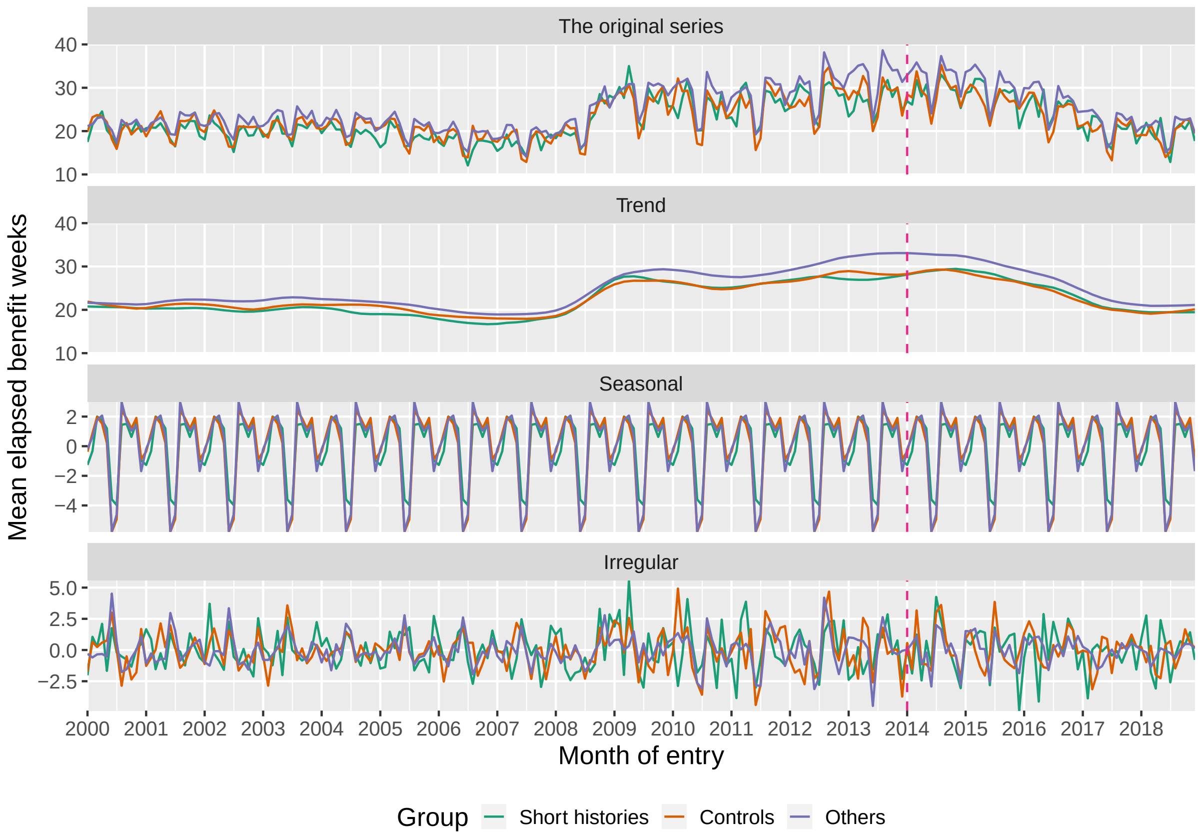

Mean duration in unemployment for fresh UI entitlements. Sample selection and treatment/control group definitions are the same as with the targeted reform, but the follow-up is only two years, and only UI is tracked.

7.3 Placebo reform

Two placebo checks, one for each reform, were performed. For the targeted reform, the latest prior pair of years was used. For the universal reform, the year-pairs 2010/2011 and 2011/2012 are the most comparable with the smallest number of other reforms. Using more distant years was necessary in this case because other year-pairs had other legal reforms applying sharply for new entitlements at the year-turn.[25]

A procedure identical to the respective main setup was repeated in both cases. The process included the same steps of sample selection, weighting, and estimation. Placebo effects on the mean duration for both tests are summarised in Table 5. In both cases, the placebo effects are close to zero when using both weights and controls, consistent with a (conditional) parallel trends assumption. However, the point estimate for the placebo against the universal reform is more sensitive to the weights and controls, and

[25] The issue is that many reforms specifically applied to spells starting from January, treating starts in December differently from those in January. Such reforms are typically less of an issue for the targeted reform's setup as long as they applied similarly to the unemployed with short and slightly longer work histories.

the estimate when not using either is relatively far from zero. This suggests that the additional balancing is necessary for the year-turn setup.

Table 5: Placebo tests for the reforms.

Reform	Unweighted, without controls	Unweighted, with controls	Weighted, without controls	Weighted, with controls
Targeted reform: years 2011–2012	0.63 (1.05)	-0.66 (1.00)	-0.20 (1.07)	-0.50 (1.00)
Universal reform: turns of year 2010/2011 and 2011/2012	-1.338 (1.270)	-0.796 (1.271)	-0.129 (1.363)	0.043 (1.362)

Effects on unemployment duration (in FTE benefit weeks). Bootstrapped standard errors in parentheses.

7.4 Alternative setups

Several alternative setups were used as further specification checks. As most involve many details, a more thorough analysis of each is saved for Appendices K, L and O.

The first check (Appendix L) was against the universal reform, using older cohorts alone. As acknowledged in Section 1, the reform was only *near*-universal: those aged 58 when entitlement was determined kept the old entitlement after 2017.[26] The default samples thus excluded the oldest cohorts. As an alternative setup, those aged 57 were used as the treated (entitlement cut) and those aged 58 the controls (entitlement not cut); full years 2016 and 2017 were the pre- and post-periods. The point estimate of the effect of the universal reform in this check was a -2.5 (s.e. 1.6) weeks change in unemployment duration. While the result is imprecise due to the sample restrictions required, the point estimate is quite close to the results in Table 3.

The second check (Appendix K) examined how the short-history group responded to the later reform. Two alternative setups were used for this question.[27] First, the setup for the earlier targeted reform was replicated for entire years 2013 (base year) and 2017 (post-reform year). From 2013 to 2017, entitlements were cut from 100 to 60 weeks for the short-history treatment group, and from 100 to 80 weeks for the comparison group with slightly longer histories. The DiD estimate for the duration of unemployment from this setup is $+3.0$ weeks (s.e. 0.81).

As a second part of this check, consecutive years were compared: 2016 and 2017. The short-history group was chosen as before, but this time, those with 5–19 years of work history constituted the comparison group.[28] The DiD estimate for this setup was $+1.23$ (s.e. 0.58).

[26] Those aged 59 or older at the time typically had unlimited entitlement, studied by, for example, Kyyrä and Pesola (2020a).

[27] As the short-history group is a relatively small minority of UI recipients, examining its responses as part of the main year-turn specification proved infeasible due to issues with power and the weighting methods.

[28] This change was made to reduce measurement error and is justified in more detail in the appendix.

To place both of these results in context, suppose that the baseline results from the universal reform holds on both of the above control groups with longer work histories. Since the control groups experienced a 20 week cut to the entitlement in both of the above cases, this would mean the reform cut their mean unemployment duration in both cases by roughly 3 weeks. The first auxiliary DiD estimate of $+3.0$ would then mean that even a -40 week cut to the entitlement caused no change among the short-history group. The second auxiliary result of $+1.2$ weeks would, in turn, mean that a -20 week cut caused a $-3.0 + 1.2 = -1.8$ week change in the mean unemployment time among the same group. The results are thus quite different.

As both the auxiliary setups suffer from multiple issues detailed in Appendix K, both results are likely to be only approximate. Both findings are consistent with a weaker response among the short-history group, but they cannot be used to conclude that the group never reacts at all. This leaves open the possibility that the non-response to the first reform was driven partly by the group targeted and partly by the administrative implementation.

Finally, both reforms were extended to include re-entries into unemployment (appendix O). For each entitlement, cumulative UI weeks over a follow-up of 2 years and 10 months was now the measure of unemployment duration, even if there were longer gaps between these periods of unemployment. The results were qualitatively in line with only considering the first spell: the average cumulative duration of spells only responded weakly in the targeted reform. In contrast, the universal reform caused a larger, statistically significant change. Additionally, for the targeted reform, the exit spike remained at the agency switch time rather than the new UI expiration time.

7.5 Changes in benefit levels

Simultaneously with the targeted reform, a four-week initial benefit increase was phased out. This increase had until then been available to those with at least three years of work history, i.e., those *not* treated by the entitlement cut. At face value, this might suggest there were two competing reforms: an entitlement cut applied to the treated and a reduction in initial benefits to the control group.

Three factors attenuate the probability that this phaseout was important for our results. First, because of the short duration, the median mechanical impact of the change on total benefits per spell in the control group was only -1.2%. A mechanical decomposition of different benefit rule changes on each group is given in Appendix S. Many other such changes had a larger effect on empirical benefits, but these other changes applied symmetrically to both groups.

Second, earlier research can be used to assess the likely pattern of changes when they do occur. Uusitalo and Verho (2010) studied a reform in 2003 in Finland that introduced similar but longer schedules of increased benefits. The daily level of this increase was comparable to the increase available until the targeted reform, but the duration and targeting were quite different. The 2003 reform targeted those with at least 20 years of work history, and the increase had a maximum duration of 30 weeks. The authors find that the impact of this increase appeared to dissipate after the increase period was over. If the shorter increase phased out with the targeted reform in 2014 also had significant effects, one would expect to find them around the time the individual exhausts the increase duration rather than much later.

The estimated hazards in Section 6 and Appendix Q imply that the difference in exit hazards between the groups stayed stable during and around the first four weeks. If the increase were an important driver for the exit rates of the controls in the first place, one would expect early-stage hazards to change differently for the controls from the treated.

Third, the impact of the increase on early-stage hazards was directly investigated using two alternative setups: one when the increase was introduced in 2010 and one when the increase was phased out. For the first case, the setup was the same as for the targeted reform, and only the control and treatment groups were flipped: short-history individuals saw no change, while those with slightly longer histories had their initial benefits increased. The difference in the exit hazard for the first 20 weeks was quite stable. Appendix S shows the resulting hazard estimates and documents the second case, which required a very different sample but showed qualitatively similar results.

7.6 Anticipatory effects

The official government proposal for the targeted reform was introduced in August 2013; it relied upon agreements publicly reached in March 2012. Technically, some short-history individuals who anticipated the entitlement cut could try to time their entry into unemployment in December 2013 rather than January 2014 to maximise their entitlement.

Such self-selection appears unlikely. First, persons who voluntarily quit their jobs (excluded from the samples used) have a 3-month waiting period for benefits, eliminating any gains from this type of self-selection. Second, the empirical patterns do not suggest increased bunching of treatment group entries around the turn of the year 2013–2014. Weekly entries around various year-turns are plotted in Figure 16.

8 Discussion

According to the difference-in-differences estimates obtained here, a targeted reduction in the maximum potential duration of UI by 20 weeks (20%) did not significantly affect the average unemployment duration among those targeted, the unemployed with short work histories. In contrast, a later universal reduction of the same amount reduced the mean

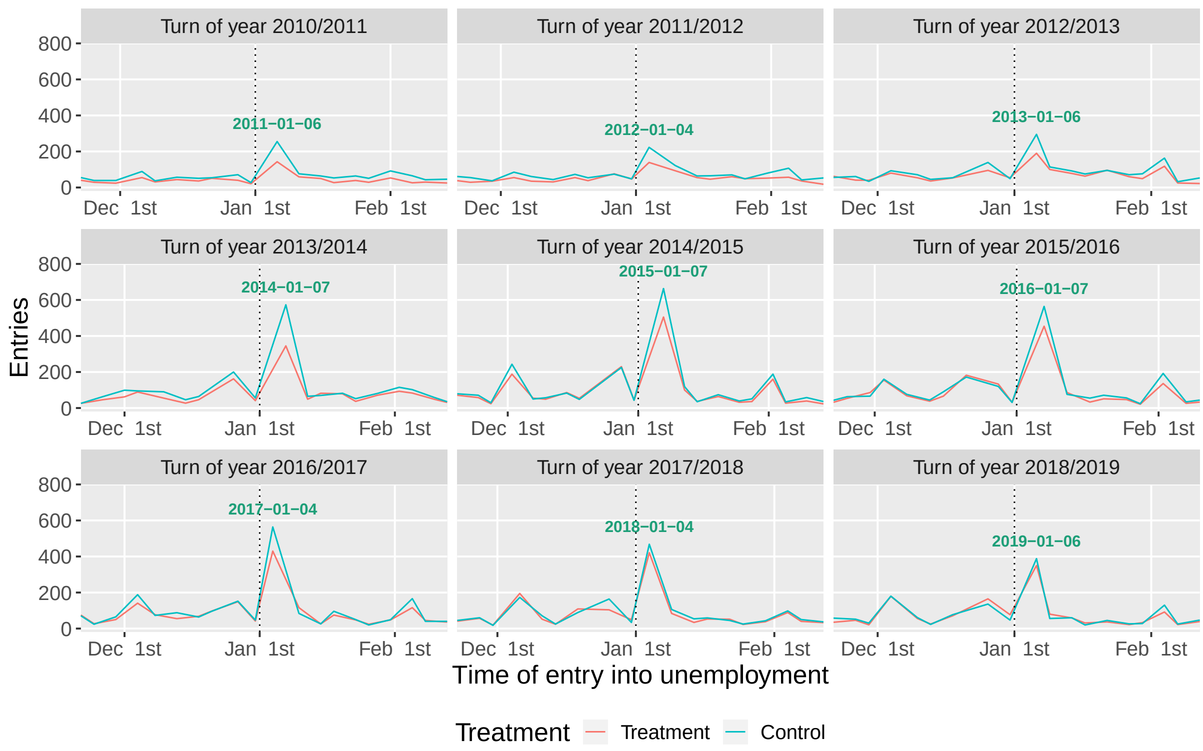

Figure 16: Weekly entries into unemployment by treatment (work history) status around the turn of the year. The end date for the peak entry week for each turn of the year is highlighted.

length of unemployment by roughly 3 weeks.

The latter result is very similar to earlier results for Finland by Kyyrä and Pesola (2020b) and comparable to international estimates. In their literature review, Tatsiramos and Van Ours (2014) present that one extra week of entitlement is expected to prolong unemployment by about 0.2 weeks. Both the earlier results for Finland and the present results for the universal reform imply an elasticity of unemployment to entitlement duration of about 0.5. This is not far from the literature consensus estimate of 0.4 offered by Schmieder and von Wachter (2016).

The earlier results for Finland were obtained using differences in entitlement amongst those who re-enter unemployment and either start a new entitlement or consume their old residual entitlement. This setup required various sample restrictions. In this paper, very similar results have been obtained among new entitlements starting around the turn of the year and also for those aged 57 or 58 (subsection 7.4). Corroborating the earlier Finnish findings across different strategies and groups is helpful for at least two reasons. First, it increases confidence in the external validity of the findings, i.e., that the *average* effect of cutting entitlement by one week is likely to be about 0.16 weeks shorter spells of unemployment. Second, these results also provide a reasonable benchmark for the short-history group. The comparison of the results, in turn, implies that the weak response to the targeted reform was likely specific to the target group rather than a feature of the Finnish labour market.

There are two caveats. First, the average effect may not hold for all groups, such as those with short work histories, who are still targeted by shorter entitlements. Second, while results on time in unemployment were similar across studies, results on the quality of re-employment differ.

Neither reform had a significant effect on wages or duration of the next job within a year of an exit from unemployment. This contrasts with the results by Kyyrä and Pesola (2020b), who found insurance duration cuts to have adverse effects. The different results might arise because the papers examined different groups: Kyyrä and Pesola relied on groups with relatively specific bands of recent employment for identification, while this paper used those entering unemployment around the year-turn for the universal reform. The results on this topic in previous literature are also varied. Card, Chetty, and Weber (2007), Lalive (2007), Le Barbanchon (2016) and Schmieder, von Wachter and Bender (2016) find that extending the UI entitlement has minor negative effects or no effects on subsequent job quality. Nekoei and Weber (2017) report minor positive effects on wages, while de Groot and van der Klaauw (2019) report mixed results within a single paper.

The empirical results do not explain *why* those with short work histories respond less than others to cuts in the entitlement. One possible driver is that these individuals have not faced long unemployment spells before (see Table 1 in Section 4): most enter UI for the first time and have only had relatively short UA spells before. Mueller, Spinnewijn, and Topa (2021) show that many unemployed have overly optimistic and persistent beliefs about their job-finding probabilities. Such beliefs might be particularly prevalent among those who have not experienced long-term unemployment before. If persons in the early stages of their careers commonly overestimate the ease at which they can find a job before UI is exhausted, they might underreact to changes in the entitlement, thinking it does not concern them.

Another hypothesis is that less experienced cohorts might be more likely to use education as a fallback strategy to exit unemployment. This could make them less worried about exhausting their UI benefits. This hypothesis gains no support from the available evidence: among the short-history group roughly 7% of UI spells, at most, end in observable transitions to education (see Appendix H), compared to about 3% of all UI spells. The magnitudes are too low to offer a credible explanation for the weak response.

The weak response studied here is related to the length of the *overall* history. Even among individuals with shorter work histories, observed recent employment varies significantly. Persons with less recent employment may be more responsive, as Le Barbanchon, Rathelot, and Roulet (2019) found was the case in France, *and* those with shorter work history less so, as shown in this paper and earlier by Van Ours and Vodopivec (2006) for Slovenia. The available evidence thus suggests that if the entitlement is to vary based on prior employment and the aim is to increase employment, conditioning it on recent employment rather than the entire work history is probably more effective.

In the universal reform, most of the reduction in mean time in unemployment comes from higher exit rates in the first ten weeks of unemployment. In particular, merely moving the exit spike at exhaustion by 20 weeks only has a small effect on average durations, because only a small share of spells is left at the end of entitlement (whether at 100 or 80 weeks). While cutting the entitlement may still help reduce spells that would otherwise be very long, such a reform might not be particularly effective at incentivising those who remain unemployed for long.

The evidence on the exit spike in this paper suggests that the true exit spike is much flatter when time from exit to next job is adjusted for. It appears that frictions related to switching to a new benefits agency, rather than the drop in benefits, are causing observed exits from the benefits system to bunch at the exhaustion time. If enough data is available, the adjustment presented in this paper can be used to test whether similar frictions might also be driving the spike in other countries.

References

Boone, Jan and Jan C Van Ours. "Why is there a spike in the job finding rate at benefit exhaustion?" In: *De Economist* 160.4 (2012), pp. 413–438.

Card, David, Raj Chetty, and Andrea Weber. "The spike at benefit exhaustion: Leaving the unemployment system or starting a new job?" In: *American Economic Review* 97.2 (2007), pp. 113–118.

de Groot, Nynke and Bas van der Klaauw. "The effects of reducing the entitlement period to unemployment insurance benefits". In: *Labour Economics* 57 (2019), pp. 195–208.

Greifer, Noah and Elizabeth A Stuart. "Matching Methods for Confounder Adjustment: An Addition to the Epidemiologist's Toolbox". In: *Epidemiologic reviews* 43.1 (2022), pp. 118–129.

Hainmueller, Jens. "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies". In: *Political Analysis* 20 (Aug. 2011).

Kyyrä, Tomi and Hanna Pesola. "Long-term effects of extended unemployment benefits for older workers". In: *Labour Economics* 62 (2020), p. 101777.

— "The effects of unemployment benefit duration: Evidence from residual benefit duration". In: *Labour Economics* 65 (2020), p. 101859.

Kyyrä, Tomi, Hanna Pesola, and Jouko Verho. "The spike at benefit exhaustion: The role of measurement error in benefit eligibility". In: *Labour Economics* 60 (2019), pp. 75–83.

Lalive, Rafael. “Unemployment Benefits, Unemployment Duration, and Post-Unemployment Jobs: A Regression Discontinuity Approach”. In: *American Economic Review* 97.2 (May 2007), pp. 108–112.

Lalive, Rafael, Jan van Ours, and Josef Zweimüller. “How Changes in Financial Incentives Affect the Duration of Unemployment”. In: *The Review of economic studies* 73.4 (2006), pp. 1009–1038.

Le Barbanchon, Thomas. “The effect of the potential duration of unemployment benefits on unemployment exits to work and match quality in France”. In: *Labour Economics* 42 (2016), pp. 16–29.

Le Barbanchon, Thomas, Roland Rathelot, and Alexandra Roulet. “Unemployment insurance and reservation wages: Evidence from administrative data”. In: *Journal of Public Economics* 171 (2019), pp. 1–17.

Moffitt, Robert. “Unemployment insurance and the distribution of unemployment spells”. In: *Journal of econometrics*. Journal of Econometrics 28.1 (1985), pp. 85–101.

Mortensen, Dale T. “Unemployment Insurance and Job Search Decisions”. In: *Industrial and Labor Relations Review* 30.4 (1977), pp. 505–517.

Mueller, Andreas I., Johannes Spinnewijn, and Giorgio Topa. “Job Seekers’ Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias”. In: *American Economic Review* 111.1 (Jan. 2021), pp. 324–63.

Nekoei, Arash and Andrea Weber. “Does Extending Unemployment Benefits Improve Job Quality?” In: *American Economic Review* 107.2 (Feb. 2017), pp. 527–61.

Schmieder, Johannes F. and Till von Wachter. “The effects of unemployment insurance benefits: New evidence and interpretation”. In: *Annual Review of Economics* 8 (2016), pp. 547–581.

Schmieder, Johannes F., Till von Wachter, and Stefan Bender. “The Effect of Unemployment Benefits and Nonemployment Durations on Wages”. In: *American Economic Review* 106.3 (2016), pp. 739–777.

Tatsiramos, Konstantinos and Jan C Van Ours. “Labor market effects of unemployment insurance design”. In: *Journal of Economic Surveys* 28.2 (2014), pp. 284–311.

Uusitalo, Roope and Jouko Verho. “The effect of unemployment benefits on re-employment rates: Evidence from the Finnish unemployment insurance reform”. In: *Labour Economics* 17.4 (2010), pp. 643–654.

Van Ours, Jan C and Milan Vodopivec. “How shortening the potential duration of unemployment benefits affects the duration of unemployment: Evidence from a natural experiment”. In: *Journal of Labor economics* 24.2 (2006), pp. 351–378.

Zhao, Qingyuan and Daniel Percival. “Entropy Balancing is Doubly Robust”. In: *Journal of causal inference* 5.1 (2016), pp. 41–55.