Incentive Effects of Job Security Provisions

Clemens Noelke (corresponding author)

David E. Bell Postdoctoral Fellow

Center for Population and Development Studies

Harvard University | 9 Bow Street | Cambridge, MA 02138

Tel: (617) 496-8040 | Fax: (617) 495-5418 | Email: cnoelke@hsph.harvard.edu

Jason Beckfield

Professor of Sociology

Department of Sociology

Harvard University | 550 William James Hall | 33 Kirkland Street | Cambridge, MA 02138

Tel: (617) 496-3895 | Fax: (617) 495-5794 | Email: jbeckfie@wjh.harvard.edu

1 Figures, 5 Tables, 1 Appendix (for reviewers, online publication).

Word count: 7440

Acknowledgements: Previous versions have been presented at the 2010 Meeting of the American Sociological Association, at the Harvard Center for Population and Development Studies Seminar Series and the Swedish Institute for Social Research (SOFI). We thank participants at these events for valuable comments. All remaining errors are our own.

Abstract

This study explores the role of job security provisions as determinants of worker behaviour, specifically workplace absence and overtime work. We contrast two theoretical frameworks, both of which predict that stricter job security provisions result in lower overtime work and higher absence rates. The empirical analysis implements a quasi-experimental research design based on a British job security provisions reform that extended dismissal protection among a representative sample of low-tenure workers. Using a Difference-in-Differences approach, we find that overall absences among those affected by reform increased by 9% and absences due to sickness increased by 35% relative to pre-reform levels. Moreover, we find evidence of increased absences that respondents classify as "other leave/holiday" and a reduction in overtime work.

Keywords: labor market institutions, labor law, labor supply, flexibility, United Kingdom, economic growth

JEL classification: J22 Labor Supply, J24 Labor Productivity, J42 Monopsony, J65 Severance Pay, K31 Labor Law

Introduction

Job security provisions limit the ability of employers to dismiss unproductive workers and adjust staffing levels to current market conditions. Because they restrict employer flexibility, job security provisions and employment protection legislation more generally (including regulations on temporary work and collective dismissal) have been identified as central causes of macroeconomic dysfunction, in particular poor job growth, in OECD countries (Blanchard and Katz 1997; Bradley and Stephens 2007; OECD 1994, 2004; Pontusson 2005; Scharpf and Schmidt 2001; Siebert 1997). This view continues to motivate deregulatory reforms in OECD countries, even though its theoretical and empirical foundations have always been contested (Avdagic and Salardi forthcoming; Baccaro and Rei 2007; Freeman 2005; Glyn et al. 2003; Nickell and Layard 1999; Nickell 1997).

Because research is still split on whether job security provisions are indeed an important cause of macroeconomic dysfunction, this study seeks to make a theoretical and empirical contribution to further our understanding of the labour market effects of job security provisions. First, while previous research has focused on demand side effects, for example on the overall level of employment or hiring and firing dynamics, still little is known about supply side effects of job security provisions. The following study asks whether job security provisions affect worker behaviour, in particular patterns of workplace absence and overtime work. Second, we implement a quasi-experimental research design, exploiting a British reform that selectively increased job protection among low tenure employees. Below, we discuss advantages of this approach relative to pooled cross-sectional time-series analysis, on which earlier research on labour market effects of job security provisions is based.

Prior work on supply side effects has suggested that, by reducing the likelihood of job loss, job security provisions create incentives for workers to invest into specific skills (Estevez-Abe, Iversen, and Soskice 2001) or undertake risky but potentially innovative projects (Acharya, Baghai, and Subramanian 2010). However, workers might also exploit the bargaining power that job security provisions confer by lowering their work effort or working less than their contractually required hours, i.e. being absent more often and working less overtime, which represents the current consensus view among economists (Ichino and Riphahn 2004, 2005; Lindbeck, Palme, and Persson 2006). Below, we outline an alternative theoretical framework that emphasizes labour market frictions. These frictions may induce employees to supply more labour than they would prefer, i.e. work overtime more often and be absent from work less frequently even for legitimate causes, such as sickness. In this framework, protection against dismissal reduces inequitable and possibly inefficient overemployment.

The empirical analysis exploits a British job security provisions reform that selectively increased dismissal protection among recently hired employees (see Marinescu 2009) to test whether job security provisions are a cause of workplace absence and overtime work. A reform enacted in June 1999 reduced the tenure threshold at which workers could claim Unfair Dismissal from 24 to 12 months. Dismissal protection for individuals with 12-23 months of tenure increased as a consequence. Workers with 24-35 months of tenure, for whom dismissal protection was unchanged, constitute the control group. Using data from the British Labour Force Survey on treated and control groups before and after reform, we implement a Differences-in-Differences estimator. We use (coarsened) exact matching and non-linear treatment group specific trends to test for the sensitivity of our results to time-varying observed and unobserved confounders.

Theory and Hypotheses

The current mainstream view in economics is that job security provisions lower work effort from otherwise optimal levels (Ichino and Riphahn 2005). Employees take advantage of the bargaining power that job security provisions confer by working less, up to the point that the loss in output equals the costs of dismissal. By reducing the probability of dismissal, stricter job security provisions result in higher absence rates and lower rates of overtime work. This perspective is consistent with research in economics indicating that greater job security, due to job security provisions reforms (Lindbeck et al. 2006; Olsson 2009), crossing a tenure threshold (Ichino and Riphahn 2004), or upturns in labour demand (Arai and Thoursie 2005; Askildsen, Bratberg, and Nilsen 2005; Bratberg and Monstad 2012), is associated with higher absence levels.

However, this evidence is equally consistent with the following mechanisms based on labour market frictions and information problems. When changing jobs following dismissal or quits, individuals are faced with substantial economic and social costs, including periods of diminished or no income, lost access to job-related social networks, loss of social status, reduced health, happiness and well-being, as well as mobility costs (Brand 2006; Gangl 2006; Jahoda, Lazarsfeld, and Zeisel 1975; Sullivan and von Wachter 2009; Winkelmann and Winkelmann 1998; Young 2012).[1] Employees therefore do not immediately quit, if employers ask them to put in extra hours without adequate compensation, deny a vacation request, or request that employers take vacation rather than sick days. Because changing employers is costly, employers have some degree of monopsony power that they can exploit to demand more hours from employees (Manning 2003; Stewart and Swaffield 1997).[2]

Consistent with this argument, survey data indicate that across 15 European countries around half of dependent employees would like to reduce their work hours, by on average 3.7 hours (Bielinski, Bosch, and Wagner 2001). For both men and women, actual hours exceed desired hours, with the UK exhibiting the largest discrepancy among men (Bielinski et al. 2001). Stewart and Swaffield (1997) show that over a third of British manual workers would work fewer hours at the prevailing hourly wage (see also, Manning 2003:228), which is attributed to lack of job security and lack of alternative job opportunities. Golden and Gebreselassie (2007) similarly report substantial overemployment among U.S. workers. In this framework, job security provisions redress the balance of power between employer and employees and allow employees to resist employer demands for extra hours and adjust hourly labour supply to individually optimal levels. Employees should be more likely to take sick leave if they are sick rather than going to work sick (i.e. reduce levels of "presenteeism"), take the vacation days they are entitled to in order to recover, or deny requests for overtime.

Job security may also diminish oversupply of labour that results from employer information problems. Employers are limited in their ability to directly measure employee performance (Sorensen and Kalleberg 1981) and therefore rely on easily observable signals of productivity such as hours worked. Hours worked are commonly treated as a direct signal of productivity or a signal of productivity-related characteristics, such as work ethic. Employees therefore have a strong incentive to signal high productivity by being punctual or working overtime and avoid signals associated with low productivity (absenteeism, rejecting requests for overtime), since these may be grounds for dismissal. In consequence, they supply more labour than would be individually optimal, including overtime, not taking vacation days, and coming to work sick. Again, job security provisions, by lowering the risk of dismissal, reduces oversupply of work hours to individually optimal levels, resulting in higher absence rates and a lower incidence of overtime work.

Job Security Provisions in the UK and the 1999 Reform

Compared to other Western European countries such as Germany, France, Sweden or Italy, dismissing a permanent employee is not very costly for employers in the UK (OECD 2004). However, dismissals are considerably more regulated compared to the employment-at-will doctrine governing employment relationships in the U.S.[3] Apart from severance pay and notification requirements, the right to claim "Unfair Dismissal" in the UK entitles employees in the event of dismissal to obtain a written statement from employers detailing the cause for dismissal, and to contest dismissal in front of an Employment Tribunal. For dismissals to be fair, they have to have a just cause. In the UK such causes include employee capability and qualification to perform the job, employee misconduct, or economic necessity (1996 Employment Rights Act). When evaluating a claim, Tribunals have to take into account the specifics of each case and reach a verdict that is equitable for both parties. On average, 40-60 thousand claims are filed each year (Ministry of Justice, *Employment Tribunal and EAT Statistics*, various issues).

The June 1999 reform lowered the eligibility threshold to claim Unfair Dismissal from 24 months to 12 months of tenure. This creates a "natural experiment", with a control group and a treatment group observed before and after an intervention. The treated are those with 12-23 months of tenure, who became eligible to claim Unfair Dismissal on June 1, 1999. Workers with 24 or more months of tenure are unaffected by reform. To keep the control group homogenous yet sizeable enough, we restrict it to workers with 24-35 months of tenure.

The reform effectively increases job security among the treated by imposing non-negligible economic costs in the event of dismissal. While expected monetary costs are small, administrative costs accrue with certainty. Employers are required to keep documentation in order to justify dismissal as fair and to engage in judicial proceedings with Employment Tribunals or try to settle claims before they go to trial. Consistent with reform increasing job security among the treated, Marinescu (2009) shows that the firing hazard of workers with 12-23 months of tenure declined substantially after reform.

Methodology

Prior research on the effects of job security provisions in comparative political economy has relied on pooled cross-sectional time-series analysis (for example, Avdagic and Salardi forthcoming; Bradley and Stephens 2007; OECD 2004). In the following, we adopt a quasi-experimental approach, exploiting variation in job security provisions caused by a British reform, and implement a Differences-in-Differences estimator (see Angrist and Pischke 2009, for an overview) using micro data from the British Labour Force Survey. Since we rely on data from a single source, concerns about harmonization and comparability that arise when pooling data from different countries are effectively eliminated. Second, the use of micro data gives us considerable flexibility in terms of defining analysis sample and dependent variables, and allows us to conduct effect heterogeneity analyses across demographic subgroups. Third, we can use individual level covariates to adjust for compositional changes over time. Fourth, we do not rely on an indicator of job security provisions that brings with it its own set of problems: The widely used OECD indicator of job security provisions set aggregates different policies and regulations (severance pay, notification requirement, advance notice periods, etc.) into one indicator. On the one hand, this involves some arbitrary judgments (see Venn 2009, for discussion). On the other hand, it is not clear how policy makers should interpret the resulting estimates. Is severance pay generating the effect, or notification requirements, or their interaction? Rather than abstract movements on a synthetic indicator, we focus on the effect of a concrete and politically feasible reform.

There are additional, important considerations that make the quasi-experimental Differences-in-Differences approach attractive: It addresses a key problem of casual inference, omitted variable bias, in a transparent and effective way. First, the approach is conceptually rooted in the counterfactual framework of causality (Gangl 2010; Heckman, LaLonde, and Smith 1999; King, Keohane, and Verba 1994), which allows us to clearly define counterfactual states of interest and transparently state identifying assumptions. Second, the plausibility of the key identifying assumption of the differences-in-differences estimator ("parallel trends", see below) can be tested, and we can draw on established remedies should we find it to be suspect. Related to this point, given the large number of observations in our data, we can fit many more parameters to control for observed and unobserved confounders than is possible in pooled cross-sectional time series analysis; and we can enter covariates as dummy variables, i.e. we do not need to make any assumptions about their functional form. Finally, problems induced by unobserved confounders are most likely less severe if treated and controls groups are taken from the same context and only differ by 12 months of tenure than if treated and control groups were taken from different countries.

While studies using pooled cross-sectional time-series analysis conduct a direct test of whether job security provisions explain variations in labour market performance across OECD countries, our design yields a historically and geographically contingent effect estimate. While this effect may or may not be generalizable, it still encodes useful, policy-relevant information for the country concerned. Moreover, the quasi-experimental approaches maps out an alternative path to testing general theories about the effects of institutions. If similar analyses are repeated across different countries and time periods, we obtain different estimates of the effects of changes in job security provisions under varying conditions. In a second step, conceptually akin to multilevel analysis, one can then explain what accounts for variation of treatment effects across contexts and ultimately test a general model explaining the heterogeneous effects of job security provisions across countries (Gangl 2010:41).

Data

We use data from the British Labour Force Survey (BLFS), a representative survey of British households living at private addresses in the UK. Designed as a rotating panel, households are surveyed for five consecutive quarters. The dataset used in the main analyses is restricted to observations between June 1996 and May 2002, i.e. 36 months before and 36 months after the June 1999 reform (including the month of reform). The sample is further restricted to dependent employees born in the UK aged 25-54 with 12-35 months of tenure working 16+ hours per week. The right to claim unfair dismissal does not apply to self-employed or family workers, who are therefore excluded from the main analyses. Employees on non-permanent contracts are also excluded from the analysis.[4] 29% of observations are based on proxy interviews, e.g. by spouses. We include observations from all five waves.[5] The analysis includes all observations with non-missing values on the dependent variables and following covariates: age, gender, education level, region, month, proxy-response, and wave (see Table A1, for sample means).

The dependent variables are constructed from variables measuring labour supply during the "reference week", i.e. the week preceding the interview. Respondents are first asked about usual hours worked in their main job, followed by consistency checks. They are then asked about hours actually worked, again followed by consistency checks. If respondents worked less than their usual hours, they are asked to state a reason for the discrepancy. 17% of employees have worked less than their usual hours because their working hours vary. We count these individuals as not having been absent from their job. Another 23% of employees report having worked less than their usual hours for different reasons: 11% state that they have been absent due to having taken "other leave/holiday", 5% have been absent due to a bank holiday, and 4% have been absent due to "sickness or injury". The remaining 3% are distributed over nine other causes of absence.[6]

From these variables, we derive the following dependent variables: Total absence, i.e. missing some or all of one's hours for any reason besides working on a flexible work schedule (=1, 0=everyone else), missing some/all hours due to sickness (=1, 0=everyone else), missing some/all hours due to vacation (=1, 0=everyone else) and missing any amount of time due to other reasons (bank holiday, "other"…) (=1, 0=everyone else). In supplementary analyses, we differentiate, for each cause of absence, between either missing some but not all of one's usual hours or missing all of one's usual hours. Finally, overtime work is measured by a dummy variable indicating whether or not a respondent has worked more than his/her usual hours during the reference week.

Identification and Estimation

To estimate the effect of the June 1999 reform on workplace absence and overtime work, we adopt a Differences-in-Differences approach (Angrist and Pischke 2009; Gangl 2010; Imbens and Wooldridge 2009) and estimate variants of the following linear model:

$$(1)\ Y_{it} = \alpha + d_t + \gamma D_i + \ \partial D_i PST_t + \boldsymbol{X}\beta + \varepsilon_{it}$$

α is a constant, d_t are calendar month fixed effects, D_i is a dummy variable indicating 12-23 months of tenure, PST_t is an indicator variable for post-reform observations (i.e. date >= June 1999). ∂ estimates the coefficient of an interaction between treatment and post-reform indicator, and can be interpreted as the average treatment effect on the treated (ATT). D_i controls for time-constant, treatment-group specific unobserved factors, while the d_t control for unobserved time-varying shocks to the outcome variable that are common to both treated and controls. $\boldsymbol{X}$ is a matrix of flexibly specified covariates that are assumed to be uncorrelated with ε_{it} conditional on the other included variables and β is a vector containing their regression coefficients. The covariates are dummy variables for age (grouped into five-year intervals), gender, education level (dichotomized: college vs. no college), region, month of survey, year, proxy vs. self-response, and wave (Table A1 in the Appendix contains information on the sample means of all variables except for calendar month dummies). For ∂ to estimate the causal effect of reform, the key identifying assumption is that the (counterfactual) change in outcomes in the treatment group that would have occurred in the absence of reform is identical to the change in outcomes that is observed in the control group.

Equation 1 is a restricted version of the following event-study model (Cascio 2009), which substitutes in place of $D_i PST_t$ in equation 1 a set of k indicator variables D_{it}^k that are equal to 1 for the treated ($D_i = 1$) in year k and zero otherwise. k denotes the year relative to the year prior to reform, for which $k = 0$:

$$(2)\ Y_{it} = \alpha + d_t + \gamma D_i + \sum_{k=-2}^{k=3} \theta^k D_{it}^k + \boldsymbol{X}\beta + \varepsilon_{it},$$

To identify this model, we omit D_{it}^0, the indicator pertaining to the year prior to reform. The parameters θ^kestimate the expected difference in outcomes between treated and controls relative to the year prior to reform.[7] Compared to equation 2, equation 1 constrains the pre-reform differences to be zero ($\theta^k = 0\ for\ all\ k < 0$) and the post-reform differences to be constant ($\theta^k = \partial\ for\ all\ k \geq 1$).[8]

Equation 2 allows us to study pre- and post-reform effect dynamics. If post-reform effect dynamics are found to be minor or not deemed relevant, constraining $\theta^k = \partial\ for\ all\ k \geq 1$ is defensible for efficiency gains. However, if we observe a trend in the pre-reform θ^k, this can be taken as evidence that control group and treatment group differ in terms of underlying trends prior to reform and the estimate of ∂ in equation 1 will be biased without further adjustments.

We use two strategies independently and in combination to address potential biases. First, we enrich equation 1 with linear time trend variables:

$$(3)\ Y_{it} = \alpha + d_t + \gamma D_i + \partial D_i PST_t + \pi_1(l_\mathrm{t} - l_0)D_i + \pi_2(l_\mathrm{t} - l_0)D_i PST_t + \boldsymbol{X}\beta + \varepsilon_{it}$$

While the d_t control non-parametrically for time shocks, such as changes in labour demand affecting both groups equally, (treated) low tenure workers may be more sensitive in their absence behaviour to changes in labour demand. π_1 adjusts for linear departures of the change over time among the treated from these common time shocks. π_2 allows this trend to break after reform. The trend variables (defined at the calendar month level) are centred to equal 0 in the month prior to reform. We will also add higher order polynomial terms of each trend variable as further specification checks. The resulting specifications adjust flexibly for time-varying, treatment-group specific unobserved factors that could generate non-parallel, non-linear trends before and after reform.

Second, non-parallel trends across the two groups may stem from changes in the distribution of observed covariates (Abadie 2005). For example, if the sample distribution of an observed confounder, such as gender, changes differentially over time across treated and controls, this may result in non-parallel trends. To illustrate this issue, Table A1 list the sample means for key covariates that were used in the matching analysis (see following paragraphs). Column 4 shows the difference in sample means between treated and controls, which appear small: The average absolute difference in means across all covariates was 0.3 percentage points (Table A1) with a maximum absolute difference of 1.5 percentage points (education). Furthermore, column 5 reports an estimate of ∂ from equation 1 using the covariates as dependent variables.[9] While for most covariates, we observe no differential changes after reform, we obtain significant coefficients in a number of instances. If the covariates, whose distribution changes differentially over time, are also (independently or interactively) correlated with the outcome variable, our resulting estimate of the reform effect will be biased.

To address this issue, we adopt Coarsened Exact Matching (CEM), a statistical matching technique, to reduce imbalances in the distribution of covariates between treatment and control groups in terms of covariate means and, importantly, in terms of all nonlinearities/interactions (Blackwell et al. 2009; Iacus, King, and Porro 2011). CEM sorts the data into strata defined by unique values of coarsened covariates and then deletes all strata that contain either treated or control observations only. Coarsening the covariates is typically necessary to retain a sufficient number of observations within strata. We match on the same set of covariates $\boldsymbol{X}$ included in equation 1 (dummies for age, gender, college degree, region, month of survey, proxy response) as well as calendar time, coarsening the 36 calendar month into six year dummies.

Table 1 indicates that 17527 strata of the 95980 strata defined by unique combinations of the coarsened covariate values contain both treated and control observations. Dropping the strata (observations in strata) without both control and treated observations reduces sample size by 64%. The resulting dataset is perfectly balanced in terms of observed covariates. For each covariate, matching reduces the difference in means between controls and treated to zero (Table A1). Moreover, after matching the global imbalance statistic L_1 is equal to zero, indicating that the multivariate empirical distributions of the covariates in treatment and control group exactly coincide in the matched dataset (Table A1). After restricting the data to the matched observations, we estimate equations 1, 2, and 3, omitting $\boldsymbol{X\beta}$. To obtain the ATT for the full (original) sample, we use weights supplied by the CEM software.[10]

Table 1. Results from Coarsened Exact Matching.

Number of strata		95980
Number of matched strata		17527
Number of Observations	Control	Treated
All	62743	81448
Matched	24364	27103
Unmatched	38379	54345
Multivariate L_1 distance		
before matching		0.7297
after matching		0.0000

Source: British Labour Force Survey, own calculations.

To illustrate the effectiveness of the matching strategy, Table A1 (column 9) reports estimates of ∂ analogous to those in column 5. For all covariates, the resulting estimate is equal to zero. Figure A2 illustrates for two covariates that the differences in covariate means for treated and controls changed somewhat in each year, changes that could potentially bias our causal effect of interest . However, in the matched sample these differences are zero in every year. The same holds true for all other covariates.

Finally, ∂ may be confounded by unobserved discrete changes around reform that are specific to either treatment or control group, such as a reform enacted at the same time, or in the year preceding or following reform. Marinescu (2009) discusses labour market reforms implemented around June 1999, including a minimum wage reform in April 1999 and reforms of parental and dependent care leave and sex discrimination regulations. However, with the potential exception of the minimum wage reform, there is no reason to expect that these reforms have differentially affected labour supply among the treated and controls. Marinescu's (2009) as well as our specification checks (see below) indicate that these reforms had no differential effect on the treated.

Results

Table 2 shows descriptive statistics for treated and control group before and after reform. We tend to observe lower absence rates for the treated (12-23 months of tenure) compared to controls (24-35 months of tenure). The last column reports a Differences-in-Differences estimate calculated as the pre-post difference in means for treated minus the pre-post difference in means for controls. These estimates suggests that after reform the difference in outcomes between treated and controls has increased by 0.9 percentage points. Most of this increase appears to be due to higher rates of vacation taking. Regarding overtime work, we observe higher rates among the treated before and lower rates after compared to controls, suggesting that reform has diminished the incidence of overtime work by 0.8 percentage points.

Table 2. Sample Means for Dependent Variable by Treatment Group and Reform Status

	Before Reform (June 1996 – May 1999)		After Reform (June 1999 – May 2002)		$\partial = (T_1 - T_0) - (C_1 - C_0)$
	Treated (T_0)	Control (C_0)	Treated (T_1)	Control (C_1)	
Total absence	21.9	23.4	23.3	24.0	0.9
Sickness absence	4.0	4.4	4.2	4.2	0.3
Other leave/holiday	9.9	11.0	10.9	11.1	0.9
"Other absences"	8.0	8.0	8.2	8.6	-0.3
Overtime	13.8	13.4	12.0	12.4	-0.8

Note: Based on full analysis sample. "Other absences" include: bank holidays, maternity or paternity leave, training course away from workplace, started new job or changed jobs, ended job or did not start new one, laid off or work interruption due to bad weather, laid off or work interruption due to labour disputes, laid off or work interruption due to economic or other reasons, other personal or family reasons, other reasons. Source: British Labour Force Survey, own calculations.

Figure 1 displays estimates of θ^k obtained from Equation 2 with 90% confidence intervals both the full and the matched sample. The estimates can be interpreted as the expected change in the difference in outcomes relative to the year preceding reform. Table 2 indicates, on average, lower absence rates among the treated prior to reform. Figure 1 (top left) indicates that this difference became more negative in the years leading up to reform. This pre-reform trend breaks in the year of reform, which is consistent with a causal effect of reform. The pattern in the matched data is also consistent with this interpretation.

Table 3 reports the corresponding estimates of the average treatment effect of reform on the treated. For the full and matched sample, we report estimates of ∂ from the baseline equation (equation 1, columns 1 and 5, "Baseline") and from equation 3 that adds to the baseline specification a linear trend for the treated as well as a linear trend for the treated for the post-reform period only (columns 2 and 6, "Linear Spline"). Estimates in columns titled "Quadratic Spline" ("Cubic Spline") are further adjusted for 2^{nd} (2^{nd} and 3^{rd}) order polynomial terms of both trend variables.

For overall absence rates, the estimates in Table 3 (first row) indicate a statistically significant reform effect of 1.0 percentage points which is almost identical to the DD estimate in Table 2 (row 1, column 1). Once we adjust for pre- and post-reform linear trends specific to the treated (column 2), the effect estimate increases to 1.9 percentage points. This increase is consistent with the sharp discontinuity observed around reform in Figure 1 (top left). It appears that the differential pre-reform trend resulted in a downward bias of the treatment effect. Further adjusting for non-linear trends does not change coefficient estimates substantially but causes confidence intervals to widen. Results from the matched sample are consistent with results from the full sample.

Figure 1. Changes in Outcome Differences Between Treated and Controls Relative to Year Preceding Reform.

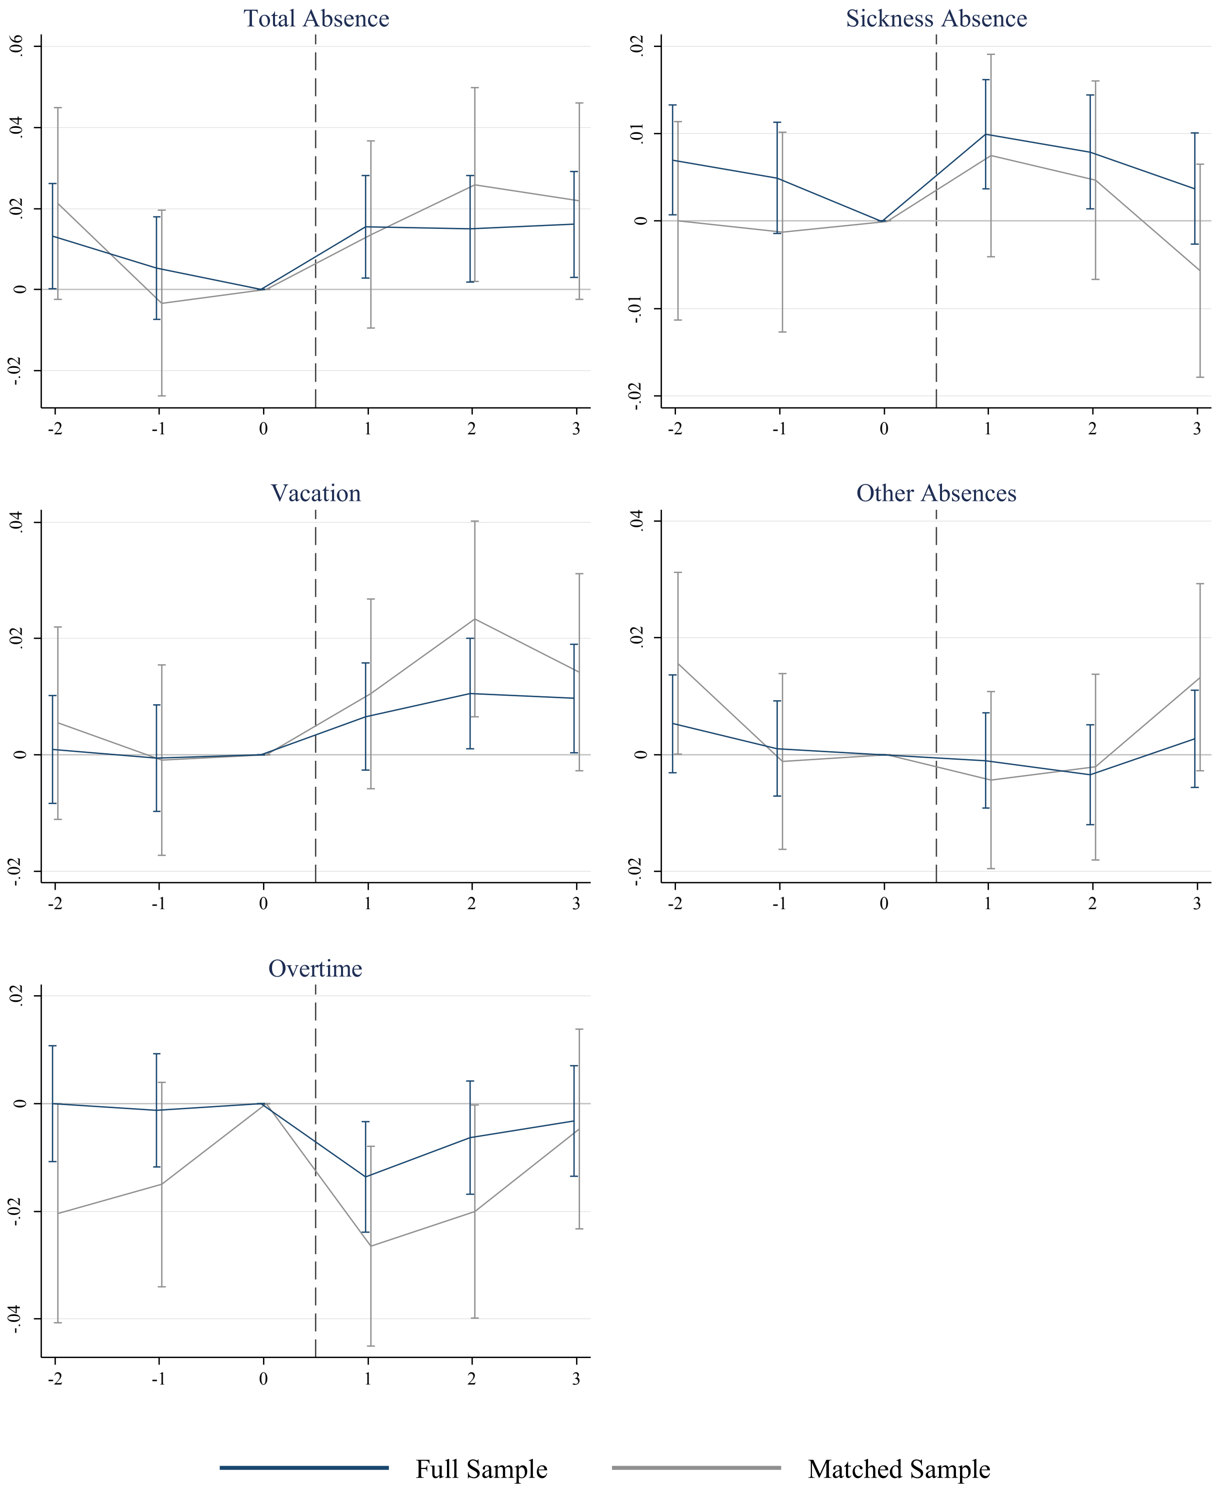

Notes: The figures display estimates of θ^k from Equation 2 with 90% confidence intervals from both the full and matched sample. Reform is implemented at the beginning of year k=1 (June 1999). Source: British Labour Force Survey, own calculations.

For sickness absence, we observe a similar pre-reform trend pattern (Figure 1, top right) as in the case of overall absence rates. This trend breaks in the year reform is implemented, indicating a positive effect of reform in years $k = 1$ and $k = 2$. Effects weaken in the course of the post-reform observation period, and turn negative in the matched sample in the third year after reform. Once differences in relative (linear) trends are adjusted for, we obtain an effect estimate of 1.4 percentage points (row 2, column 2). Adjusting for higher-order trends or restricting the sample to matched observations does not change this conclusion.[11]

Regarding vacation-taking, Figure 1 (middle left) indicates that the difference in outcomes between treated and controls does not change prior to reform, but increases in the year reform is implemented, which is consistent with a causal effect of reform. The baseline estimate in Table 3 indicates that that reform increased the incidence of vacation-taking by 0.9 percentage points (or 9% relative to average pre-reform levels among the treated). Adjusting for trends in this instance is over-controlling, because there is no indication of divergent pre-reform trends. Effect sizes increase slightly if we restrict the sample to matched observations. Hausman tests indicate that none of the other estimates differs significantly from the baseline estimate.

Figure 1 (middle right) also charts trends for absences due to other reasons, mainly absences due to bank holidays as well as numerous other factors (see Footnote 5). We observe a trend in the difference in outcomes starting in year $k = -2$ and lasting until year $k = 2$ and no indication of an immediate reform effect. The estimated reform effects (Table 3, row 4) are small, negative in tendency, and never reach statistical significance.