Received: 4 September 2023 | Accepted: 29 January 2024

DOI: 10.1111/ijsw.12651

ORIGINAL ARTICLE

INTERNATIONAL JOURNAL OF SOCIAL WELFARE

Capping welfare payments for workless families increases employment and economic inactivity: Evidence from the UK's benefit cap

Aaron Reeves[1,2] | **Mark Fransham**[1,2] | **Kitty Stewart**[3] | **Mary Reader**[4] | **Ruth Patrick**[5]

[1]Department of Social Policy and Intervention, University of Oxford, Oxford, UK

[2]International Inequalities Institute, LSE, London, UK

[3]Centre for Analysis of Social Exclusion, LSE, London, UK

[4]Department of Economics, Stanford University, Palo Alto, California, USA

[5]School for Business and Society, University of York, York, UK

Correspondence
Aaron Reeves, Department of Social Policy and Intervention, University of Oxford, Oxford, UK.
Email: aaron.reeves@spi.ox.ac.uk

Funding information
Nuffield Foundation, Grant/Award Number: FR-23208

Abstract

In this paper, we examine the labour market effects of lowering the UK's benefit cap in 2016. This policy limits the total amount a working-age non-disabled household with no-one in employment can receive in social security. We treat the sharp reduction in this benefit cap as a natural experiment, comparing those at risk of being capped and those who were not before and after the cap was lowered. Drawing on data from ~500,000 individuals, we find that this reform reduced unemployment compared to those not at risk of being capped. The reform also increased economic inactivity, partly because the cap harmed mental health but also because those at risk of being capped were eligible to claim disability-related welfare payments that made them exempt. Limiting total monthly welfare payments of low-income families may increase employment for some but it can also push others out of the labour market altogether.

KEYWORDS

disability, economic inactivity, mental health, welfare reform

INTRODUCTION

The United Kingdom (UK) is one of a very small number of countries that superimposes an upper limit on the total amount of financial support that a household can receive from all forms of social security, leading to reductions in monthly payments to claimants. This 'benefit cap' has been in place since 2013 and is designed to incentivise low-income families to seek employment. The benefit cap appears to have been somewhat successful, with early estimates suggesting the cap increased employment among those subject to the policy (Kaur et al., 2014; Tonutti, 2018). More recently, a government-commissioned evaluation of the effects of lowering the cap in 2016 (thereby reducing the amount of money claimants received) has complicated the picture. This analysis found some increases in employment but also observed a large number of people who remained subject to the cap because they had not entered employment (Griggs et al., 2023).

There are two ways in which these studies do not provide a full picture of the employment effects of the cap. First, they largely focus on the transition from unemployment to employment and paid limited attention to

Abbreviations: DWP, Department for work and pensions; LFS, Labour force survey; OECD, Organisation for economic co-operation and development; PIP, Personal indepedence payments; UK, United Kingdom; WTC, Working tax credits.

whether the benefit cap pushed some into economic inactivity. This matters because families can become exempt from the cap if they claim some form of disability-related social security, some of which remove the requirement to be available for paid work. Many of those affected by the benefit cap already have poor health and so capped families may have been incentivised to claim disability-related support to which they were already entitled but had not applied for, potentially becoming economically inactive as part of their claim (Geiger, 2017). In addition, being affected by the cap has a negative effect on mental health (Reeves et al., 2022), which may also push people away from the labour market (Frijters et al., 2014; García-Gómez et al., 2010) and have made them newly eligible for disability-related support (over 50% of Employment Support Allowance claimants had 'mental and behavioural disorders') (Mind UK, 2023).

Second, existing work on the reduction of the cap in 2016 was conducted with a relatively small, longitudinal survey collected after the reform was implemented and which unfortunately but understandably had high rates of attrition (Griggs et al., 2023). The earlier report provides helpful insights but is unable to provide a clear picture of what happened before and after the cap was lowered, among which groups, and how this varied across different regions. This is important because, before the 2016 changes, the cap largely affected people in London, which was also the region where the cap seemed to have the largest impact on employment (Kaur et al., 2014). After 2016, the cap affected more people outside of London—in places that were less affluent and economically productive (Leyshon, 2021)—and this may have altered the employment effects of the policy.

Part of the reason the benefit cap is important is because by design it disproportionately affects families who have substantial caring commitments. The majority of those affected have three or more children, many of whom are headed by single parents (Department for Work and Pensions [DWP], 2022). These households may be less able to respond to work incentives than other groups. One consequence of this could be to incentivise people to become economically inactive whilst remaining attached to the social security system, especially if those parents have poor health.

In this paper, we take up these questions by examining the labour market effects of lowering the cap in November 2016. Unlike earlier work, we examine these issues by treating a change in the level of the benefit cap in late 2016 as a natural experiment (Dunning, 2012). We draw on a large, repeated cross-sectional survey to identify those at risk of being subject to the cap and those who are not. We then follow these groups over time. Using a variety of causal identification strategies (including difference-in-differences models and interrupted time series analysis) we show that lowering the level of the cap (thereby increasing the number of people who were at risk of being affected, as well as the size of the income loss for those already affected) had a dual impact on those affected. While some individuals responded to being capped by returning to work, there was a larger group of individuals who responded by becoming economically inactive. In the second stage of our analysis, we unpack who is becoming economically inactive, finding that there is a rise in the proportion of people claiming disability-related social security and that the mental health effects of the cap are likely pushing more people to become economically inactive post-reform. Although we adopt a different methodology, our findings are consistent with but also extend earlier work (Griggs et al., 2023). The dual impact of the benefit cap that we uncover has implications for how governments design policies intended to incentivise employment, especially because they may inadvertently exacerbate health and labour market inequalities between those already in work and those who are currently economically marginalised.

THE BENEFIT CAP AND WELFARE REFORM IN THE UK

The benefit cap was announced as part of a broader set of reforms introduced by the 2010 Conservative-led Coalition Government to reduce spending while increasing 'fairness' in the benefit system (Osborne, 2010). The explicit justification, set out by Chancellor of the Exchequer George Osborne in announcing the policy in October 2010, was to ensure that 'no family should get more from living on benefits than the average family gets from going out to work' (Osborne, 2010). The cap for couples was therefore set at the estimated median earnings for working households after tax and national insurance, with the cap for single people set at 70% of the couple rate, broadly in line with Organisation for Economic Co-operation and Development (OECD) equivalence scales.[1]

Once implemented in April 2013, the roll-out was relatively rapid. By the end of that year, around 28,000 families were subject to the cap every month, although many of these families did not know what the policy was or that they would be affected (Finlay et al., 2013). The cap was fixed in cash terms, so as housing costs rose, more families were drawn into the cap. A review of the policy conducted 1 year after April 2013 lauded the reform

[1]No other adjustments were made for household size or composition.

(DWP, 2014) in part because it was encouraging people into work (Kaur et al., 2014).

The benefit cap was so popular with the general public (Finlay et al., 2013) that the Conservative Party included a commitment in their 2015 manifesto to lower its level; and after they won a majority in that year, the cap was indeed lowered in November 2016 (as noted above). This increased both the number of households affected and the average impact. By May 2017 around 68,000 families had their benefits capped each month, up from 20,000 in August 2016, with 49% capped by more than £50 per week, up from 42%.[2]

The cap does not affect everyone in receipt of state benefits. The policy is explicitly targeted at families with little or no paid work, so it does not apply to those receiving Working Tax Credit (WTC) which implies a regular amount of employment each week. Similarly, where households are claiming Universal Credit (UC), which is gradually replacing the tax credit system, they are exempt if anyone in the household earns a salary equivalent to 16 hours of work per week paid at the minimum wage. Other groups are also exempt, including pensioners and families with at least one adult receiving financial support because of a disability that stops them from working. The cap disproportionately affects women and children. In 2018, more than 70% of capped families were single parents and most of these were headed by women (∼90%) (DWP, 2022). Over 93% of capped families had children and the majority were larger families with three or more children (DWP, 2022).

The loss of income experienced by capped families is significant. The average reduction is approximately £2600 per year, and is slightly higher for families with children (DWP, 2022). This is approximately a 10% cut in total family income (more details on who is affected are in Web Appendix 1, Supporting Information). This puts household budgets under pressure. The immediacy of the financial impact of the cap can be temporarily reduced by Discretionary Housing Payments, which can (at the discretion of local government) provide transitional payments to households subject to the benefit cap, thereby alleviating the immediate impact of the cap on household incomes. This means it may be a few months before families facing the cap see their household incomes fall to the fullest extent entailed by the cap. We take account of this in our interpretation.

While one of the intended outcomes was for affected households to move to a different property with lower rents, this has happened very rarely in practice (Griggs et al., 2023), likely because cheaper properties are scarce or because of a reluctance to move far from schools and social networks (DWP, 2022). This leaves capped households with three remaining options: they can find work, reduce spending, or claim disability-related social security.

THE DUAL IMPACT OF THE BENEFIT CAP ON LABOUR MARKET ACTIVITY

The loss of income to families affected by the cap could have two divergent effects on labour market behaviour. The first and most obvious consequence is that those affected find employment in order to escape the cap and increase their household income overall (Eissa & Hoynes, 2004; Jensen & Blundell, 2022). This is the behavioural response desired by the government. As noted above, there is some existing evidence that a small number of people affected by the cap did indeed move into employment after being subject to it (Kaur et al., 2014); this is consistent with evidence from other settings in which increased benefit generosity reduced labour supply (González, 2013; Jensen & Blundell, 2022). Our first hypothesis then is that the benefit cap will increase employment.

There is a second consequence of the benefit cap on labour market status, and this is our second hypothesis. The policy may push people into economic inactivity while keeping them attached to the social security system. This may occur if people begin claiming disability-related social security, thereby securing an exemption from the cap. Such new claims may happen for two reasons. First, a person may start claiming for something they are already eligible to receive. According to the Labour Force Survey (LFS), before the cap was lowered in 2016, around 26% of people at risk of being capped had some kind of health problem, with around 18% of the sample considered to be disabled (according to the definition used in government data). However, around 87% of those who were (a) at risk of being capped before the cap was lowered and (b) regarded as disabled under the Equality Act were not claiming any form of disability-related social security. In other words, a sizeable group of people were at risk of being capped and could have potentially claimed some form of disability-related social security but chose not to do so (Geiger, 2020). Second, capped households may start claiming if the policy itself harms their health, including their mental health, making it harder to engage in work-search activity (Hussain et al., 2020). There is evidence that the benefit cap did

[2]Most recently, the cap was increased for the first time since it was established as part of wider inflation linked increases to benefit levels in April 2023. From April 2023, the cap has increased to £25,000 (for families in London) and £23,000 for families outside of Greater London.

cause harm to mental health. After the cap was lowered in 2016, the number of people experiencing mental health problems increased substantially: by the end of 2018, over 30% of people at risk of being capped were reporting depression-like symptoms (Reeves et al., 2022). If the cap is found to be increasing economic inactivity via disability-related social security claims, it would point to a counterproductive impact of the policy in terms of its intended outcomes, and would stand in contrast to other evidence on the effect of cuts in benefit generosity.

The labour market effects of the cap may depend on demographics characteristics, such as age, sex and ethnicity. In particular, older people are often inactive because of health conditions (GOV.UK, 2023) and so given the health-related mechanisms outlined above we might (hypothesis 3) expect the benefit cap to have a larger impact on older individuals.

Crucially, the impact of the benefit cap on employment is likely to be contingent on local labour market conditions. Moving into suitable employment may be more difficult for lone parents with young children (many of whom are at risk of being capped) (Chzhen & Bradshaw, 2012; Hussain et al., 2020), particularly in less affluent parts of the country. Our fourth hypothesis, then, is that the impact of the cap will vary regionally and will be more likely to increase economic inactivity in less affluent areas.

BENEFIT CAP AS A NATURAL EXPERIMENT

We use the introduction of the more restrictive benefit cap in November 2016 to examine the causal effect of this type of limit on welfare payments on labour market activity. We treat this policy change as a natural experiment, exploiting the fact that the lower cap both expanded the number of households exposed to the policy and increased the financial loss associated with being capped for those households already affected. We focus on the point at which the cap was lowered, rather than the cap's initial introduction in 2013, because the increased coverage makes it easier to identify capped individuals in survey data, and because the 2016 reform extended the reach of the policy beyond London, allowing us to explore differential regional impact. This specific policy change, then, allows us to test whether the families most likely to be affected by the reform—in particular, lone parents with children—respond to this loss in income by finding work or by moving away from the labour market. To date, there is only limited evidence on whether limiting the total monthly value of welfare payments affects labour market participation (Griggs et al., 2023), and this paper attempts to fill this gap.

DATA AND METHOD

We use a large, repeated cross-sectional household survey from the UK, the LFS between January 2015 and December 2018. This is a stratified random sample of private addresses which interviews 90,000 people quarterly, and is frequently used to generate official statistics. In a supplementary analysis, we also use the Family Resources Survey (FRS) ($n \sim 20{,}000$), which has very good data on household income but does not have measures of mental health. We therefore use the LFS for most of our analysis and use the FRS as a sensitivity analysis, as described below.

To estimate the impact of the benefit cap on employment we need to identify those who are at risk of being capped after the policy was changed in 2016 (an intention to treat approach) and compare them to those who are not. The LFS does not, unfortunately, contain a direct measure of whether people are capped. Similar to earlier work on this topic (Reeves et al., 2022), we define families at risk of being capped in terms of whether they meet the following criteria: aged 16–65, in rented accommodation, either a lone parent (with any number of children) or a two-parent family which contains at least three dependent children, and receiving housing benefit and at least one other form of social security (e.g., Universal Credit, Income Support or Jobseeker's Allowance, and Child Tax Credit). We do not exclude people who are employed because labour market activity is our dependent variable. People who meet all of these criteria but who are also in receipt of Working Tax Credits or Universal Credit are excluded because they are exempt from the cap. Lone parents and large families in rented accommodation are the focus of this analysis because these are the main risk factors for being capped prior to the pandemic (DWP, 2022). It is possible that those at risk of being capped also includes those who might cut their hours, lose their eligibility for WTCs, and therefore become subject to the cap. We do not focus on that group here.

The control group is comprised of everyone else in our data set except for: (1) those whose earnings are above the median, (2) those employed in professional occupations and (3) those with a university degree. We exclude these individuals because very few people exposed to the cap are in any of these categories and this therefore increases the comparability between the treatment and the control group ($n = 489{,}014$). We do, however, recognise the imperfections in this contrast and so we also explore the stability of our results when we restrict the control group to specific sub-groups.

Our 'treatment' group (those at risk of being capped) is larger than the proportion of capped households in the

population. According to our definition, around 0.76% of households in our sample were at risk of being capped ($n = 6524$) while approximately 0.25% of households were subject to the cap in official statistics. Using this intention-to-treat approach is necessary because it will allow us to see the behavioural responses of those at risk of being capped: if we focused on those actually capped, rather than those *at risk* of being capped, we would not be able to observe changes in employment status which lead to the cap being lifted. Importantly, however, the composition of this at risk group is similar to that of the group actually subject to the cap as shown in official statistics (Web Appendix 2, Supporting Information). In this respect, our definition is successful. Our 'at risk' sample is largely made up of women (~95%) in their mid-30s, who are lone parents (~97%) and economically inactive (~73%). There are some differences, however. Our sample has a lower share of people in London, a higher share of lone parents, and a lower average number of children.

Another possible source of bias in how we construct the treatment group lies in whether our group of 'at risk' individuals changes after the reform is implemented in ways that might be correlated with employment status. We explore this possibility but find the compositional characteristics are largely stable before and after the reforms were implemented and those changes we do observe were consistent with administrative data (see Web Appendix 2, Supporting Information).

Our main dependent variable is current economic activity. This is split into three categories according to International Labour Organisation definitions: employed, unemployed, and economically inactive. We create three binary variables which examine specific contrasts: not employed (1) versus employed (0); economically inactive (1) versus not economically inactive (0); and unemployed (1) versus not unemployed. In each case, the baseline combines the other two possible categories of economic activity. We approach these variables in this way because if, as we hypothesise, some of the unemployed are moving into economic inactivity, then focussing only on the move from unemployment to employment would potentially over-estimate the reduction in the proportion of the unemployed. We are interested in all transitions, not only those transitions within the labour market (e.g., moving between unemployment and employment). We also explore the data in a multichotomous form with all three categories kept separate. As a final step, we consider the data on economic inactivity in more detail by examining whether there is any change after the 2016 reform in the proportion of people claiming disability-related benefits.

Additionally, we explore whether the labour market effects of the benefit cap vary geographically by estimating the same difference-in-differences model on the affluent (East Midlands, Eastern, London, South East, Scotland and South West) and then the less affluent parts of the country (North East, North West, Northern Ireland, Yorkshire and Humberside, West Midlands, and Wales), as defined by their gross disposable household income.

We also use a measure of mental health in a mediation analysis which tests whether the policy's impact on economic activity is mediated by its impact on mental health. The measure of mental health which we use is derived from a battery of health questions in which respondents are asked to affirm whether they have (no time period is specified) 'depression, bad nerves or anxiety'. People are coded as 1 if they describe themselves as having depression, bad nerves or anxiety and 0 otherwise.

Statistical analysis

In the few months following the lowering of the benefit cap in November 2016, the numbers being capped increased from ~20,000 households to ~70,000 households, exposing many more families to the cap. In addition, those families already capped faced further reductions in their incomes when the cap was lowered. We examine the labour market effects of those at risk of being capped before and after November 2016, and compare those at-risk of being capped with those who experienced a low risk according to our indicator.

We start by estimating a series of OLS difference-in-differences models with the following 2 × 2 specification:

$$\text{EconAct}_{i,t} = \alpha + \beta_1 \text{AtRisk}_{i,t} + \beta_2 \text{Policy}_{i.t} + \beta_3 \text{AtRisk}_{i,t} \times \text{Policy}_{i,t} + X_{i,t} + \epsilon_{i,t},$$

where i denotes individuals and t the time-period in which the data were collected. *EconAct* is a vector of 3 binary variables which compare types of economic activity (e.g., employed vs not employed). α is the constant (which in the model reports the probability of experiencing a particular type of labour market activity, it varies across the different dependent variables, before the policy change and among those at a low risk of being capped). *AtRisk* is a dummy variable which is 1 if the respondent meets the criteria described above for being at risk of being capped ($n = 6524$) and 0 otherwise ($n = 489{,}014$). *Policy* is 1 if an individual was interviewed during or after November 2016 ($n = 257{,}985$) and 0 otherwise ($n = 237{,}553$). *AtRisk* × *Policy* is an interaction term which captures those who are at risk of being capped and

who are interviewed after the cap has become more restrictive. $\beta_z X_{i,t}$ is a vector of control variables. These include age (measured in years), age squared, gender, the subnational region in which the respondent lives, whether they self-report being 'white' in a question about ethnicity, whether they have other health problems aside from 'depression and anxiety', whether they are a renter or not, and their education (7 categories ranging from university degree to no qualifications). ε is our error term. We also estimate multinomial logistic regression models which retain the three categories in the dependent variable as a sensitivity check. Our standard errors are clustered at the level of the treatment period, a conservative approach (Brewer et al., 2018).

In seeking to identify the causal effect of this policy on economic activity we assume that, in the absence of changes to the benefit cap, the trends in various forms of economic activity amongst the 'at risk of being capped' group would have been identical to the trend for the 'not at risk' group. We conduct extensive checks of the parallel trends assumption. This is particularly important in this case because there are baseline differences between those at-risk of being exposed to the cap and those not at-risk of being unexposed. Specifically, we visualise the trends for both groups over time but we also formally examine this assumption using a variety of techniques that have been recently recommended (Kahn-Lang & Lang, 2020), such as exploring whether results are consistent when adjusting for the predicted linear trend and when adjusting for interactions between covariates and the intervention (see Web Appendix 3, Supporting Information). We also re-estimate our main models using an interrupted time series design to test whether the trends in economic activity diverge after November 2016, when the cap becomes more restrictive. This confirms the parallel trends assumption as well as reinforcing our main results. In short, despite the differences between the two main groups, our extensive checks suggest that the parallel trends assumption is plausible in this case. As there is no variation in the timing of exposure to the policy that we can detect, there is no need to account for this (Callaway & Sant'Anna, 2021). However, we do check our results are stable when using Sant'Anna and Zhou's doubly-robust difference-in-differences estimator (Sant'Anna & Zhao, 2020).

Another part of our empirical strategy is to examine what lies behind any changes in economic inactivity. We do this in two ways. First, we look at the composition of the economically inactive in terms of whether they claim disability-related social security over this period. Here, we hypothesise that there will be an increase in economic inactivity and that this rise in economic inactivity will be linked with a rise in the proportion of people claiming disability-related social security. Second, we consider whether the mental health effects of the benefit cap mediate the relationship between the policy change and economic activity. In particular, we hypothesise that an increase in mental health problems may cause a rise in economic inactivity among those affected. Testing mediation is not straightforward, however. It requires additional assumptions which are challenging to meet with observational data and harder to test. As a result, we deploy the approach developed by Imai and colleagues to estimating mediation (Imai et al., 2011), which enable us to estimate how sensitive these estimates are to violations to the sequential ignorability assumption, key additional assumptions in mediation analysis (Imai et al., 2011).

Sensitivity analyses

We explore whether our results are sensitive to design choices. We describe these details below in the appendix but these include changing the composition of the 'control' group. Here we focus on three contrasts, comparing families at risk of being capped with three other groups who are not subject to the cap: (1) larger families in rented accommodation but not in receipt of benefits; (2) larger families who are home owners; and (3) those who receive at least one form of social security. We also explore whether our results are explained by other welfare reforms being implemented over this period (the main ones being the two-child limit and the roll-out of Universal Credit) and we discuss what these were and how they might affect our analysis in detail in an appendix (see Web Appendix 4, Supporting Information). None of the other reforms differentially affected labour supply incentives in the way that the benefit cap did.

We also use a separate dataset, the FRS, to develop an alternative way of identifying our treatment group, those at risk of being capped. As the LFS lacks a measure of total income from the government, our indicator of those at risk is inevitably blunt. The FRS contains more detailed measures of benefit receipt, allowing us to more accurately (albeit still imperfectly) identify who is in fact subject to the cap. We create a statistical model in the FRS data which predicts whether individuals are likely to be capped or not (see Web Appendix 12, Supporting Information for full details) and use this to predict the probability of being capped for LFS respondents. We then re-estimate our models using this alternative way of identifying the treatment group. We use the FRS for this sensitivity analysis rather than for our main estimates because the FRS sample size is considerably smaller than that in the LFS and there is also no indicator of mental health.

RESULTS

Did the benefit cap affect employment and economic inactivity?

The first question we answer is whether rates of unemployment and economic inactivity changed after the benefit cap was lowered. Table 1 shows that when we look at the proportion of people who are not employed (i.e., either unemployed or economically inactive) there is a small negative albeit non-significant effect. Once we include control variables, we find that being outside of the labour market falls (and therefore employment rates rise) by ∼0.90%-points (95% CI: −2.20 to 2.02). However, this masks two divergent trends which are seen very clearly in columns 3–6: those exposed to the more restrictive benefit cap are both more likely to become economically inactive (4.68%, 95% CI: 2.62 to 6.74) and less likely to be unemployed (−4.77%, 95% CI: −3.72 to −5.82). This is more clearly seen in Figure 1, which visualises results from these regression models but estimated using multinomial logistic regression. The results are almost exactly the same. There is a slight increase in employment after the reform alongside a decline in unemployment and a rise in economic inactivity. The rise in economic inactivity among those at-risk of being capped is striking because it is the exact opposite of the trend observed in the control group, which saw a decline in economic inactivity. This rise in economic inactivity is small as a proportion of all economically inactive people but it is a big effect among those subject to the cap (it has the same impact on economic inactivity rates for a mature worker as making that person 4 years older). This confirms hypothesis two and contradicts hypothesis one.

We also explore the timing of these changes. Did the decline in unemployment and the rise in economic inactivity occur after the policy was introduced in 2013, or was it only in 2016 when the cap was lowered? In Figure 2, we show difference-in-difference estimates using quarterly data rather than a binary measure of before and after the policy. This allows us to see when precisely the impact of the benefit cap emerges. It is clear that in both cases the rise in economic inactivity and the decline in unemployment only emerges after the benefit was lowered in late 2016.

This decline in economic inactivity may vary by demographic characteristics, could be particularly acute among older workers. We find no variation by gender or ethnicity (Web Appendix 5, Supporting Information) but did see evidence that older people (41 or older) at risk of being capped are more likely to move into economic inactivity than younger people, confirming hypothesis three.

We also hypothesised that there would be regional variation in the impact of these reforms (Web Appendix 6, Supporting Information). In affluent parts of the country, we see imprecisely estimated increases in economic inactivity but we do see stronger associations with reductions in unemployment and increases in employment. In contrast, in poorer parts of the country we find very clear

TABLE 1 The introduction of the benefit cap and economic activity.

	Not in employment (unemployed or economically inactive)		Economically inactive		Unemployed	
	(1)	(2)	(3)	(4)	(5)	(6)
Difference-in-differences: Capped individuals compared to uncapped individuals after the reform	−0.0074 (0.012)	−0.00090 (0.011)	0.041** (0.012)	0.047** (0.011)	−0.049** (0.0054)	−0.048** (0.0054)
Change over time for the non-capped individuals	−0.011** (0.0013)	−0.012** (0.0012)	−0.0025* (0.0013)	−0.0038** (0.0011)	−0.0084** (0.00059)	−0.0080** (0.00058)
Difference between capped and non-capped individuals at baseline	0.58** (0.0078)	0.51** (0.0070)	0.40** (0.0075)	0.36** (0.0068)	0.18** (0.0035)	0.16** (0.0035)
Constant	0.36** (0.00091)	1.69** (0.0094)	0.31** (0.00088)	1.55** (0.0091)	0.051** (0.00041)	0.14** (0.0047)
Controls for covariates		Y		Y		Y
Number of individuals	538,790	538,771	538,790	538,771	538,790	538,771

Note: Standard errors are reported in parentheses. In terms of a regression equation, the α is reported as the row 'constant', $\beta_1 \text{AtRisk}_{i,t}$ is 'Difference between capped and non-capped individuals at baseline', $\beta_2 \text{Policy}_{i,t}$ (capturing before and after the policy) is 'Change over time for the non-capped individuals', and $\beta_3 \text{AtRisk}_{i,t} \times \text{Policy}_{i,t}$ is the interaction effect and is reported in 'Difference-in-differences'. Data comes from the Labour Force Survey. Control variables include: age (measured in years), age squared, gender, the government office region in which the respondent lives, whether they self-report being 'white' in a question about ethnicity, whether they have other health problems aside from 'depression and anxiety', whether they are a renter or not, and their education (seven categories ranging from university degree to no qualifications).

*$p < 0.05$; **$p < 0.01$.

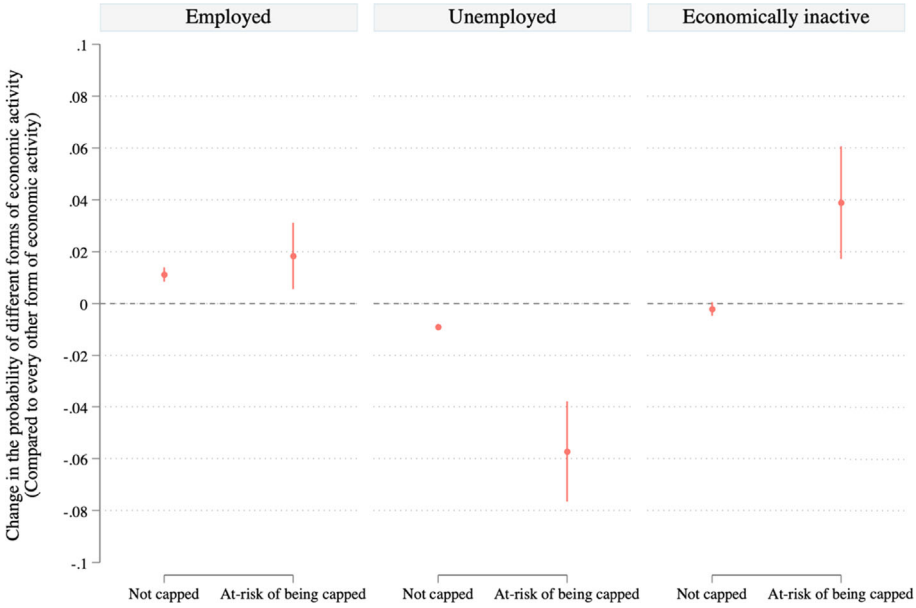

FIGURE 1 The impact of introducing the more restrictive benefit cap on various forms of economic activity. Vertical lines are 95% confidence intervals. The models reported in this figure are derived from a multinomial logistic regression model but the substantive implications are the same as those reported in Table 1.

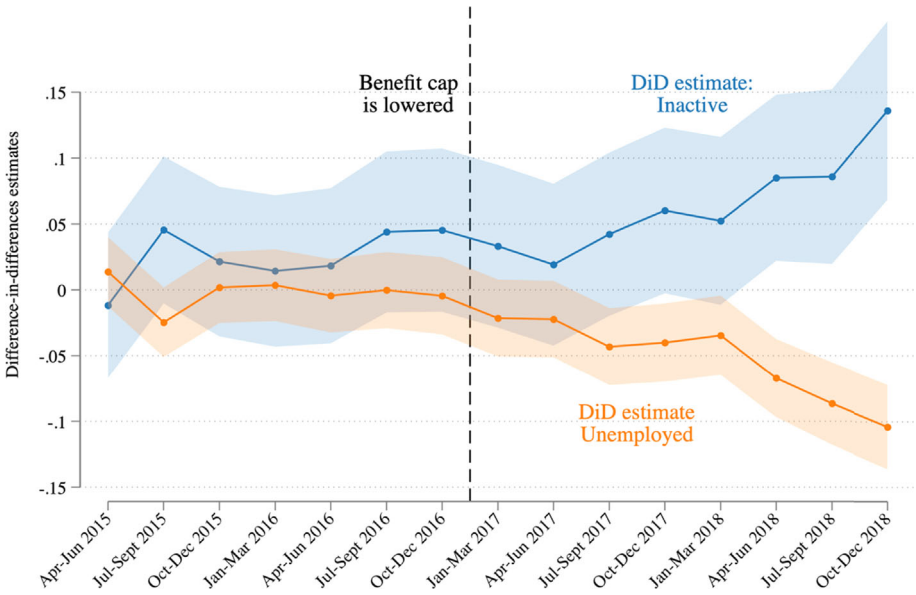

FIGURE 2 The timing of the labour market changes associated with the introduction of a more restrictive benefit cap. Shaded area represents the 95% confidence intervals.

rises in economic inactivity. This suggests therefore that the benefit cap is pushing people from unemployment into economic inactivity in poorer parts of the country. In more affluent parts of country, there is some evidence that the benefit cap may be increasing employment, confirming hypothesis four.[3]

[3]We also examine the impact of these reforms on hours in Web Appendix 7, Supporting Information.

Sensitivity tests

We conduct a series of sensitivity analyses which explore whether some of the assumptions we have made in the main analysis affect our findings. First, we re-estimate our models using the doubly-robust difference-in-differences procedure developed by Sant'Anna and Zhao (2020), finding that our results are almost exactly the same (Web Appendix 8, Supporting Information). Second, we exploit

various exclusions to the benefit cap to see whether restricting the households included in the 'control' group changes our results. We focus, as described above, on those that are more similar to the capped group than the unrestricted comparison group used in the main analysis. Each model is in the same direction and the difference-in-differences estimates are of approximately the same size (see Web Appendix 9, Supporting Information). We also explore whether our results change if we remove people who were potentially affected by the two-child limit, given this policy was introduced only 5 months after the cap was lowered and reduces the amount of support larger families can claim even if they are in work, thereby reducing the financial incentive for capped households to move into employment. The interaction between these two policies could make it difficult to isolate the effect of the benefit cap, but in practice we find that removing these individuals does not alter our findings, perhaps because relatively few households were affected by the two-child limit at the start (Web Appendix 10, Supporting Information). Finally, other policy changes in the years preceding this reform could have impacted employment outcomes (such as the roll-out of Universal Credit from 2013 or the reductions in the local housing allowance in 2011). To account for this, we therefore compare our capped households to those who are claiming any other form of social security (capturing a range of reforms that were all intended to incentivise employment), finding that our results are consistent (Web Appendix 11, Supporting Information).

Second, we report on the results from our FRS model of benefit cap risk, which attempts to address the problem of identifying capped families in the LFS. We use the predicted probabilities from our FRS model to analyse the association between the benefit cap and economic activity in a variety of ways, described in detail in Web Appendix 12, Supporting Information. Here, we report briefly the results of using the probability of being capped (see Web Appendix 12, Supporting Information). We conduct a difference-in-difference analysis similar to that presented in Figure 1; the results, shown in Web Appendix 13, Supporting Information, are very similar in that they indicate a statistically significant reduction in unemployment and a statistically significant increase in inactivity for people at risk of being capped, after the introduction of the lower cap. All the results from the FRS model are consistent with our findings with the simpler risk model from the LFS.

Did the benefit cap affect the proportion of people claiming disability-related social security?

Households at-risk of being subject to the benefit cap can become exempt if they claim disability-related social security. We therefore check whether the introduction of the more restrictive cap in November 2016 was also associated with a rise in the proportion of households at-risk of being capped claiming disability-related social security. Here, we estimate the same difference-in-differences models as above (adjusting for covariates) except that now we change the dependent variable to whether respondents are claiming a form of disability-related social security or not. The key coefficient from this model suggests that the proportion of people at-risk of being capped claiming disability benefits increased by 3.67%-points (95% CI: 2.34% to 4.99%). Recall that the estimated rise in the proportion of people at-risk of being capped who were now economically inactive was around 3.59%-points (95% CI: 1.37 to 5.80). Of course, not everyone who claims disability-related social security is economically inactive and some forms of disability-related social security can be combined with paid work. In our sample of households at-risk of being capped, however, over 90% of those receiving disability-related social security were economically inactive, including recipients of Personal Indepedence Payments (PIP). Moreover, our data contains information on the reasons for people becoming economically inactive. The only reported reason that increases after the policy was introduced is long-term sickness. None of the other reasons see any change at all (such as family commitments or becoming a student), suggesting that the benefit cap is driving economic inactivity almost solely because of health-related reasons. Overall, then, this suggests that a significant proportion of the rise in economic inactivity was potentially driven by people making new claims for disability-related social security which made them exempt even if those specific forms of support could, in theory, be combined with paid work.

Where are these new claims coming from? We know from earlier work that there was a rise in mental health problems among those at-risk of being capped over this period. Was there also a rise in other health problems? We find no clear change in health problems that are not mental health problems (see Figure 3, panel 1: $p = 0.582$). We do not observe any changes in the probability that people report having past health problems that limited their daily activity (see Figure 3, panel 2: $p = 0.218$). There is no change in the probability that people report having past health problems that limited their daily activity and that have lasted for more than 1 year (see Figure 3, panel 3: $p = 0.411$). And finally, there is no change in the probability that people with health problems expect those problems to last longer than 1 year into the future (see Figure 3, panel 4: $p = 0.615$). In other words, although there is some evidence to suggest that more people are claiming disability-related benefits there is no evidence that

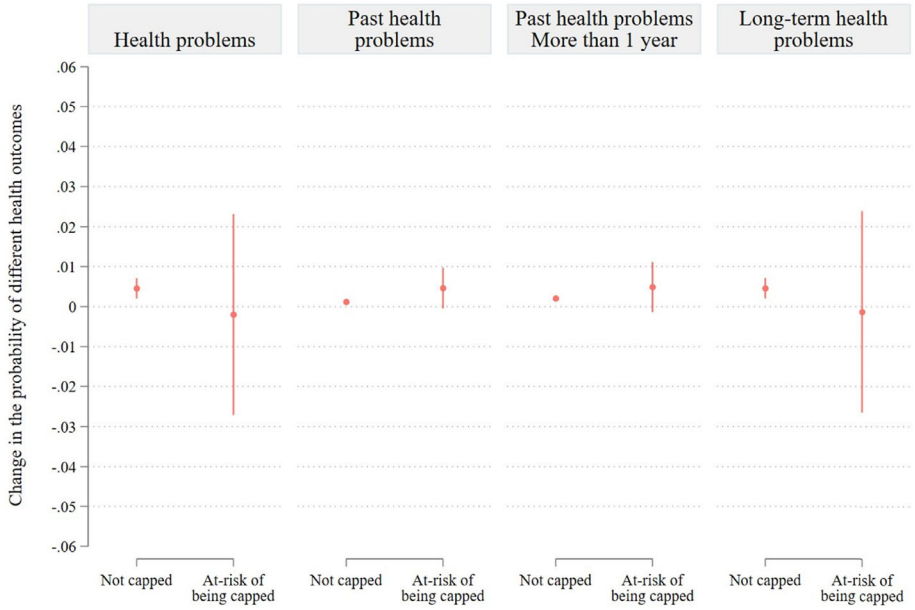

FIGURE 3 Did the benefit cap increase health problems? The outcome variable 'health problems' includes all health problems except those related to mental health. Past health problems is 1 if respondent has had health problems that limit daily activities. Past health problems more than 1 year is 1 if respondent has had health problems that limit daily activities for more than 1 year. Long-term health problems is 1 if respondent believes their health problems will last for more than 1 year.

(outside of mental health) the physical health of these claimants has changed much at all. Therefore, these new claims appear to be people who were eligible for disability-related social security but who were not claiming such benefits before they were capped (for more on take up of benefits see Bennett, 2023).

Did mental health mediate the association between the benefit cap and economic activity?

Having established that the introduction of a more restrictive benefit cap in 2016 increased employment but also increased economic inactivity (in part because it led to more people claiming disability-related benefits), we now explore whether mental health mediates the relationship between the benefit cap and economic inactivity (whatever the reason). Previous work has already shown that the policy reform led to higher rates of mental ill health among those who were affected by the policy (Reeves et al., 2022). We therefore test whether mental health sits on the pathway between the policy and labour market outcomes.

First (see Figure 4), we consider whether the association between the changes to the benefit cap and the rise in economic inactivity are mediated by the effects of the benefit cap on mental health. We find that there is some mediating impact of mental health, in that the policy change affects mental health and mental health influences economic inactivity. More precisely, consistent with our main models, we find that the total effect of implementing this more restrictive benefit cap policy is associated with a ~4%-point (95% CI: 2.41 to 6.34) increase in the proportion of people who are economically inactive. This is consistent with our main results reported above. The average causal mediation effect of this policy on economic inactivity through its effect on mental health is 0.632%-points (95% CI: 0.050 to 1.214). This suggests that around 15.3% of the total effect of the policy on economic inactivity is mediated via its effect on mental health. The average direct effect of the policy (i.e., not via mental health) is 3.304%-points (95% CI: 1.139 to 5.878). This suggests that the association between reducing the benefit cap and changes in economic inactivity are partially but not entirely mediated by their impact on the mental health of those at-risk of being affected. This is consistent with the evidence we have already reported, which suggests that some of the new disability claims (which can include mental ill health) were people who were previously eligible but who were not claiming before they were capped.

Given the challenges to proving mediation, we conduct a sensitivity analysis in order to determine how substantial any unobserved confounding of this mediation effect would need to be in order to push the coefficient to 0 (Web Appendix 14, Supporting Information). We do this by simulating the correlation (ρ) between the

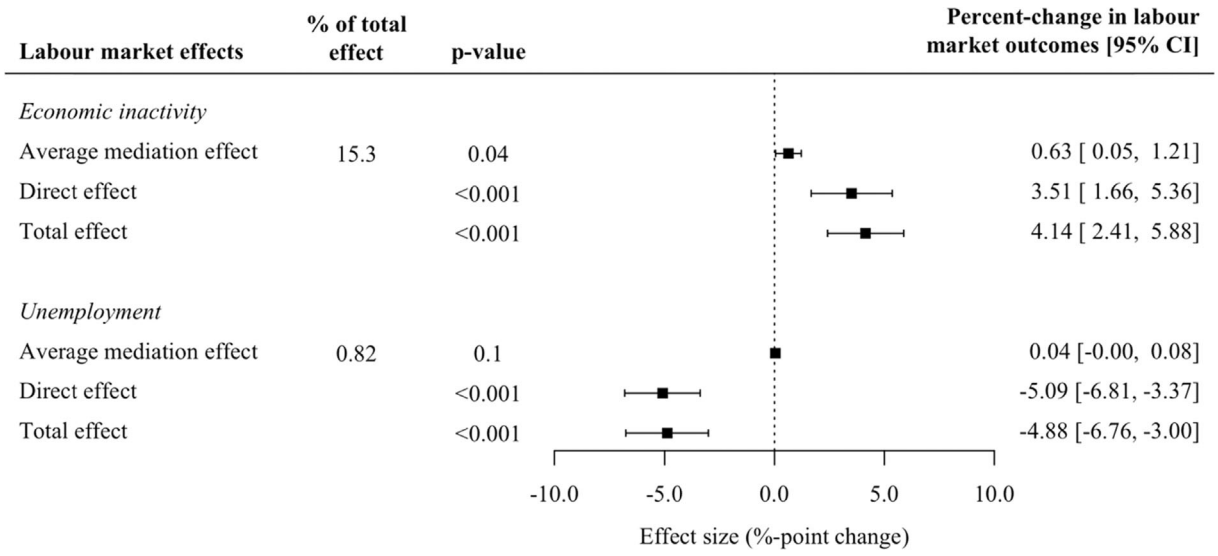

FIGURE 4 Does mental health mediate the association between the benefit cap and labour market activity?

residuals in the outcome and the mediator regressions. If there exists some unobserved pre-treatment confounders which affect both the mediator and the outcome, we expect that the sequential ignorability assumption is violated (a crucial assumption in identifying mediation effects) and ρ is no longer zero. We therefore explore what happens to the estimated average causal mediation effect as ρ goes above and below zero. We find evidence that our estimated mediation effect is robust, as it would require the correlation between the residuals of these models to be ∼0.2. Or, to put this another way, the total explained variance of the model would need to double to remove our estimated mediation effect.

Second, we consider whether the association between the changes to the benefit cap and the reduction in unemployment are mediated by the effects of the benefit cap on mental health. Recall that unemployment here is contrasted to those who are employed and those who are economically inactive, in order to avoid over-estimating the impact of the policy reform on those in the labour market. Again, we find that there is some mediating effect, but it is much smaller for unemployment than for economic inactivity. For example, consistent with our main models, we find that the total effect of the implementation of this more restrictive benefit cap policy is associated with a ∼5%-point (95% CI: −3.03 to −6.76) reduction in the proportion of people who are unemployed. This is consistent with our main results reported above. The average causal mediation effect of this policy on unemployment through its effect on mental health is 0.040%-points (95% CI: −0.0031 to 0.083). This suggests that less than 1% of the total effect is mediated via its effect on mental health. This is not significantly different from zero and it is very small. The average direct effect of the policy (i.e., not via mental health) is −4.92%-points (95% CI: −3.37 to −6.81). This suggests that the association between reducing the benefit cap and changes in mental health is only a very minor part of explaining the reduction in unemployment. Our sensitivity analysis (see Web Appendix 14, Supporting Information) also confirms this, suggesting that it would take a much smaller degree of confounding to remove the mediating relationship observed here for mental health.

DISCUSSION

In this paper, we take advantage of a reduction in the benefit cap to explore whether the reduction in income that follows being capped—which was intended to incentivise employment—affected labour market activity. We find that implementing a more restrictive cap both (a) reduced unemployment and (b) increased economic inactivity (which appears to be linked with more people claiming disability-related benefits), corresponding to but also deepening existing evidence (Griggs et al., 2023). It is thereby simultaneously pushing some people towards work and other people away from it. Indeed, the net effect of the policy on employment rates is basically zero, suggesting that most of the reduction in unemployment is actually driven by increases in economic inactivity. Finally, we also find evidence that part of the reason this policy change led to economic inactivity is because it harms mental health, leaving people in a position where their work-readiness has declined and where they are now eligible to claim disability-related benefits.

Our findings have implications for broader debates about how the design of social security systems affects labour market outcomes (Eissa & Hoynes, 2004). At one level, we find evidence consistent with standard economic models of how people respond to economic incentives, namely that reductions in social security payments can incentivise people into employment, presumably because the reduction in incomes increases financial pressure (Jensen & Blundell, 2022). In this specific case, working at least 16 h per week also provides a way for people to escape the cap. This creates a set of clear incentives which seem to have encouraged some people to return to work. However, our findings also illustrate that such reductions can also have unintended impacts on economic inactivity. In particular, increasing financial strain on households has two consequences. On the one hand, the financial pressures lead some (likely those who were already eligible) to claim disability-related social security. On the other hand, the reduction in income harms their mental health and actually push some of those affected further away from the labour market (again leading them to claim disability-related benefits) (Hussain et al., 2020). While the net effect of these two consequences is a very slight increase in employment, it is important to unpack how the bifurcation of capped households may deepen inequalities. It may be that those who are incentivised into work are those who were already able to take up work because their family situations allowed it, while those being moved away from employment are those who have competing responsibilities (often childcare) that make it difficult for them to return to work right now. These may well be people who want to work in the future when their circumstances allow, and there is a chance that the policy may be making it less likely that these people will return to the labour market in that future moment.

Crucially, the regional variation in the impact of the benefit cap helps us understand how structural economic conditions shape the effectiveness of these kinds of incentives. We find that in poorer parts of the country the benefit cap seems to be pushing unemployed people into economic inactivity whereas in affluent areas there is some evidence that the cap is incentivising people into work. This variation is particularly salient because the early trials of the cap were conducted when the policy largely affected people in London. In other words, those early findings of higher employment rates may have been driven largely by the fact that the affected people were living in a dynamic and high-wage part of the British economy. As the reach of the cap has extended to other parts of the UK, its impact on employment may have diminished to such an extent that we find no net impact on employment at all. Clearly, the economic background against which such policies operate is crucial.

There are a number of limitations to our analysis. First, we do not have within our data an entirely exchangeable comparison group for those affected by the policy, and so there are important differences at baseline between those at risk of being capped and those who are not at risk of being capped. This creates uncertainty regarding our estimate of the causal effect of this policy change (Dunning, 2012). Second, we cannot perfectly identify those who are (or could have been) subject to the cap using the two datasets which underpin our analysis. We address this limitation by adopting an intention-to-treat approach, which likely leads to conservative estimates of the causal effect. We have also conducted a series of sensitivity analyses which explore different ways of identifying those at risk of being capped. These reveal that our main estimates may indeed be conservative and that the impact of the benefit cap could be even larger than our results suggest. At the same time, however, our construction of the at risk group over-estimates the proportion of lone parents, who might be less responsive than others to the labour market incentives of the benefit cap (Chzhen & Bradshaw, 2012), and this might lead to under-estimating the impact of the cap on employment. Together, this creates some uncertainty about the exact magnitude of the impact of this policy.

While the benefit cap was very popular with British voters, it is not clear that it has accomplished its stated goals. It has not contributed to a significant reduction in government spending, a fact that some seemed to acknowledge even while the policy was being developed (Chakelian, 2021). It has led to a reduction in unemployment, but has also pushed an equally large number of people into economic inactivity. Our results suggest that this is in part because the policy harmed mental health but also because many of the people who were at-risk of being capped were already eligible for disability-related social security. The UK's benefit cap, then, bears some similarities to other policies which attempt to limit spending on social security in that it does motivate some to return to work but it also simultaneously pushes other further away from the labour market. As a result, this policy has left economically marginalised families struggling to make ends meet (Edmiston, 2021), has harmed mental health, and pushed some further away from the labour market. This means the policy is failing against its own objectives, while also causing real and significant harms.

ACKNOWLEDGEMENTS

We would like to thank Eleni Karagiannaki for comments on an earlier version of this paper.

FUNDING INFORMATION

Aaron Reeves, Kitty Stewart, Ruth Patrick all express appreciation to the Nuffield Foundation for funding this research (Grant/Award Number: FR-23208). The Nuffield Foundation is an independent charitable trust with a mission to advance social well-being. It funds research that informs social policy, primarily in Education, Welfare and Justice. It also funds student programmes that provide opportunities for young people to develop skills in quantitative and scientific methods. The Nuffield Foundation is the founder and co-funder of the Nuffield Council on Bioethics and the Ada Lovelace Institute. The Foundation has funded this project, but the views expressed are those of the authors and not necessarily the Foundation. Visit www.nuffieldfoundation.org.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available in UKDA at https://beta.ukdataservice.ac.uk. These data were derived from the following resources available in the public domain: Labour Force Survey, https://beta.ukdataservice.ac.uk/datacatalogue/series/series?id=2000026.

ORCID

Aaron Reeves https://orcid.org/0000-0001-9114-965X
Mary Reader https://orcid.org/0000-0002-2154-1813

REFERENCES

Bennett, F. (2023). Take-up of social security benefits: Past, present – and future? *Journal of Poverty and Social Justice*, *1*, 1–24. https://doi.org/10.1332/17598273Y2023D000000005

Brewer, M., Crossley, T. F., & Joyce, R. (2018). Inference with difference-in-differences revisited. *Journal of Econometric Methods*, *7*(1), 21070005. https://doi.org/10.1515/jem-2017-0005

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

Chakelian, A. (2021). "It polled off the charts": The benefit cap is for popularity not saving money, claims former Tory welfare minister. New Statesman. https://www.newstatesman.com/politics/welfare/2021/12/it-polled-off-the-charts-the-benefit-cap-is-for-popularity-not-saving-money-claims-former-tory-welfare-minister

Chzhen, Y., & Bradshaw, J. (2012). Lone parents, poverty and policy in the European Union. *Journal of European Social Policy*, *22*(5), 487–506.

Dunning, T. (2012). *Natural experiments in the social sciences*. Cambridge University Press.

DWP. (2014). *The benefit cap: A review of the first year (Cm 8985)*. DWP.

DWP. (2022). Benefit cap: Number of households capped to February 2022. https://www.gov.uk/government/publications/benefit-cap-number-of-households-capped-to-may-2020/benefit-cap-number-of-households-capped-to-may-2020

Edmiston, D. (2021). Plumbing the depths: The changing (socio-demographic) profile of UK poverty. *Journal of Social Policy*, *1–27*, 385–411. https://doi.org/10.1017/S0047279421000180

Eissa, N., & Hoynes, H. W. (2004). Taxes and the labor market participation of married couples: The earned income tax credit. *Journal of Public Economics*, *88*(9–10), 1931–1958.

Finlay, S., Marshall, B., & Hill, L. (2013). *The benefit cap: Public perceptions and pre-implementation effects (850)*. DWP.

Frijters, P., Johnston, D. W., & Shields, M. A. (2014). The effect of mental health on employment: Evidence from Australian Panel Data. *Health Economics*, *23*(9), 1058–1071. https://doi.org/10.1002/hec.3083

García-Gómez, P., Jones, A. M., & Rice, N. (2010). Health effects on labour market exits and entries. *Labour Economics*, *17*(1), 62–76. https://doi.org/10.1016/j.labeco.2009.04.004

Geiger, B. B. (2017). Benefits conditionality for disabled people: Stylised facts from a review of international evidence and practice. *Journal of Poverty and Social Justice*, *25*(2), 107–128.

Geiger, B. B. (2020). Has working-age morbidity been declining? Changes over time in survey measures of general health, chronic diseases, symptoms and biomarkers in England 1994–2014. *BMJ Open*, *10*(3), e032378.

González, L. (2013). The effect of a universal child benefit on conceptions, abortions, and early maternal labor supply. *American Economic Journal: Economic Policy*, *5*(3), 160–188.

GOV.UK. (2023). *Economic inactivity*. https://www.ethnicity-facts-figures.service.gov.uk/work-pay-and-benefits/unemployment-and-economic-inactivity/economic-inactivity/latest/

Griggs, J., Jones, H., Basi, T., & Piggott, H. (2023). *Evaluation of the lower benefit cap: Findings from the mixed method longitudinal evaluation (1026)*. DWP.

Hussain, M. A., Ejrnæs, M., & Larsen, J. E. (2020). Are benefit reductions an effective activation strategy? The case of the lowest benefit recipients in Denmark. *Journal of Social Policy*, *1–19*, 569–587. https://doi.org/10.1017/S0047279420000318

Imai, K., Keele, L., Tingley, D., & Yamamoto, T. (2011). Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review*, *105*(4), 765–789.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning: With applications in R*. Springer Science & Business Media.

Jensen, M. F., & Blundell, J. (2022). Income effects and labour supply: Evidence from a child benefits reform. Available at SSRN 4053318.

Kahn-Lang, A., & Lang, K. (2020). The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications. *Journal of Business & Economic Statistics*, *38*(3), 613–620. https://doi.org/10.1080/07350015.2018.1546591

Kaur, K., Mosley, R., Nutley, T., & Thompson, J. (2014). *Benefit cap: Analysis of outcomes of capped claimants (11)*. DWP.

Leyshon, A. (2021). Economic geography I: Uneven development, 'left behind places' and 'levelling up' in a time of crisis. *Progress in Human Geography*, *45*(6), 1678–1691. https://doi.org/10.1177/03091325211011684

Mind UK. (2023). Benefits assessments making people with mental health problems more unwell, Mind research reveals. https://www.mind.org.uk/news-campaigns/news/benefits-assessments-making-people-with-mental-health-problems-more-unwell-mind-research-reveals/

Osborne, G. (2010). *Spending review 2010: Presented to parliament by the chancellor of the exchequer by command of her majesty*. TSO.

Reeves, A., Fransham, M., Stewart, K., & Patrick, R. (2022). Does capping social security harm health? A natural experiment in the UK. *Social Policy & Administration*, *56*(3), 345–359. https://doi.org/10.1111/spol.12768

Sant'Anna, P. H. C., & Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, *219*(1), 101–122. https://doi.org/10.1016/j.jeconom.2020.06.003

Tonutti, G. (2018). The impact of welfare reform in London—Using benefits data to track the effect of government policies on employment outcomes. *International Journal for Population Data Science*, *3*(2) https://www.researchgate.net/publication/325885230_The_impact_of_welfare_reform_in_London_-_using_benefits_data_to_track_the_effect_of_government_policies_on_employment_outcomes

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Reeves, A., Fransham, M., Stewart, K., Reader, M., & Patrick, R. (2024). Capping welfare payments for workless families increases employment and economic inactivity: Evidence from the UK's benefit cap. *International Journal of Social Welfare*, 1–14. https://doi.org/10.1111/ijsw.12651