Moral Hazard among the Employed: Spillover effects

Mette Ejrnæs, Stefan Hochguertel, and Andrew C. Johnston*

February 24, 2025

Abstract

Unemployment insurance (UI) is known to create moral hazard among the unemployed, but does it also do so among the employed? We leverage a Danish reform that quasi-randomly induced some workers to enroll in coverage at an age threshold. Using this natural experiment to address selection, we find that UI coverage significantly increases unemployment among the employed. Coverage also has spillovers, reducing employment among coworkers and spouses who were not directly affected by the reform.

JEL Codes: J24, J32, J45, H55, M52

*Ejrnæs, University of Copenhagen: mette.ejrnes@econ.ku.dk; Hochguertel, Vrije Universiteit Amsterdam: s.hochguertel@vu.nl. Johnston, UC-Merced: acjohnston@ucmerced.edu. The authors are grateful for financial support from the Economic Policy Research Network (EPRN) and from the Danish National Research Foundation (grant DNRF-134) through the Center for Economic Behavior and Inequality at the University of Copenhagen. All errors our own.

1 Introduction

Designing optimal unemployment insurance (UI) requires a full accounting of moral hazard—the tendency of insurance to increase risky behavior. While extensively studied among the unemployed (Schmieder et al., 2012; Mitman and Rabinovich, 2019; Landais et al., 2021; Dahl and Knepper, 2022), moral hazard may extend to employed workers as well, potentially reducing work effort and increasing entry into unemployment. If so, the total effect of insurance on employment, output, and welfare may be substantially larger than what is found by studying moral hazard among the unemployed alone.

In this paper, we measure moral hazard among the employed.

Despite its potential importance, measuring moral hazard among employed workers presents unique challenges. In most labor markets, workers have uniform insurance coverage, precluding the possibility of an informative control group. Denmark, however, offers an exception to this rule. In the Danish system, UI is voluntary, and workers enroll in coverage for a monthly fee. This feature produces a useful control group when paired with a reform that induces quasi-random enrollment.

We leverage a Danish reform that induced some workers to enroll in UI coverage at an age cutoff. The reform lowered the age at which workers needed to enroll in UI to qualify for early retirement from age 50 to 40. This change triggered a substantial influx of workers to UI who otherwise would have likely enrolled later in life. We harness the induced timing of insurance coverage, using a difference-in-differences (DID) design with full-population administrative data.

While many studies exploit variation in benefit generosity or duration among unemployed claimants, ours examines variation in coverage that affects the employed as well.

Our research design leverages the sharp and sustained impact of the reform on insurance coverage. In the year of the reform, coverage jumps by 6 percentage points (about 8 percent) for exposed workers, and the increase is stable over time. We use reform-induced

coverage as a natural experiment. Treated and control groups appear similar before the reform, attributable to the age-based cutoff that allows us to compare treated workers to slightly younger workers that were not affected by the reform. The simple, stable change in coverage, moreover, makes interpretation straightforward.

To validate the design, we examine potential threats to identification. The primary threat would be if dynamic individual factors coincided with the reform's timing and threshold. This is unlikely. First, we observe similar trends in treatment and control before the reform. If time-varying factors were driving the results, we would expect to see dynamic factors differing before the reform.

Second, unemployment effects are precisely timed with a unique coverage delay. In Denmark, UI coverage begins 12 months after initial enrollment. We observe no change in unemployment at the time of reform-induced enrollment. The unemployment effect only materializes 12 months later when coverage actually takes effect, as we would expect from exogenous enrollment decisions. This pattern contrasts with endogenous enrollment decisions, where we observe unemployment rising *before* insurance coverage begins, indicating dynamic selection in endogenous enrollment.

These two findings suggest that the reform is orthogonal to time-varying individual factors at the cutoff. Consequently, the observed differences between treatment and control groups can be attributed to the effect of insurance rather than selection into it.

Having established the validity of our natural experiment, we apply our event-study difference-in-differences design to estimate the causal effects of UI coverage on key labor market outcomes. The data reveal substantial impacts. First, UI coverage increases unemployment benefit receipt, with the effect growing steadily over time. Four years after coverage begins, reform-affected cohorts have a 1.1 percentage point increase in receipt (XX percent change, significant at the 0.001 level). We see declines in employment of similar magnitude. Coverage reduces yearly earnings conditional on employment. These declines amount to 1,000 fewer Danish Kroner per year on average, a 2 percent decline.

We also find evidence that coverage has spillovers on other workers. For example, when already-insured workers work in a setting were few colleagues are exposed to the reform, they have no increase in unemployment. In contrast, when they work in a setting where many colleagues are exposed, their unemployment increases by 0.5 percentage points, despite no change in their coverage. These suggest that moral hazard is socially spread whereby it changes the workplace norms and behaviors, influencing even those not directly affected by the policy change.

Likewise, when we observe that when already insured workers are single, the reform has no effect on their unemployment. But when they are married to someone exposed to the reform, their own unemployment rate increases by 0.5 percentage points, despite no change in their own insurance rate. This again suggests a social transmission of moral hazard, where effort decisions are shaped by the insurance status of others in addition to their own.

Distortions to employment have broader social significance because employment generates positive externalities for society. When employed, workers generate tax revenues, add value to firms, and generate consumer surplus through production. We calculate the economic loss associated with moral hazard among the employed in our setting...

This paper contributes to a line of work examining the effect of insurance among the employed. Several authors have shown spikes in layoffs when workers become UI eligible (see, for example, Christofides and McKenna, 1996; Brébion et al., 2022; Van Doornik et al., 2023). Another branch shows that workers reduce effort with greater UI generosity (Ejrnæs and Hochguertel, 2013; Lusher et al., 2022; Ahammer et al., 2023). A growing body of work finds that cuts in unemployment benefit duration reduce inflows into unemployment and lead to separations timed to maximize benefit receipt. Tuit and van Ours (2010) show that duration cuts in the Netherlands reduced unemployment inflows. Similarly, Hartung et al. (2022) argue that the success of the German Hartz reforms was primarily due to the reduction in long-term benefit duration, which disincentivized

separations. Most recently, Gudgeon et al. (2024) provide compelling evidence that unemployment inflows are strategically timed so that workers become pension eligible at the time of UI expiration.

Our paper contributes to this literature in several ways. First, we provide a transparent instrument for insurance take-up that enables us to measure moral hazard among the employed. We are not aware of prior studies that can estimate the effect of insurance on the extensive margin, only on the margins of additional generosity or particular timing. Second, we examine the effects of insurance on a wide range of important life outcomes leveraging the trove of Danish registrar data. Third, we test for spillover effects from insurance onto spouses and siblings to see how insurance decisions are shaped socially and unemployment effects potentially compound through social networks.

2 Institutional Setting

Denmark has a unique unemployment insurance system that provides a productive setting to understand the effects of UI on employed workers. This section covers the institutional arrangements of the program from 1980 to 1998 and focuses on the UI system, the early retirement program embedded within it, and a significant reform in 1992 that provides exogenous variation in UI enrollment.

2.1 Insurance Mechanisms

Denmark is one of the few countries where UI coverage is voluntary (along with Sweden, Finland, and Iceland). Danish UI, rooted in the Ghent system, is organized around 35 private industry-specific UI funds Holmlund1999. These non-profit organizations do not have selection restrictions for applicants, and they finance UI benefits through membership

Table 1: UI parameters

period	Max duration	Employment qualifications
1980-1993	unlimited	26 weeks within last 3 years
1994	7 years	26 weeks within last 3 years
1995	7 years	52 weeks within last 3 years
1996-1998	5 years	52 weeks within last 3 years
1999-2009	4 years	52 weeks within last 3 years
2009-	2 years	52 weeks within last 3 years

fees, payroll taxes, and government subsidies.[1]

During the study period, Denmark had generous UI benefit durations by international standards: Until 1996, UI would pay benefits for 84 months. This was then reduced to 60 months and further decreased to 48 months in 1998 (see Table 1). The 1990s also saw the introduction of mandatory activation programs for UI benefit recipients, starting 12 months after they first enrolled in the UI program.

Individual worker contributions to UI are independent of earnings, and employers typically do not observe their workers' insurance status.[2]

UI benefits are calculated as 90% of previous earnings, subject to a nominal minimum and maximum. In 1995, the annual ceiling was 132,000 DKK (approximately $22,000 USD), which was about half of the median of white-collar worker salary.[3] To be eligible for benefits, individuals must have been UI fund members for at least 12 months, except for those who have just completed their education.

Jobless individuals not covered by UI fund benefits, including those who have exhausted the maximum benefit period, can receive social assistance. The amount of social assistance depends on spousal income and individual circumstances and is considerably

[1] On average, workers pay about one-third of the actual premium, with the rest being subsidized Lentz2009.

[2] Parsons2003 report that the contribution paid by an individual amounted to about 3,600 DKK for an employed wage worker for 1995, excluding administration costs, which can vary substantially across UI funds.

[3] Sources: The maximum amount of unemployment benefit is stated by "Direktoratet for Arbejdsløshedsforsikring".

lower than UI for most UI fund members. To receive social assistance, individuals must register as unemployed [with the XXX office] and show that they are actively searching for a job. These are the same requirements as those for a jobless person covered by UI.[4]

The voluntary nature of the UI system leads to people enrolling when they expect they need insurance most. As a result, enrollment rates vary over time, with age, and across birth cohorts. Figure 1 presents a heat map of UI entry rates, where columns represent observation years and rows represent birth cohorts. The cells display the variation in enrollment rates across cohorts for each year, with different levels of entry colored from green (low) to red (high). Older individuals are less likely to enroll in UI, and higher entry rates are observed during periods of high unemployment, like 1981 and 1993. However, there are strong patterns that cannot be explained by either age, cohort, or year effects alone. The most salient pattern has to do with incentives in the early retirement system and a reform that took place in 1992.

2.2 The Early Retirement System and its 1992 Reform

The Danish early retirement (ER) system, introduced in 1979, aims to relieve 'worn-out' blue-collar workers. All workers, however, can access the ER system if they enroll in UI as fund members. The ER system is separate from the compulsory old-age retirement pension system, which allows workers to retire starting at age 67.

Until 1992, UI fund members aged 60 and older qualified for ER if they had been enrolled in the UI system for the last 10 years. This typically led to a spike in enrollment at age 50. During the study period, there was no additional premium associated with benefiting from the ER system, meaning that ER was available at zero marginal cost for

[4]The "active labor market policy" was one that shifted from the previous "passive" (public welfare), where unemployed individuals could receive unemployment benefits with few requirements. After 1994, to counterbalance generous and long-lasting benefits, unemployed individuals were required to actively search for new jobs, participate in job-training, and engage in other skill development to receive benefits [CITATION?]. As a result, Denmark spends 1 percent of GDP on encouraging the unemployed to find work.

interested UI members. ER benefits correspond to UI benefits and averaged about 68 percent of previous earnings in 1992. These benefits are generally higher than the flat-rate old-age pension and are not means-tested. Once an individual enters ER, other labor market activities are largely precluded, with only small-scale labor market participation being permitted (up to 200 work hours per year).

In 1992, a significant reform of the ER system took place. The reform required continued UI membership for at least 20 years (instead of 10) before early retirement, effectively decreasing the latest age for joining a UI fund from 50 to 40. Individuals between 40 and 50 in 1992 were required to join the UI fund that year and remain members until age 60 to collect ER benefits. See Figure **??**(a) in Appendix **??**.

The reform affected different age groups differently, providing exogenous variation in UI enrollment rates by cohort. The comparison between the age-based rules and the empirical entry rates (Figure 1) provides prima facie evidence that the ER incentive significantly impacts UI coverage rates, as the entry rates increased substantially for the affected cohorts.[5]

3 Data and Descriptives

3.1 Data Base and Samples

We use longitudinal administrative register data on the universe of individuals born between 1931 and 1958 residing in Denmark between 1980 and 1998. Our sample is limited to individuals aged 25–59, resulting in 16 million person-year observations from 1.1 million different individual people.

The data combine administrative records from various government authorities, allowing us to track individuals and their households over time at annual frequency. We focus on those initially uninsured who join the insurance system because of the reform.

[5]There was another reform in 1999, outside our observation period.

Figure 1: Heat Map of UI Fund Entry

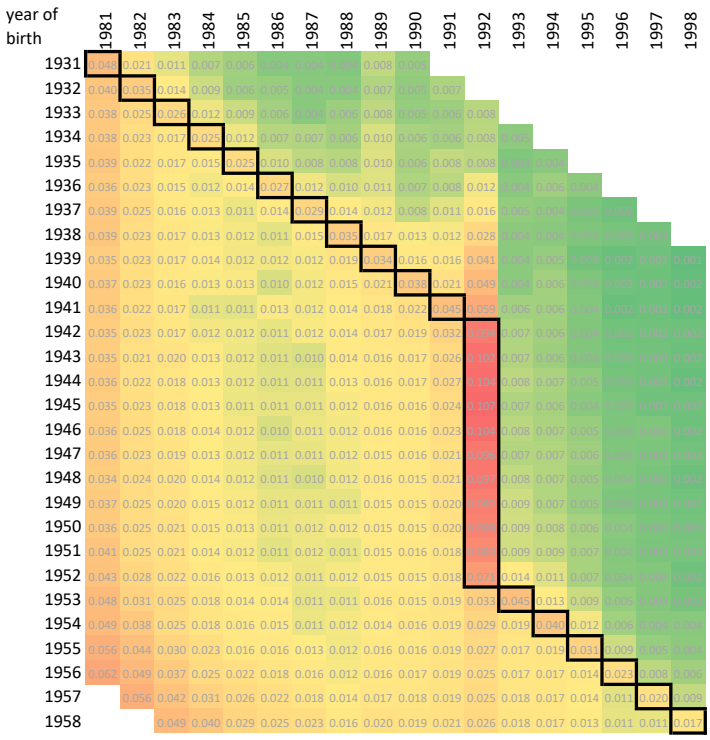

Note: This figure shows empirical entry probabilities of workers into the UI system, per year-of-birth and calendar year cell, conditional on not having been participating in the UI system in the year before. The color scheme aids in obtaining a quick visual impression of low (green) and high (red) probabilities. Data source: register data, Statistics Denmark.

We exclude individuals who were ever self-employed during the sample period.

For our difference-in-difference analysis, we use a strictly balanced sample of individuals with nine observations each. This balanced sample retains 5.1 million person-year observations from 570,000 unique persons. Our key variables include:

1. *UI Fund Membership:* We use information on membership in a UI fund, which is voluntary in Denmark.

2. *Unemployment Degree:* This is a continuous variable ranging from 0 to 1, representing the fraction of the year a person is unemployed. It reflects both the incidence and length of unemployment spells.

3. *Employment Degree:* Derived from Statistics Denmark's measure of total labor market experience, calculated from mandatory pension contribution information.

4. *Earnings:* We use yearly earnings, a comprehensive measure from all employment sources, aggregated within the year and constituting the major part of taxable income for individuals.

5. *Mobility Measures:*

    - Geographical mobility: moves between Danish regions.
    - Industry mobility: changes in industry classifications.

6. *Health Measure:* We use the incidence of sickness benefit receipt as a proxy for health outcomes. These benefits are paid by the municipality of residence and are uniform for all, regardless of UI status.

3.2 Descriptives

We present summary statistics in Table 2. Our main sample is 63.1 percent male and 97.1 percent Danish nationals, with an average age of 45.8 years. About one-third

have primary school as their highest educational attainment, two-fifths are vocationally trained, and about one in five has an academic degree. The average wage is 299,000 DKK (2005 prices). Four out of five are homeowners, 75 percent are married or partnered, and three in five have children. More than four-fifths are insured in one of the UI funds.

We also explore employment dynamics for several subsamples: (1) Individuals who sign up for UI at the reform notch, (2) individuals who sign up for UI not at the reform notch; (3) individuals who are always UI insured; and (4) Individuals who are never UI insured. These subsamples allow us to compare the labor market dynamics associated with UI take-up.

Table 2: Summary Statistics and Samples

Sample →	Large	Balanced	[1]	[2]	[3]	[4]
Subsample →			joins UI	joins UI	always	never
Variable ↓			at ER	not at ER	in UI	insured
age	42.7	45.8	45.5	44.9	45.8	47.1
Education						
elementary	0.307	0.303	0.175	0.267	0.333	0.237
high school	0.024	0.020	0.027	0.031	0.016	0.029
vocational	0.406	0.426	0.339	0.429	0.460	0.244
short further edu.	0.041	0.042	0.045	0.047	0.036	0.079
bachelor/college	0.156	0.148	0.360	0.172	0.099	0.299
master/PhD	0.052	0.050	0.044	0.040	0.045	0.101
Male	0.588	0.631	0.579	0.572	0.645	0.629
Danish	0.966	0.971	0.980	0.969	0.970	0.972
Experience	16.8	20.0	21.5	19.3	19.9	19.7
Unemployment	0.054	0.039	0.010	0.029	0.047	0.008
Unemployment Degree	0.062	0.044	0.010	0.030	0.054	0.008
Non-employment	0.059	0.039	0.010	0.031	0.048	0.010
Wage (1000 DKK_{2005})	276	299	345	313	286	344
Region						
Copenhagen area	0.230	0.222	0.248	0.279	0.203	0.279
North Zealand	0.132	0.137	0.174	0.174	0.122	0.183

Continued on next page

Table 2: Summary Statistics and Samples – *continued*

Sample →	Large	Balanced	[1]	[2]	[3]	[4]
Subsample →			joins UI	joins UI	always	never
Variable ↓			at ER	not at ER	in UI	insured
South Zealand	0.112	0.113	0.103	0.107	0.116	0.107
Funen Island	0.086	0.086	0.075	0.072	0.092	0.065
Aarhus area	0.156	0.154	0.138	0.131	0.163	0.126
South Jutland	0.084	0.084	0.085	0.076	0.086	0.078
Mid Jutland	0.111	0.113	0.105	0.097	0.118	0.095
North Jutland	0.089	0.090	0.072	0.064	0.098	0.067
Houseowner	0.769	0.808	0.858	0.787	0.804	0.813
Partnered	0.710	0.753	0.782	0.744	0.744	0.799
Children in HH	0.613	0.613	0.681	0.621	0.602	0.629
UI fund member	0.830	0.829	0.574	0.601	1.000	0.000
UI fund entry	0.017	0.021	0.112	0.111	0.000	0.000
Spouse UI fund entry	0.027	0.027	0.047	0.041	0.022	0.027
year	1989.1	1990.8	1991.6	1990.2	1990.9	1989.8
Nobs in 1,000	16,268	5,153	452	520	3,702	479
N indiv in 1,000	1,078	573	50	57	411	53

Figure 2 illustrates the impact of the 1992 reform on UI fund membership rates across different birth cohorts. Cohorts born in 1952 and earlier, who were 40 or older at the time of the reform, exhibit a sharp increase in UI fund membership in the year of the reform. The graph thus depicts the kind of variation we use to study the effect of insurance coverage, with those at or above the new cutoff responding most strongly to the changed incentives.

Before we present a causal design, Figure 3 shows descriptively how unemployment degree changes around UI enrollment 'events' for different groups over a nine-year period (-4 to +4 years from enrollment). The dashed red line represents workers enrolling precisely at the early retirement eligibility threshold (age 50 before the reform, 40 after). The navy dotted line shows the never-insured group, where year zero is when they cross the early eligibility threshold. These two groups have similar levels and trends before enrollment, but after coverage takes effect, unemployment increases in the enrolling group by 2.5

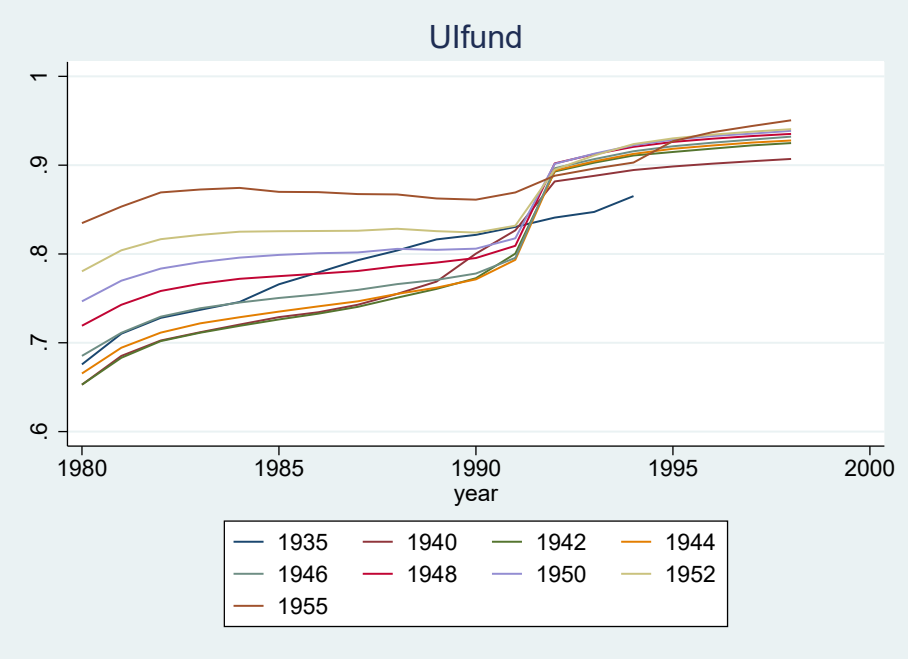

Figure 2: Caption

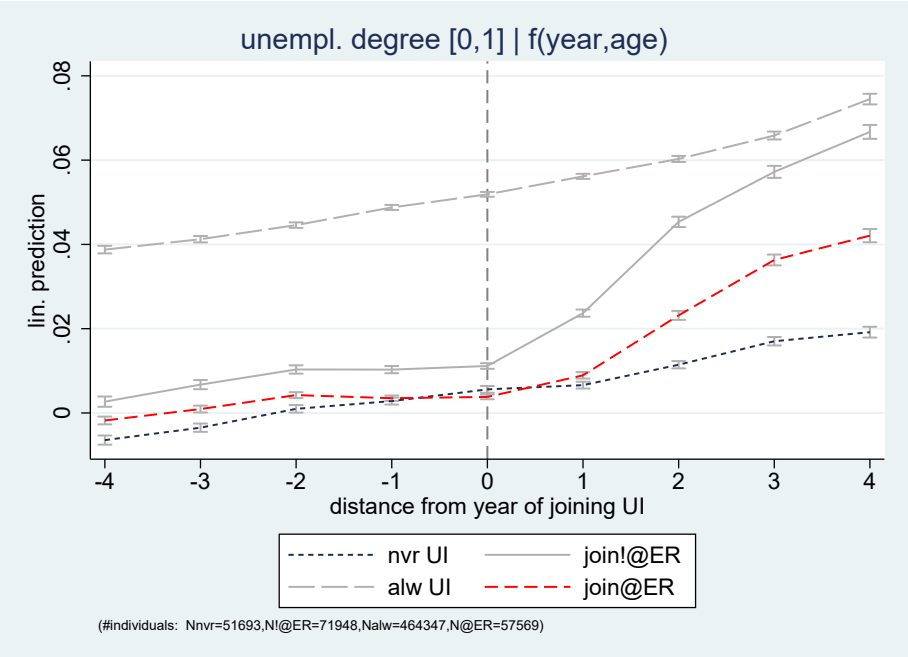

Figure 3: Caption

percentage points. The solid gray line represents those enrolling endogenously not at the early-retirement threshold. This group has higher baseline unemployment, with rates increasing after enrollment but before coverage, suggesting dynamic selection into coverage as workers anticipate potential unemployment. The top dashed gray line shows always-insured workers, with year zero again being when they cross the early retirement threshold. This group exhibits much higher baseline unemployment rates and slightly different trends compared to the other three groups. The y-axis represents the unemployment degree, ranging from 0 to 1, controlling for year and age effects.

While this descriptive analysis is interesting, it may not tell an accurate causal story. Namely, workers can anticipate the timing of the early retirement threshold, so in principle there could be anticipatory effects. In the next section, we zoom in on variation arising from an unanticipated reform to estimate the causal effects of coverage. This approach allows us to overcome potential endogeneity issues and provide more robust estimates of the impact of UI enrollment on unemployment.

4 Design and Results

To improve causal inference, we focus on variation in coverage generated by the 1992 reform. The reform sheds light on causal mechanisms because it generated large, unanticipated variation in UI coverage with quasi-random variation arising from the age threshold.

We implement this design by comparing individuals exposed to the reform to a comparison group of workers who were not exposed to the reform and have not yet crossed the early retirement threshold. Once a member of the comparison group crosses the threshold, they are dropped from the analytic sample, following Callaway and Sant'Anna (2021) and Goodman-Bacon (2021).

We model the outcome variable Y_{it} (outcome measure for individual i in year t) as a

function of the worker's exposure to the reform in event time while accounting for worker fixed effects and time controls:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{s=1988}^{1990} \delta_s D_i \times \mathbf{1}(t = s) + \sum_{s=1992}^{1996} \beta_s D_i \times \mathbf{1}(t = s) + X_{it}\zeta + \varepsilon_{it}. \quad (1)$$

Here, D_i indicates treatment based on age in 1992. It is one for each worker that was at least 40 years old, but younger than 50 in that year. The indicators $\mathbf{1}(t = s)$ refer to year dummies that equal one when a worker's outcome is observed in year t, and zero otherwise. The first sum includes pre-reform time so that the δ_s coefficients capture pre-eligible trends. The second sum includes post-reform event time so that the β_s coefficients describe the dynamic effect of the reform. We exclude a dummy for the period $s = 1991$ so that period is the omitted category and the implicit reference for comparison for the year coefficients.

If coverage increases unemployment rates, we would expect that the β_s would be positive. If coverage has no effect, we would expect them to be zero.

The sample includes all workers from birth cohorts 1952–1957 for the period 1988–1996. The treatment group consists of the 1952 birth cohort and the comparison group consists of 1953-1957 birth cohorts.[6] To avoid contamination, we drop control units once they turn 40 and therefore cross the early retirement threshold.[7] For example, we include in the sample the 1954 cohort from 1988 to 1993, when the cohort is not yet encouraged to enroll in insurance by the threshold. The sample selection ensures that the estimated effect is identified from comparisons between the treatment cohort and cohorts who are not yet treated.

[6] In principle, we could have included more treatment cohorts, but chose to focus on the youngest cohort so that the treatment and control groups were most similar, following the intuition of a regression discontinuity approach.

[7] We use individuals in the comparison group up to the year they turn 39.

Figure 4: Unemployment Insurance: Cohort graphs

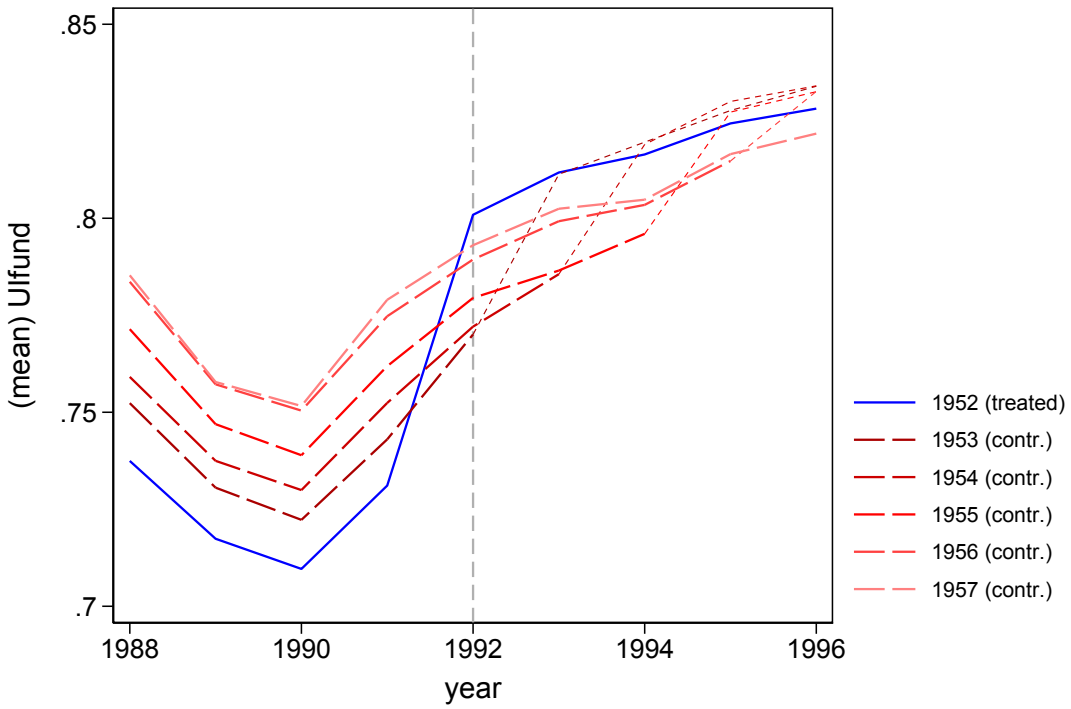

Note:

.pdfReform1992.doh:/workdata/704508/Data/figures/

To illustrate our estimation sample, we have depicted the relevant birth cohorts in Figure 4. The solid blue line shows the treatment cohort (the 1952 cohort), while the red long-dashed lines show the observations used from the not-yet-treated control cohorts. The short dashed lines represent observations we exclude from those control group once they have approached the threshold.

It is clear from Figure 4 that the 1992 reform had a large impact on the 1952 cohort. The fraction being insured increased from around 73 percent to 80 percent between 1991 and 1992. Similarly we also see for the other cohorts an increase at the year they turn 40, which was the age threshold for these cohorts respectively (e.g. the 1953 cohort turned

40 in 1993, when their fraction of insured increased from 76 percent to 81 percent.) The short-dotted lines represent when each control cohort is dropped from the analytic sample. We also drop students and individuals on disability.

4.1 Complier Analysis

Before presenting the estimation results, we compare the demographics of compliers to the already insured and the never insured. We divide the treatment cohort (born 1952) into four groups: the always takers (the already insured), compliers (those that insure at the reform), never takers (those that never insure), and defiers (those that leave insurance at the reform). The always taker group is the largest, and they account for 73.9 percent of the population.

The compliers are defined as those who sign up for insurance for the first time in the year of the reform. They make up 7.5 percent of the sample. The never takers are 17.2 percent. Finally the defiers are those who were insured before the reform and then left in the reform year. This group is small and makes up less than 1 percent of the sample. We present descriptive statistics for the four groups in Table 3, where all characteristics are measured before the reform in 1991.

The results in Table 3 reveal several interesting patterns about the characteristics of compliers compared to other groups.

First, compliers have the lowest unemployment intensity (0.02) compared to always takers (0.09), never takers (0.06), and defiers (0.15). This suggests that the reform primarily induced individuals with relatively stable employment to join the UI system. Second, compliers have the highest average years of education (12.84 years) among all groups. They also have the highest proportion of individuals with intermediate and high education (67%). This indicates that the reform was particularly effective in encouraging more educated workers to enroll in UI.

Interestingly, compliers have a much higher proportion of self-employed individuals

Table 3: Complier analysis for 1952 Cohort

	(1) Always takers	(2) Compliers	(3) Never takers	(4) Defiers	(5) All
Unemp. Intensity	0.09	0.02	0.06	0.15	0.08
	(0.22)	(0.11)	(0.20)	(0.29)	(0.22)
Gender: male	0.50	0.51	0.56	0.54	0.51
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
Years of educ.	12.24	12.84	12.19	11.63	12.26
	(2.80)	(2.68)	(3.16)	(3.03)	(2.87)
Danish	0.97	0.97	0.92	0.95	0.96
	(0.16)	(0.17)	(0.27)	(0.22)	(0.19)
Experience (years)	14.19	13.43	10.44	11.61	13.36
	(4.88)	(6.22)	(6.77)	(5.30)	(5.61)
Earnings	88.02	89.64	67.07	64.64	83.52
(1,000kr$_{1979}$)	(49.36)	(65.86)	(78.70)	(54.62)	(58.57)
Intermediate and	0.60	0.67	0.56	0.51	0.60
high edu.	(0.49)	(0.47)	(0.50)	(0.50)	(0.49)
Self-employed	0.05	0.17	0.15	0.08	0.08
	(0.21)	(0.37)	(0.36)	(0.28)	(0.26)
Out of LF	0.01	0.05	0.22	0.09	0.06
	(0.10)	(0.21)	(0.42)	(0.29)	(0.23)
No obs	50191	5214	11979	533	69359

Note: means; standard deviations in parentheses.

(17%) compared to always takers (5%), but similar to never takers (15%). This suggests that the reform was particularly effective in encouraging self-employed individuals to join the UI system, a group that traditionally may have been less likely to enroll. Finally, while compliers have slightly less labor market experience (13.43 years) than always takers (14.19 years), they have significantly more experience than never takers (10.44 years) and defiers (11.61 years). This indicates that the reform primarily attracted individuals with substantial labor market attachment.

These findings suggest that the reform was particularly effective in inducing UI enrollment among a group of workers who were relatively well-educated, high-earning, and stably employed. This complier profile has important implications for interpreting the effects of UI coverage on labor market outcomes, as it suggests that the marginal enrollees induced by the reform may be less likely to use UI benefits in the short term.

X

4.2 Results

We now present our estimation results from equation 1. We plot the δ_s and β_s coefficients over event time to show the dynamic effect of insurance coverage on outcomes. The confidence intervals are estimated using robust standard errors.

4.3 The Effect of Reform on Insurance Takeup

We first show the estimated effect of the reform on insurance enrollment in Figure 6. This can be thought of as the first stage in an instrumental variable approach, where the reform serves as an instrument for the endogenous decision to enroll in insurance. To be transparent about the control group and demonstrate the robustness of our estimates, we show how the figure changes when using different control cohorts.

Cohorts closest to the treatment cohort are most similar in age, but provide the

Figure 5: The 1952 cohort:

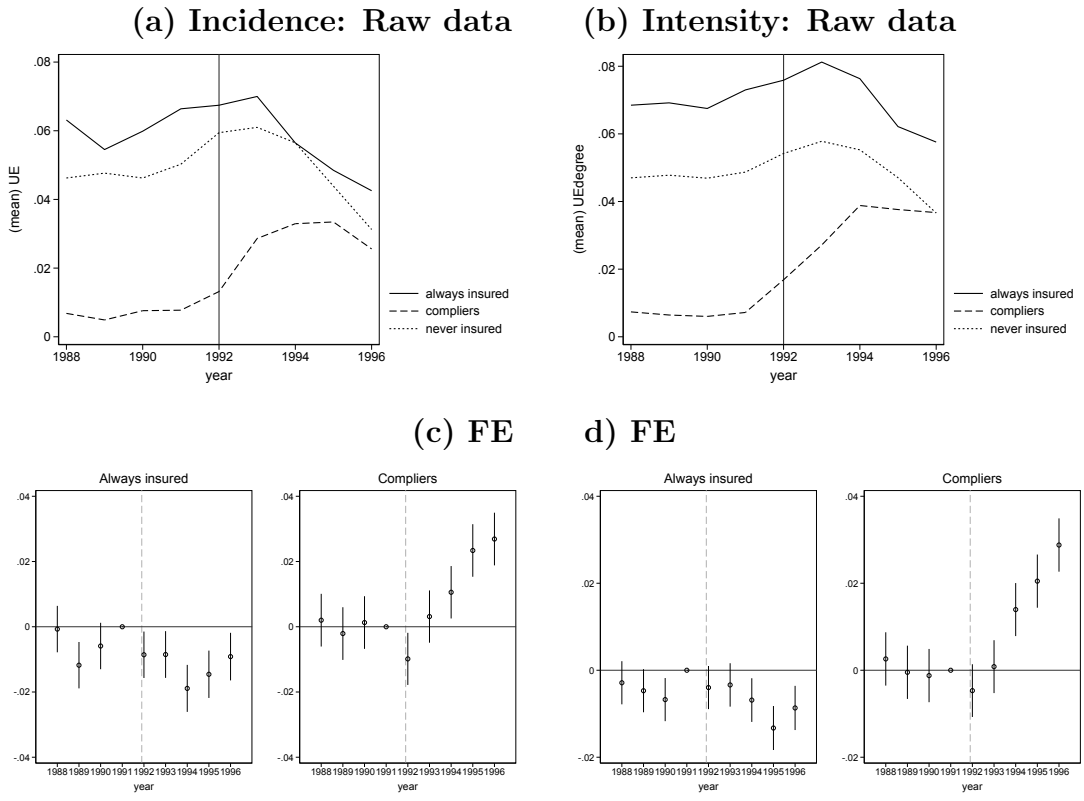

Note:

.pdfReform1992.doh:/workdata/704508/Data/figures/

shortest followup period because they soon are affected by the early retirement threshold. Thus, as we use younger and younger cohorts as the comparison group, we can estimate the effect for a longer post period. In the last figure, we estimate the event study using all five control cohorts (1953–1957). The results show similar magnitudes as we vary and eventually pool the control group.

In each case, we find that there is no differential pre-trend in the enrollment rates between treatment and control. At the time of the reform, we find a 5–6 percentage point increase in enrollment in the treatment group relative to comparison. The increase is immediate to the time of the reform and stable thereafter. The distinct timing and stability of the change aids interpretation of later results, since the reform introduces essentially a one-time stable shift in enrollment rates. This stability suggests that the reform's effect on enrollment was persistent and did not induce complex dynamic selection patterns over time.

4.4 The Effect of Reform on Unemployment Intensity

The primary focus of our paper is to understand moral hazard among the employed. We apply our difference-in-differences design to this outcome, which we present in Figure 7. Like before, we show the results while varying the control cohort to show robustness and, in the last figure, present the estimates pooling all the available control information.

Before the reform, treatment and control units have similar trends. Starting in 1993, a year after the reform, unemployment intensity begins to increase significantly, eventually reaching a 1.5 percentage point increase in unemployment intensity four years after the reform. The effects are quite large when one considers the fact that coverage only increased by about 6 percentage points. Without spillovers, the results would imply that insurance causes about a quarter of enrollees to become unemployed four years out. We will show, however, that spillovers play a key role explaining the large aggregate effects.

Consistently across the six figures in Figure 7, we find no effect in 1992, the year

Figure 6: Unemployment Insurance: Estimation results

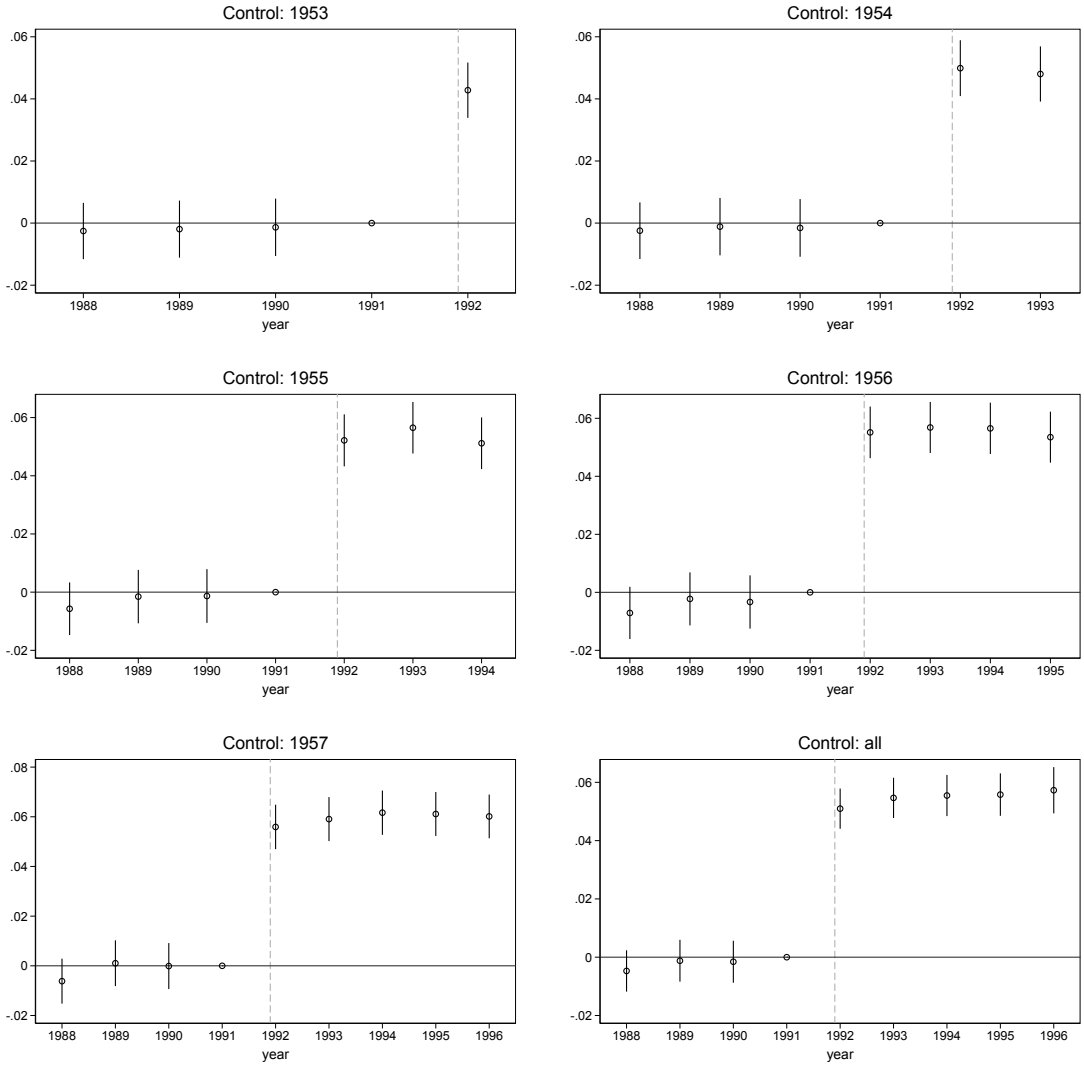

Note:

.pdfReform1992.doh:/workdata/704508/Data/figures/

enrollments are initially induced by the reform. This we find reassuring. Since coverage begins 12 months *after* initial enrollment, actual treatment is delayed by one year from enrollment. This result also supports the identifying assumption that time-varying factors are not correlated with treatment within our quasi-experimental design. For instance, if those who were induced to take up insurance in 1992 happened to be those that knew their unemployment risk was increasing for reasons unrelated to insurance, we would find unemployment rates rise somewhat *before* insurance takes effect, as we do when studying explicitly endogenous insurance decisions in the last section.

We also consider the effect of insurance on other labor market outcomes, namely earnings, another measure of unemployment, and non-employment. The estimation results are shown in Figure 8.

4.5 Robustness

To validate our research design and strengthen the causal interpretation of our main results, we conducted several robustness checks. The primary concern in any difference-in-differences design is the presence of confounding factors that might be correlated with both the treatment and the outcomes of interest. We address this concern through a series of placebo tests and alternative specifications.

Placebo Test: Never-Insured Individuals

We implement a placebo test in which we examine the "effects" of the reform on individuals who were never insured throughout the study period. If our estimated effects are truly due to the UI reform and not to other unobserved factors affecting the 1952 cohort, we should observe no effect of the reform on this group. Figure 9 presents the results of this placebo test.

As shown in Figure 9, the estimated effects on unemployment incidence for the never-insured group are consistently close to zero and statistically insignificant across all years, including the post-reform period. This lack of effect provides strong support for the

Figure 7: Unemployment Intensity: Estimation results

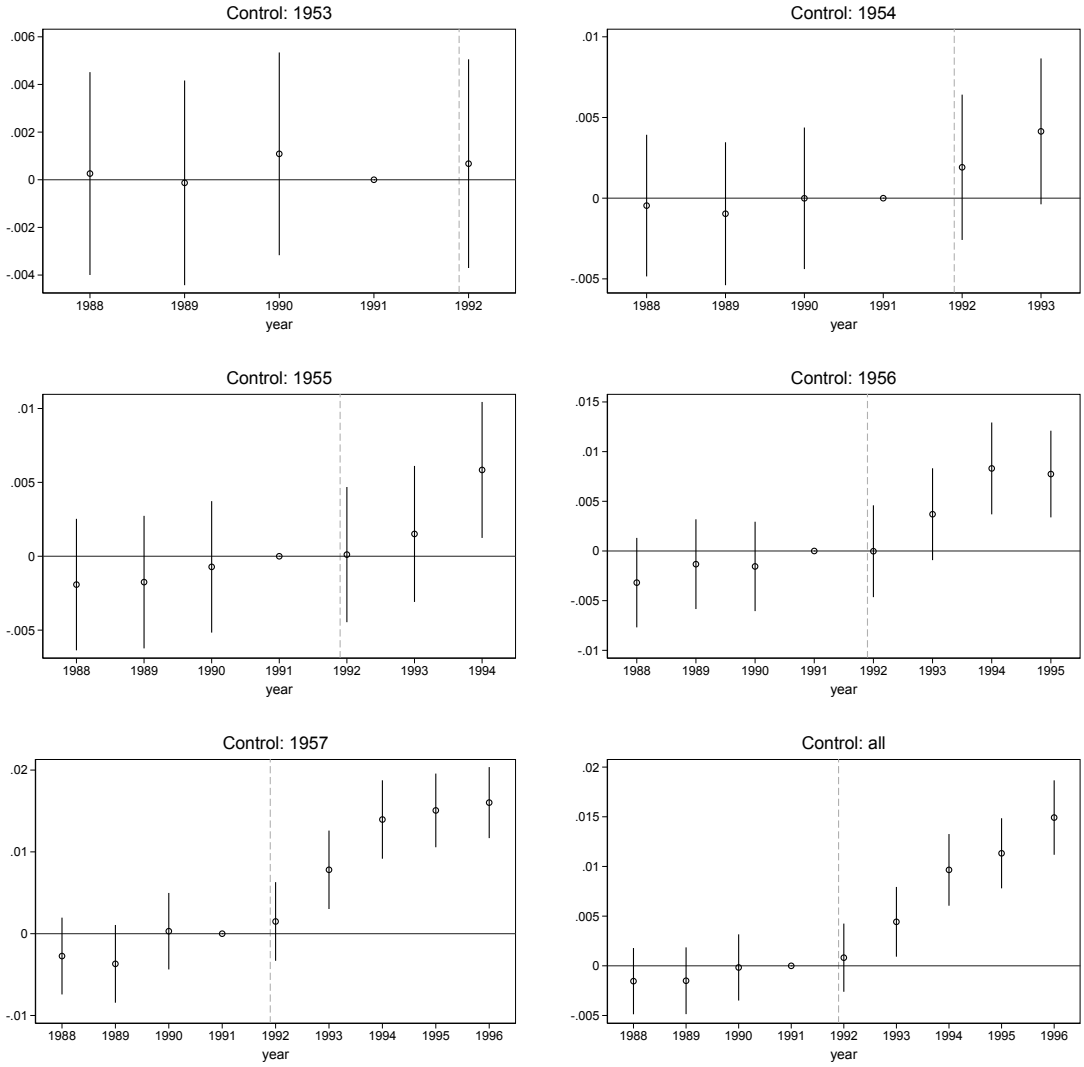

Note:

.pdfReform1992.doh:/workdata/704508/Data/figures/

Figure 8: Other outcomes: Estimation results

(a) Unemployment Intensity

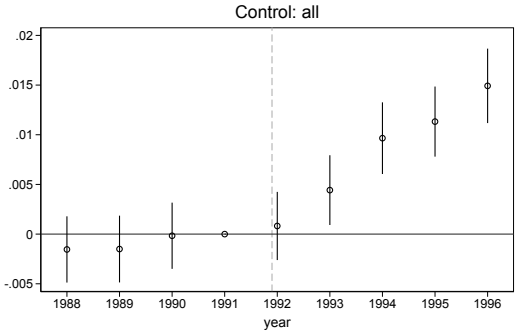

(b) Earnings in k 1979 kr

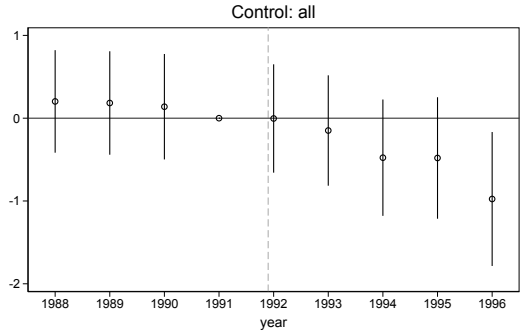

(c) Unemployment Incidence

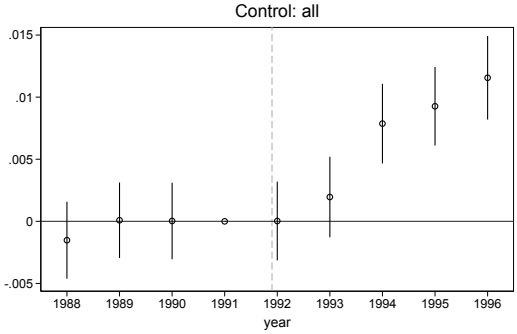

(d) non-Employment Incidence

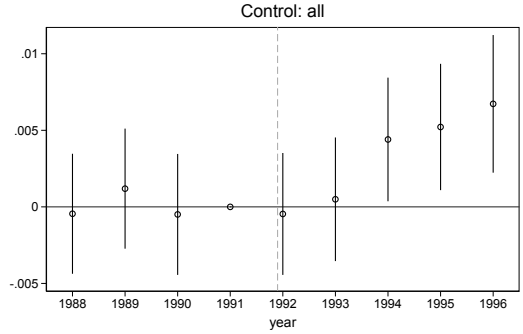

Note:

.pdfReform1992.doh:/workdata/704508/Data/figures/

Figure 9: Placebo Test: Effect of UI Reform on Never-Insured Individuals

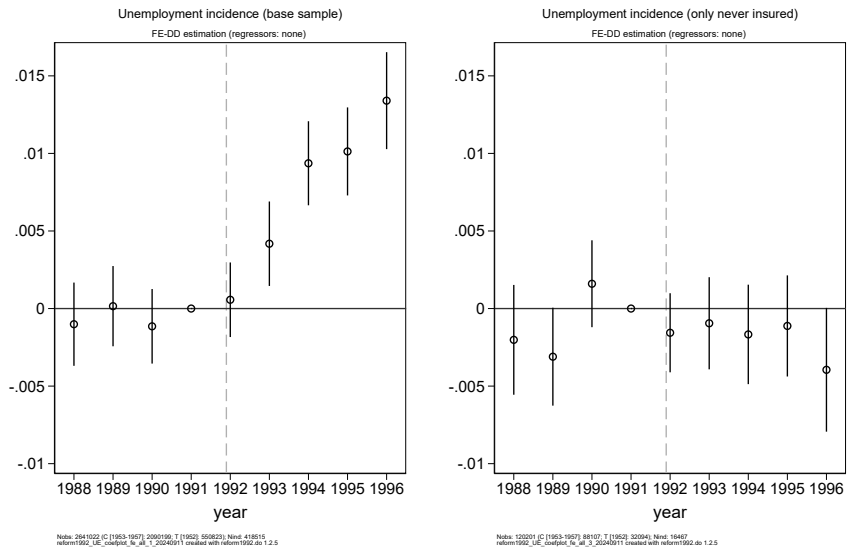

Note: This figure shows the effect of the 1992 UI reform on unemployment incidence for individuals who were never insured. The absence of significant effects in this sample supports the validity of the research design, suggesting that the observed effects in the full sample are not driven by unobserved confounding factors. The x-axis represents years, and the y-axis shows the estimated effect on unemployment incidence. Error bars represent 95% confidence intervals.

validity of our research design. It suggests that the significant effects we observe in our main analysis are indeed driven by the reform-induced insurance in UI enrollment, rather than by cohort-specific trends or other confounding factors.

Focusing on initially uninsured. Changing controls. IV estimates?

placebo among the never-insured including more treated cohorts.

placebo among a matched-control set.

5 Spillovers

While the direct effects of insurance on the insured have been well-documented, little has been done to understand the social contagion of these effects. Given that humans are

Figure 10: Main outcomes: FE Estimation results for those who were uninsured in 1988

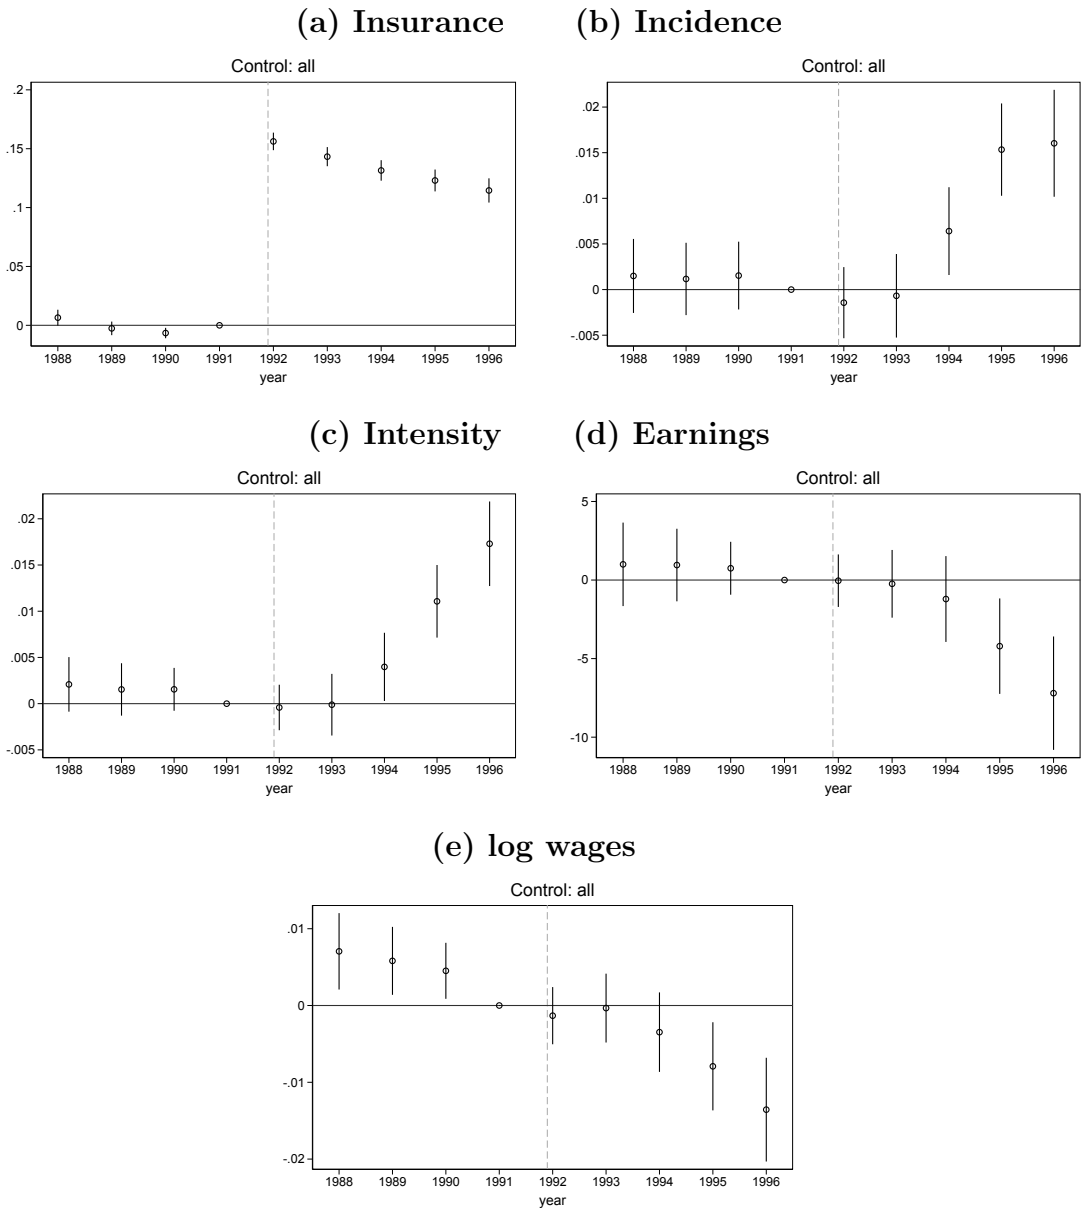

Note:

.pdfReform1992.doh:/workdata/704508/Data/figures/

inherently social creatures, heavily influenced by the norms and attitudes of their peers, it is plausible that insurance coverage could produce behavioral changes that extend beyond the insured individuals to their social networks. We examine this possibility in the context of our natural experiment.

5.1 Workplace Spillovers

We first examine spillovers through the workplace, a setting where individuals spend a significant portion of their waking hours and are likely to be influenced by their colleagues' behaviors and attitudes. We concentrate on workers who were already insured prior to the reform and thus cannot be directly affected by changes in their own insurance status. This allows us to isolate potential spillover effects from direct effects of insurance coverage.

We implement a difference-in-differences design similar to our main analysis, but with a crucial distinction: our sample consists entirely of always-insured workers. We categorize these workers based on their workplace's exposure to the reform. We define exposure as the share of a firm's employees who belong to the birth cohorts affected by the reform.

We find a striking pattern. In firms where the share of reform-affected colleagues is below average, there is no significant change in unemployment rates for already-insured workers at the reform. In firms with above-average exposure to the reform, however, already insured workers experience a 0.5 percentage point increase in their unemployment rate, despite no change in their personal insurance status. The fact that reform affects workers indirectly through peer exposure and not through insurance changes is evidence of spillover effects.

Figure 11 illustrates these contrasting outcomes. Panel (a) shows the stable unemployment trends for workers in low-exposure firms, while Panel (b) demonstrates the significant increase in unemployment for those in high-exposure firms. These results suggest that the behavioral changes induced by increased insurance coverage may propagate through workplace norms, affecting even those whose personal insurance status remains

unchanged. This spillover effect could be driven by several mechanisms, including norm changes, information sharing, and reduced peer pressure.

These findings have important implications for policy design and evaluation. They suggest that the total effect of insurance expansions may be larger than previously thought, extending beyond the directly affected individuals.

(a) Low workplace exposure

(b) High workplace exposure

Figure 11: Workplace Spillovers from Reform

Note: This figure shows spillover effects of the 1992 UI reform on unemployment incidence for always-insured workers. Panel (a) displays effects for workers in firms with below-average exposure to the reform, while panel (b) shows effects for those in firms with above-average exposure. Exposure is defined as the share of a firm's employees belonging to the reform-affected birth cohorts. The x-axis represents years, and the y-axis shows the change in unemployment incidence. Estimates are from a fixed effects difference-in-differences model. The treatment group consists of the 1952 birth cohort, while the control group includes the 1953-1957 birth cohorts. Vertical bars represent 95% confidence intervals. The results indicate no significant change in unemployment for workers in low-exposure firms, but a 0.5 percentage point increase for those in high-exposure firms, despite no change in their personal insurance status.

5.2 Marital Spillovers

Next we examine spillovers in the home through spousal relationships, another setting where individuals share significant time and mutual influence in behavior and attitudes. We again concentrate on workers who were already insured prior to the reform to isolate potential spillover effects from direct effects of insurance coverage.

Our analysis follows a similar difference-in-differences design as our workplace spillover study, but now focuses on the marital union. We compare the dynamics of two groups of always-insured workers: (1) Married workers whose spouses belong to the reform-affected birth cohorts, and (2) single workers, who serve as a comparison group not subject to spousal influence.

Figure 12 illustrates our findings. In Panel (a) we show the results for married workers with spouses in the affected cohorts. We observe a clear increase in unemployment incidence following the reform year. The effect appears to grow over time, reaching approximately 1.5 percentage points by 1996. This suggests a substantial spillover effect within marital pairs. Panel (b) displays the results for single, always-insured workers. In contrast to the married group, we see no significant change in unemployment incidence over the study period. The estimates fluctuate around zero, with confidence intervals consistently spanning zero. The mean of these points is negative.

These results add to the workplace evidence in showing strong spillovers in unemployment behavior. Always-insured individuals whose spouses were affected by the UI reform exhibit increased unemployment, despite no change in their own insurance status by construction. This effect is not observed among single workers, supporting the interpretation that the change is due to spousal influence rather than broader temporal trends.

6 Conclusion

This study provides new insights into the effects of unemployment insurance (UI) on labor market outcomes, leveraging a Danish reform that induced quasi-random enrollment in UI coverage. Our findings contribute to the literature on moral hazard in UI systems and extend our understanding of how insurance affects employed workers.

Our analysis reveals two key findings. First, insurance creates significant moral hazard among the employed, whereby induced coverage increases unemployment and reduces

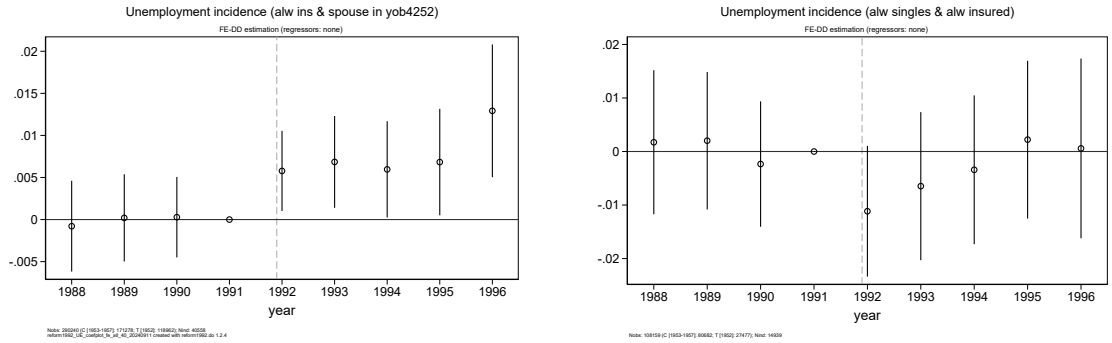

(a) Married to exposed spouse (b) Single, no exposed spouse

Figure 12: Marital Spillovers from Reform

Note: This figure shows the spillover effects of UI coverage on unemployment incidence for always-insured workers. Panel (a) displays effects for married workers whose spouses are in the affected birth cohorts (1952). Panel (b) shows effects for single workers as a comparison group. The x-axis represents years, and the y-axis shows the change in unemployment incidence. Estimates are from a fixed effects difference-in-differences model. The treatment group consists of the 1952 birth cohort, while the control group includes the 1953-1957 birth cohorts. Vertical bars represent 95% confidence intervals.

earnings. Second, insurance has substantial moral hazard spillovers. Namely, workers who were not directly affected by the reform but were exposed to peers who were see increased unemployment, but workers without exposed peers do not.

These results suggest that the total effect of insurance is larger than previously thought. It creates moral hazard among the employed, and the effects spread through social networks, even among workers whose insurance status is not affected. The results provide a more comprehensive cost-benefit analyses of insurance.

X

References

Ahammer, A., Fahn, M., and Stiftinger, F. (2023). Outside options and worker motivation. Technical report, CESifo Working Paper.

Brébion, C., Briole, S., and Khoury, L. (2022). Unemployment insurance eligibility and employment duration. In *The 34th EALE Conference 2022*.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Christofides, L. N. and McKenna, C. J. (1996). Unemployment insurance and job duration in canada. *Journal of Labor Economics*, 14(2):286–312.

Dahl, G. and Knepper, M. M. (2022). Unemployment insurance, starting salaries, and jobs. Technical report, National Bureau of Economic Research.

Ejrnæs, M. and Hochguertel, S. (2013). Is business failure due to lack of effort? Empirical evidence from a large administrative sample. *Economic Journal*, 123(571):791–830.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Gudgeon, M., Guzman, P., Schmieder, J. F., Trenkle, S., and Ye, H. (2024). When institutions interact: How the effects of unemployment insurance are shaped by retirement policies. Technical report, Mimeo.

Hartung, B., Jung, P., and Kuhn, M. (2022). What hides behind the German labor market miracle? Unemployment insurance reforms and labor market dynamics. *mimeo*.

Landais, C., Nekoei, A., Nilsson, P., Seim, D., and Spinnewijn, J. (2021). Risk-based selection in unemployment insurance: Evidence and implications. *American Economic Review*, 111(4):1315–1355.

Lusher, L., Schnorr, G. C., and Taylor, R. L. (2022). Unemployment insurance as a worker indiscipline device? evidence from scanner data. *American Economic Journal: Applied Economics*, 14(2):285–319.

Mitman, K. and Rabinovich, S. (2019). Do unemployment benefit extensions explain the emergence of jobless recoveries? Technical report, CEPR Discussion Paper No. DP13760.

Schmieder, J. F., Von Wachter, T., and Bender, S. (2012). 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.

Tuit, S. and van Ours, J. C. (2010). How changes in unemployment benefit duration affect the inflow into unemployment. *Economics Letters*, 109(2):105–107.

Van Doornik, B., Schoenherr, D., and Skrastins, J. (2023). Strategic formal layoffs: Unemployment insurance and informal labor markets. *American Economic Journal: Applied Economics*, 15(1):292–318.