Did the introduction of the benefit cap in Britain harm mental health? A natural experiment approach

Aaron Reeves, Mark Fransham, Kitty Stewart, and Ruth Patrick

CASE/221
November 2020

Centre for Analysis of Social Exclusion
London School of Economics
Houghton Street
London WC2A 2AE
CASE enquiries – tel: 020 7955 6679

Centre for Analysis of Social Exclusion

The Centre for Analysis of Social Exclusion (CASE) is a multi-disciplinary research centre based at the London School of Economics and Political Science (LSE), within the Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD). Our focus is on exploration of different dimensions of social disadvantage, particularly from longitudinal and neighbourhood perspectives, and examination of the impact of public policy.

In addition to our discussion paper series (CASEpapers), we produce occasional summaries of our research in CASEbriefs, and reports from various conferences and activities in CASEreports. All these publications are available to download free from our website.

For further information on the work of the Centre, please contact the Centre Manager, Annie-Rose Nicholas, on:

Telephone: UK+20 7955 6679
Email: a.nicholas1@lse.ac.uk
Web site: http://sticerd.lse.ac.uk/case

© Aaron Reeves
© Mark Fransham
© Kitty Stewart
© Ruth Patrick

All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Editorial note

Aaron Reeves is Associate Professor at the Department of Social Policy, University of Oxford. Mark Fransham is Departmental Lecturer at the Department of Social Policy, University of Oxford. Kitty Stewart is an Associate Director at the Centre of Analysis for Social Exclusion and an Associate Professor in Social Policy at the London School of Economics. Ruth Patrick is Lecturer at the Department of Social Policy, University of York.

We would like to thank the members of our advisory board and Eleni Karagiannaki for comments on earlier versions of this paper.

Abstract

In November 2016, the UK government announced they would be lowering the benefit cap, the total amount a family with no-one in full-time employment can receive from the government in social security. This policy change reduced financial support for large (often lone parent) families and those with high housing costs, and broke the link between needs and entitlements in the British social security system. This policy was intended to incentivise people to return to work but it may have also harmed mental health, especially because those affected by the reform may struggle to find appropriate work or move to cheaper housing. We treat this reform as a natural policy experiment, comparing those at-risk of being capped and those who were not, and then examining the risk of experiencing poor mental health both before and after the cap was lowered. The main outcome is a binary measure of self-reported mental health problems. Drawing on data from 1.4 million individuals collected between January 2015 and December 2018, we find that the prevalence of depression or anxiety among those at-risk of being capped increased by 2.6 percentage points (95% confidence interval: 1.33 to 3.88) compared with those at a low risk of being capped. This association is consistent to a variety of sensitivity tests. We conclude that lowering the total amount of financial assistance families can receive in social security may increase the risk of mental ill health and could have the unintended consequence of pushing out-of-work people even further away from the labour market.

Key words: benefit cap; welfare reform; mental health

JEL number: I14

Corresponding author: aaron.reeves@spi.ox.ac.uk.

1. Introduction

The logic of social security systems in some high-income countries has shifted in recent decades, leading to a reconfiguration of the conditions of entitlements and the structure of the financial support offered to those in precarious economic circumstances (Hills, 2011; Pierson, 1994; Scruggs and Allan, 2006; Timmins, 2017). These reforms have been motivated by a desire to reduce welfare spending but are also symptomatic of a belief that overly generous welfare regimes foster a culture of dependency among recipients (Eikemo and Bambra, 2008; Hills, 2014). One frequent manifestation of this logic has been policy interventions to make social security less generous in order to activate labour market participation. This is underpinned by the assumption that making it harder to make ends meet on benefits alone will encourage people to return to work (Hussain et al., 2020; Pierson, 1994). Beyond labour market activation, another motivation for welfare reform has been to address what some regard as unfairness in the welfare system: where it allows the 'skivers' to reap rewards funded by the 'strivers' through the tax system (Hills, 2014).

This logic and reform agenda has been central to the introduction of a 'benefit cap' in the UK (Lupton et al., 2016). Several countries have experimented with establishing limits on welfare spending. For example, block-grants to US states create an upper bound on total expenditure (Ziliak, 2015). Many US states also impose time limits on some forms of welfare, meaning families receive support for a fixed period of time irrespective of eligibility (e.g. the 60 month federal limit for TANF) (Grogger, 2002). While the UK's 'benefit cap' bears many of the hallmarks of these approaches to social security, it is also unique. Instead of restricting the number of people that can access support or the length of time help is available, it places an absolute limit on the total amount of financial support that any given family can receive from the state in a particular year across almost all forms of social security. The cap is not applied to everyone, however; there are exemptions if a family has someone in work or if a family member is disabled. These exceptions were intended to help the cap target spending reductions and work incentives on a group whose entitlements were perceived to be *prima facie* too large and who were, in principle, able to work. As Prime Minister David Cameron noted at the time: 'are [taxpayers] working hard to give benefits so people can live in homes that they [the taxpayers] can only dream of? I don't think that is fair' (Mulholland, 2010).

Despite being popular among the general public (Finlay et al., 2013), evidence to date suggests that the benefit cap has only been modestly successful in achieving its stated aims of cutting costs and incentivising employment. The amount of money saved has been relatively small, particularly when set alongside the costs of administering the programme, which are intended to judge whether those at risk of being capped are

exempt or not (Work and Pensions Committee, 2019). The minimal cost savings are partly due to the fact that few people have moved to cheaper accommodation as a result of the cap, while the cap has not been a very effective labour market activation tool either (Kaur et al., 2014). It is estimated that around one third of people affected would have moved into work anyway in the absence of this limit, with the cap leading to an increase of around five percentage points (Kaur et al., 2014; Tonutti, 2018).

In this paper, we focus on the potential *unintended* consequences of the policy – the impact on the mental health of those experiencing cuts in income as a result of the cap. Beyond a direct concern about the well-being of the individuals affected, the mental health effects of the benefit cap matter for three main reasons. First, if the mental health of parents declines then this may have negative effects for their children too, potentially adversely affecting their well-being, their educational development and their behaviour (Cooper and Stewart, 2017). Second, worsening mental health may actually push people away from the labour market, in that mental health problems can decrease the likelihood of returning to work (García-Gómez et al., 2010). If the cap does harm mental health it may have the perverse effect of deepening the degree of social exclusion felt by some of these individuals. Third, worsening mental health incurs costs for government. This is particularly salient in the UK where healthcare is almost entirely tax-financed but is also true in the US where the state system only covers a smaller proportion of the population; around 22% of all non-elderly adults with a mental illness in the US are eligible for Medicaid (Kaiser Family Foundation, 2017).

To date, however, there have been no evaluations of the mental health consequences of the benefit cap. On the one hand, negative effects may be expected, as the policy is intended to reduce family income and we know that income is correlated with mental health (Jones and Wildman, 2008; Mackenbach et al., 2005). While there are still relatively few studies that use causal or quasi-experimental techniques to investigate these links, the literature points to a causal relationship between low income and mental health, and especially between low income and maternal mental health (Cooper and Stewart, 2015, 2017). This last point is pertinent as a high percentage of capped households are larger families headed by a lone mother. However, many of the existing studies in this area look at increases rather than cuts to income, while those that have examined the mental health effects of reductions in social security have not focused on lone parents (Reeves et al., 2016). Sudden reductions in benefit income are unusual and therefore little investigated, and we cannot assume they will simply have the reverse effect of an increase in income.

At the same time, the benefit cap is potentially rather different to other cuts to social security benefits, as the cap was deliberately set initially at roughly "the average take home pay of working households" (DWP, 2014:

11). 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). If the cap is set close to median earnings, one might expect even the capped amount to provide a reasonable 'minimum income' for those families affected (Freud, 2013), and therefore the policy might be unlikely to harm mental health in the same way as a cut to benefits for those already living in hardship; indeed, in providing an incentive to move into work it could plausibly even lead to mental health improvements.

Critics, however, have noted that equating an out-of-work family's *total income* with an in-work family's *earnings* ignores differences between families in their composition and therefore spending needs. Critically, it also overlooks the additional support many in-work families are entitled to, which increase their *total income*, and help them meet their needs, for example through child benefits and housing support. Thus, depending on whether any capped families have been enjoying levels of income that exceed family needs, the cap may have relatively small (perhaps even positive) effects on mental health for these individuals, or it may be damaging.

We address this gap in our understanding 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 sample survey to identify those who are at risk of being subject to the cap and those who are not. We then follow these groups over time and 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 (and 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) increased the risk of reporting mental ill health. Our results suggest that the mental health effects of income shocks, such as reductions in social security, need to be viewed alongside the possible labour market effects. While those affected by the benefit cap may be more likely to move into work, our evidence also suggests that this policy may actually push people further away from the labour market because it undermines their well-being.

The benefit cap and welfare reform in the UK

The Conservative-led coalition government was elected in 2010 on a commitment to reduce government spending and, soon after taking office, they announced a raft of reforms aimed at minimising social security expenditures (Osborne, 2010). The benefit cap was part of this policy programme and was initially set at £500 per week (or £26,000 per year) for couples and lone parents, and an equivalent amount of £350 per week

(or £18,200 per year) for single people without children (or whose children do not live with them). The cap applied only to 'workless families'; with exemptions also offered to families in receipt of disability benefits. Implementation started in April 2013 but accelerated from July 2013 onwards. By the end of that year, around 28,000 families were subject to the cap every month.

The government's one year review of the cap was glowing, noting that "we did not fully appreciate the scale of the positive benefits of the cap" (DWP, 2014). Preliminary survey data and interviews suggested the policy had been successful at motivating capped individuals to look for work or, if they had already been seeking work, to look harder (Kaur et al., 2014). The review argued that these changes would deliver wider benefits for society, such as discouraging 'benefit dependency' and 'break[ing] intergenerational cycles of disadvantage' (DWP, 2014). The policy was popular too, with almost 75% of people in some polling data suggesting that people were in favour of the policy (Taylor-Gooby and Taylor, 2015).

A year later, in 2015, the Conservative Party included a commitment to reduce the cap further in their general election manifesto; this reduction was confirmed in the Summer Budget 2015, after the Conservatives had won the election and been returned to power with a majority. In November 2016 the cap was reduced from £26,000 per year to £23,000 per year for families in London (£15,410 for single people) and to £20,000 (13,400 for single people) outside the capital. For most families, it is receipt of housing benefit – assistance with the cost of renting a home – that pushes them into the cap. But this is not true for everyone. Couples and lone parents with around six children can also be affected by the cap even if they are not receiving housing benefit. By March 2017 around 68,000 families were subject to the cap each month. In total, more than 290,000 families had their benefit payments capped between April 2013 and November 2019. The number of capped families rose sharply as a result of the COVID-19 pandemic, as families lost employment as a result of the economic shock or found that they were pushed into the cap by temporary increases in means-tested benefits intended to provide an additional cushion during the crisis. In the three months to May 2020, the number of families that were subject to the benefit cap increased by 93% to 154,000 households (DWP, 2020a).

There are a number of exemptions from the cap. Only those of working age can be capped so pensioner households are not affected. Aside from age, most of the exemptions revolve around two issues: employment status and disability status. One exemption covers claimants in receipt of Working Tax Credit (WTC). Similar to the Earned Income Tax Credit in the US, WTC tops up wages for those in low paid work; the number of hours required for eligibility depends on whether there are children in the household. Families are also exempt if one of the adults receives financial support because of a

disability that stops them from working. There are also informal exemptions. For example, a claimant's local authority may temporarily offer them Discretionary Housing Payments, which provide transitionary support for those struggling to find work or to move to a more affordable property.

The cap has always been controversial because it disproportionately affects particular groups. Initially, for example, many capped families (over 40%) were based in London, where rents are high. Since the cap became more restrictive its effects have reached almost every part of the country. Following the introduction of the lower cap in 2016, only one quarter of capped families in 2017 were in London. Perhaps more troubling is how the cap affects women and children. More than 70% of capped families are single parents and most of these are headed by women (~90%) (DWP, 2020b). In total, over 93% of capped families include children and the vast majority are larger families with three or more children (DWP, 2020b). So acute are the gender inequalities in the policy's impact that the cap has been subject to a number of legal challenges (albeit unsuccessful) on the basis that it violates the Equality Act (Fenton-Glynn, 2015; Hollingsworth, 2015).[1]

The economic shock experienced by capped families is not trivial. On average, those affected by the benefit cap lose around £2,600 per year, slightly more for families with children (DWP, 2020b). This is approximately a 10% reduction in total family income, and would equate to a higher share of disposable income after housing costs, if part of the benefit package is covering rent and the family does not move. This potentially places families under increasing pressure in seeking to make ends meet. Official statistics show that very few people subject to the cap move to a new property – likely due to the desire to maintain social networks and to keep their children in the same schools – and so the remaining options are either to look for a job (often while managing childcare) or to reduce spending (DWP, 2020b). Early indications suggest that, between 2018/19 and 2023/24, the benefit cap will push around 400,000 children deeper into poverty (Tucker, 2019), and these estimates do not take into account the sharp increases in the numbers of capped families as a result of the effects of Covid-19 (DWP 2020).

[1] In *SG and others v Secretary of State for Work and Pensions*, the claimants argued that the benefit cap indirectly discriminates against women because it affects women more than men. A majority of the Court concluded that any disproportionate negative impact is justified, although Lady Hale in dissent explained that the prejudicial effect of the cap 'is obvious and stark [as] it breaks the link between benefit and need.' Similarly, in a legal challenge to the two child limit, the court again ruled that any prejudicial effects… are not 'too high a price to pay'. Finally, in *DA and DS v Secretary State for Work and Pensions*, the claimants argued that the benefit cap discriminated against lone parent families because they have different requirements but again the court disagreed, arguing that single parents are not different to any other persons subject to the cap (Campbell, 2020).

In this paper we use the introduction of the more restrictive benefit cap in November 2016 (which thereby increased the number of people at risk of being capped and the size of the average impact on capped families) to explore how these reductions in income affected mental health. We treat this policy change as a natural experiment to examine the causal effect of reducing welfare payments on mental health. We note that any impact on mental health may not show up in the data immediately but may instead emerge gradually as transitionary support comes to an end and people slowly find it increasingly difficult to make ends meet. We also note that there are two possible reasons why the policy may *not* have harmful effects: if it leads, as intended, to people altering their behaviour and moving back into the workplace, and this move is neutral or positive for mental health (DWP, 2014); and/or if the implicit assumption underlying the policy is correct, and a total benefits package of £20,000 is sufficient to meet families' needs (Freud, 2013). This specific policy change, then, allows us to test whether the families most likely to be affected by the reform – lone parents with children – can lose this extra income without harm. To date, we do not have good evidence on whether the benefit cap will harm mental health or not, and this paper attempts to fill this gap.

2. Data and method

We use two large-scale, repeated cross-sectional surveys from the UK, both of which are used to produce official statistics. The first (and primary) data set is the UK Labour Force Survey (LFS), which interviews 90,000 people every three months. The second data set is the Family Resources Survey (FRS). This is the highest quality survey capturing household income from all sources in the UK and contains around 20,000 families every year. As the FRS does not contain measures of mental health we focus primarily on the LFS, using the FRS for sensitivity analysis, as explained below.

The LFS does not contain a way of formally identifying those who have been capped or not and so we adopt an intention-to-treat approach: we identify those who are *at risk* of being capped and compare them with those who are not. We define families as at risk of being capped as those who are all of the following: aged 16-65, in rented accommodation, either a lone parent 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., Income Support or Jobseeker's Allowance). We exclude those who meet all of these criteria but are in receipt of Working Tax Credits, as households with a member in paid work are exempt. We focus on lone parents and large families in rented accommodation because as noted above these are the main risk factors for being capped (DWP, 2020b). The advantage of this approach to defining those at risk of being capped is that our estimates incorporate any effects on the mental health of those who responded to the policy by moving into work, as well as those who are

capped in practice. This intention-to-treat approach allows us to see the net effect of the reform.

Our main measure suggests around 0.5% of our sample were at risk of being capped, which is ~0.25% higher than what we see in the overall population. Those categorised as at risk in our data are predominantly women (~95%) in their mid-30s, who are lone parents (~97%) and economically inactive (~73%). The majority are white British but ethnic minorities are over-represented. The sample we identify as being at-risk of being capped looks quite similar to the administrative data on capped families in terms of age, gender, whether a child is present in the home, and whether or not it is a single adult family (Supplementary analysis 1). However, our sample has a lower share of people in London, a higher share of lone parents, and a lower average number of children. Our measure of capped individuals is largely stable before and after the reforms were implemented (see Supplementary analysis 1 for balance tests). There was a decline in the proportion of capped individuals in London and there was a decline in the proportion of people who were unemployed and looking for work (some of them entered unemployment and some became economically inactive).

Statistical analysis

We treat the lowering of the benefit cap in November 2016 as a natural experiment. Over the next few months, the numbers being capped increased rapidly (see Figure 1). This meant that many more families were exposed to the cap, while families that were already capped experienced additional reductions in their incomes. We therefore examine the mental health of those at risk of being capped before and after November 2016. We compare those at-risk of being capped with those who experienced a lower risk.

Figure 1 Number of households subject to the benefit cap between 2015 and 2018

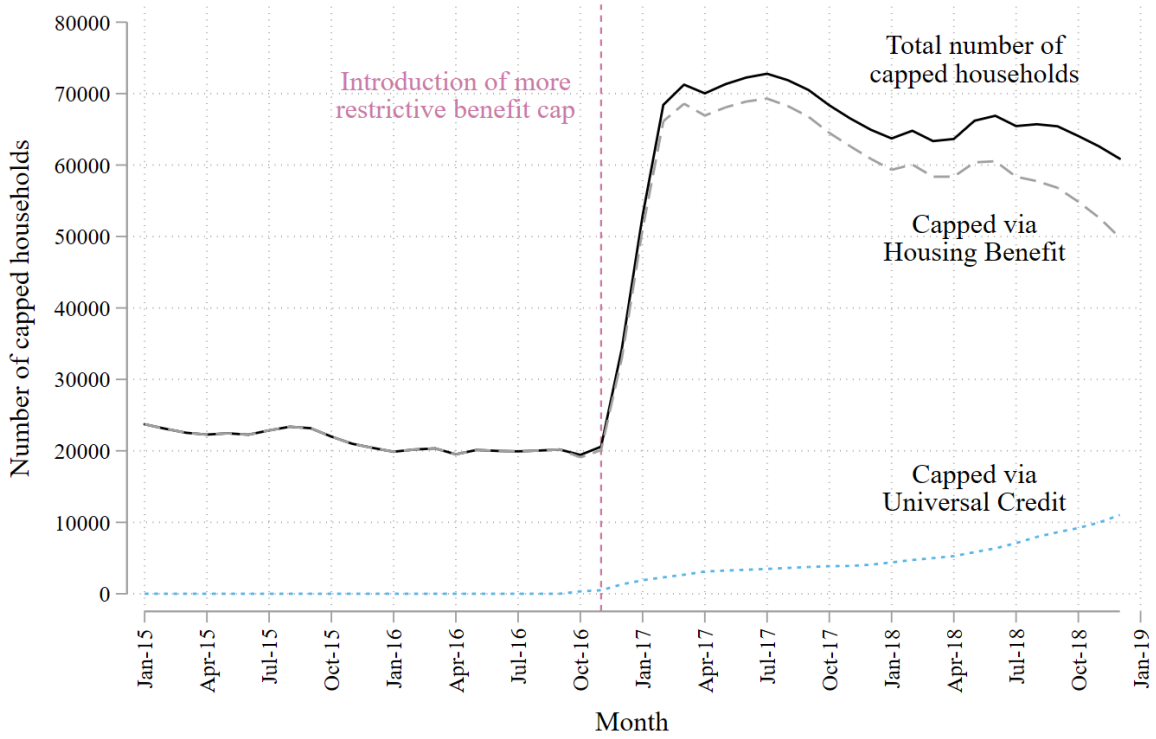

Source: Office for National Statistics. People can be capped in two different ways, both represented on the graph. Most people are capped through their Housing Benefit (the financial support they receive for their housing) but others are capped through Universal Credit, a new benefit system which combines out-of-work support, tax credits, and housing benefit.

We estimate an OLS difference-in-differences model with the following specification:

$$Health_{i,t} = \alpha_{i,t} + \beta_1 Capped_{i,t} + \beta_2 Policy_{i.t} + \beta_3 Capped \times Policy_{i,t} + \beta_z X_{i,t} + \varepsilon_{i,t}$$

Where *i* denotes individuals and *t* the time-period in which the data were collected. *Health* is a binary measure which is 1 if respondents report experiencing 'depression and anxiety' from a list of possible health problems and 0 otherwise. *a* is the constant (which in the model reports the probability of experiencing depression-like symptoms before the policy change and among those at a low risk of being capped). *Capped* is a dummy variable which is 1 if the respondent meets the criteria described above for being at risk of being capped and 0 otherwise. *Policy* is 1 if an individual was interviewed during or after November 2016 and 0 otherwise. Capped x 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),

gender (self-reported), 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, education (7 categories ranging from university degree to no qualifications), and their economic status (whether they were employed, unemployed, or economically inactive). ε is our error term.

Our coefficient of interest is β_3 – the difference-in-differences estimate. If $\beta_3 > 0$ then those at risk of being capped faced a higher probability of experiencing 'depression and anxiety' after the reform over and above any change in those with a lower risk of being capped.

We also conduct an interrupted time series analysis. This exploits the monthly data to test the parallel trends assumption and also enables us to see whether our results are sensitive to an alternative modelling strategy. We create a three-month moving average and then estimate whether the slopes before the reform are the same for those at-risk and low-risk of being capped. We also use this same model to test whether these trends diverge after November 2016, when the cap becomes more restrictive.

Sensitivity analyses

We conduct a falsification test to check whether our results are spuriously correlated with processes that should be unrelated to changes in the benefit cap. Here we use a measure of other health conditions as our dependent variable. Our theory is that in the short run these health conditions, many of which are physical conditions like experiencing difficulty in seeing or hearing, should be uncorrelated with the implementation of a more restrictive benefit cap. We also explore whether the changes to 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. This allows us to see whether the cap affected mental health more in some parts of the country than others.

We know that capped families are going to be dissimilar to non-capped families in ways that may affect their mental health and so we estimate a series of models which restrict the households in the 'control' group those that are more similar to the capped group than the unrestricted comparison group used in the main analysis. We focus on three contrasts, comparing families at risk of being capped with others who are not subject to the cap: (1) those who own their home or who are purchasing it through a mortgage (i.e., other large families), (2) those who are not receiving any benefits but who are renters and have the same household structure (i.e., other larger families in rented accommodation), and (3) those who are receiving

disability benefits and therefore exempt. We also explore a number of other restrictions to our data to examine whether the association is stable across various sub-groups. We re-estimate our main models among (1) only lone parents, (2) only households with more than two children, and (3) lone parents with more than two children.

Finally, one weakness of our main approach is that the LFS does not allow us to clearly identify those who are capped because the LFS does not contain a measure of total income from the government. The Family Resources Survey (FRS), by contrast, has more detailed measures of benefit receipt and therefore allows us to more accurately (albeit still imperfectly) assess who is being capped and who is not. We cannot use this dataset as our main source as it does not include measures of mental health, but as a sensitivity analysis we bring together the health data from the LFS with the more accurate benefit cap data from the FRS. To do this, we create a statistical model in the FRS data which predicts whether individuals are likely to be capped or not (see Supplementary analysis 2 for full details). We use this model to predict the probability of LFS respondents being capped and then re-estimate our models using this alternative measure of being at-risk of being capped. This enables us to see whether our results remain consistent across these alternative specifications.

3. Results

Did the benefit cap harm mental health among those at-risk of being capped?

We start by analysing the LFS data, comparing the probability of reporting depression-like symptoms among those at-risk of being capped and those with a low-risk, before and after the reform. Our most basic and unadjusted difference-in-differences model suggests that the prevalence of mental health problems increased among those who were at-risk of being capped after the reform was introduced and that this increase is greater than what we observe among the rest of the population. We visualise these difference-in-difference estimates in Figure 2. Our estimates suggest that among those who were at risk of being capped, the prevalence of mental ill health increased by around 2.6 percentage points (Model 1: Table 1). Adjusting for our covariates does not alter the main finding (see Model 2: Table 1) nor does excluding individuals that have experienced sustained mental health problems in the past. In other words, we are seeing an increase in the number of people experiencing depression for the first time.

Figure 2 Probability of reporting mental ill health increased more among those at-risk of being capped than everyone else after the benefit cap was introduced

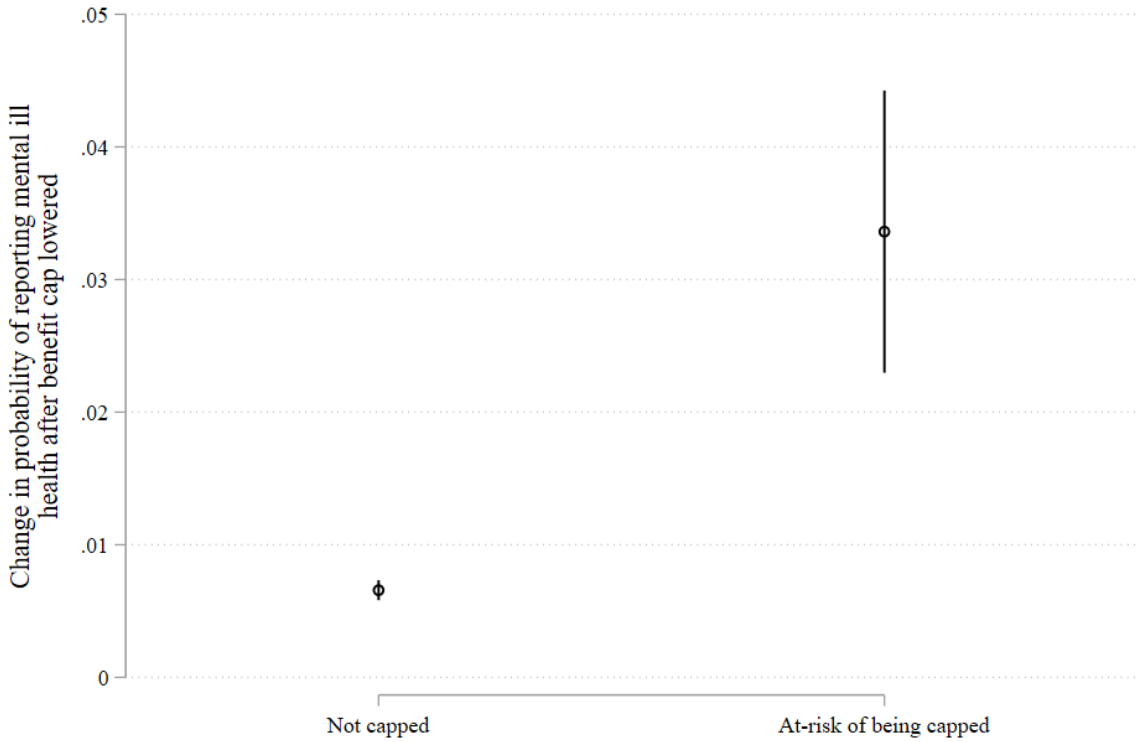

Notes: Figure based on results from Column 1 in Table 1. Data comes from the Labour Force Survey. Vertical lines represent 95% confidence intervals.

Table 1 The introduction of the benefit cap increased the prevalence of mental ill health

	Probability of reporting mental health problems		
	(1)	(2)	(3)
Difference-in-differences: Capped individuals compared to uncapped individuals after the reform	0.026** (0.0065)	0.023** (0.0061)	0.024** (0.0061)
Change over time for the non-capped individuals	0.010** (0.00056)	0.011** (0.00052)	0.011** (0.00052)
Difference between capped and non-capped individuals at baseline	0.14** (0.0042)	0.031** (0.0040)	0.031** (0.0040)
Constant (probability of depression among non-capped individuals before cap lowered)	0.069** (0.00039)	-0.090** (0.0046)	-0.089** (0.0046)
Controls for covariates		Y	Y
Restrict to those who have never had mental health problem			Y

Number of individuals	900506	900481	898294

Notes: Standard errors are reported in parentheses. Data comes from the Labour Force Survey. * p<0.05, ** p<0.01.

Next, we unpack this difference-in-differences estimate by estimating these models not only as a simple before-and-after but by examining when precisely the differences emerge. As we suggested above, we would not expect mental health problems to increase immediately after the reform was introduced. Rather, we expect this to become apparent slowly as the reduction in incomes starts to bite into family budgets and transitionary protections such as Discretionary Housing Payments begin to wane. We examine this by estimating the prevalence of mental ill health in every quarter of the data, before and after the reform. Figure 3 shows the results. We find that the level of mental ill health before the reform among those potentially affected by the cap was relatively stable, but after the reform was introduced we see a steady rise in the proportion of people reporting mental ill health. While (as Table 1 suggests) there has been a small increase in the number of people reporting mental ill health among those with a low-risk of being capped, this has been far less pronounced.

We next consider whether there is regional variation in the impact of the benefit cap (Supplementary analysis 3). The cap might be expected to have a larger impact on mental health in areas of the country where housing costs are higher, because the average reduction in income is likely to be higher. On the other hand, the reduction in incomes may be more keenly felt in poorer parts of the country, because the relative reduction might be larger. Even in London, over half (~56%) of capped families lose less than £50 per week and this is similar to the North West (59%), suggesting that this reduction might have a larger relative impact in less affluent parts of the country. Our results suggest that the mental health effects of the cuts may have been slightly larger in wealthier parts of the country, where the number of people affected is greater, but this difference is not statistically significant at conventional levels (p = 0.15).

Figure 3 Introduction of the benefit cap and the prevalence of mental ill health among those who are at-risk of being capped and those who are not, by quarter

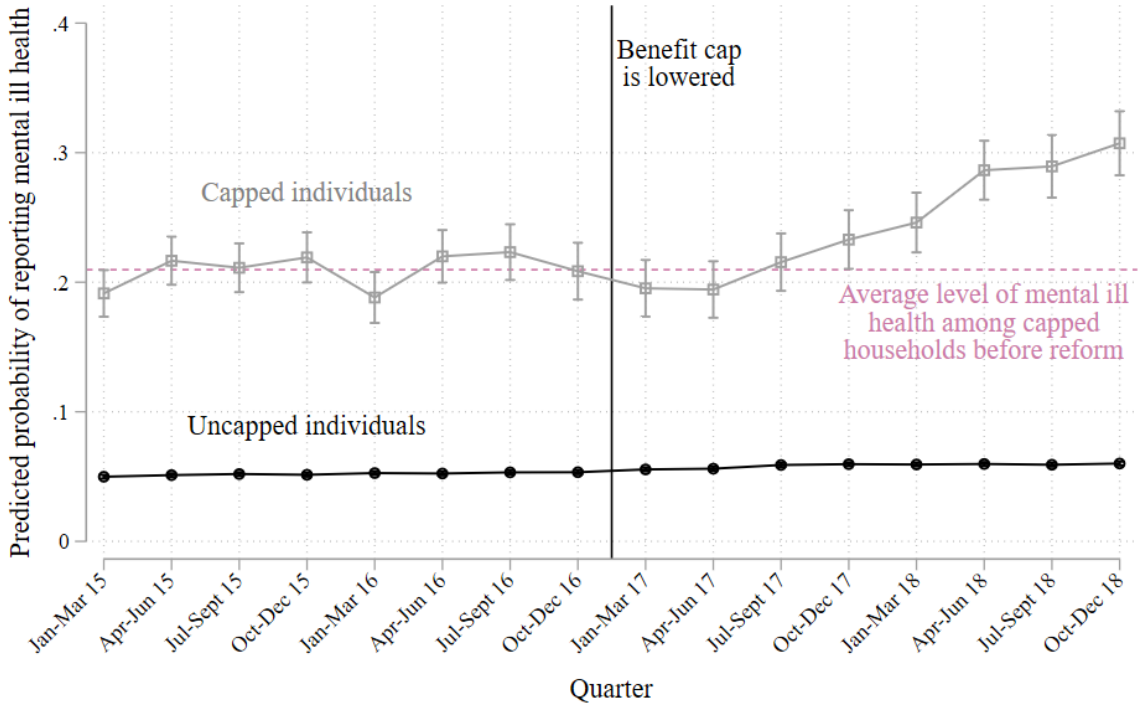

Notes: Data comes from the Labour Force Survey. Vertical lines represent 95% confidence intervals. Vertical black line indicates when the Benefit cap was lowered.

Finally, we re-estimate our main finding but now use an interrupted time series analysis. Here, we calculate the prevalence of mental ill health for both those at –risk of being capped and everyone else for every month from Jan 2015 to Dec 2018. We then calculate a three-month rolling average from these monthly observations. This analysis allows us both to test the parallel trends assumption and to examine whether there is a change in the slopes post reform. The time series models confirm the main findings from our quarterly analysis (Figure 4). We find no difference in the slopes prior to the reform but there is a very clear break in the slope after the reform was introduced, suggesting the trajectory for mental ill health clearly deviates from the rest of the population after the benefit cap becomes more restrictive (Supplementary analysis 4 for full results).

Figure 4 Interrupted Time Series Analysis of the lowering of the benefit cap

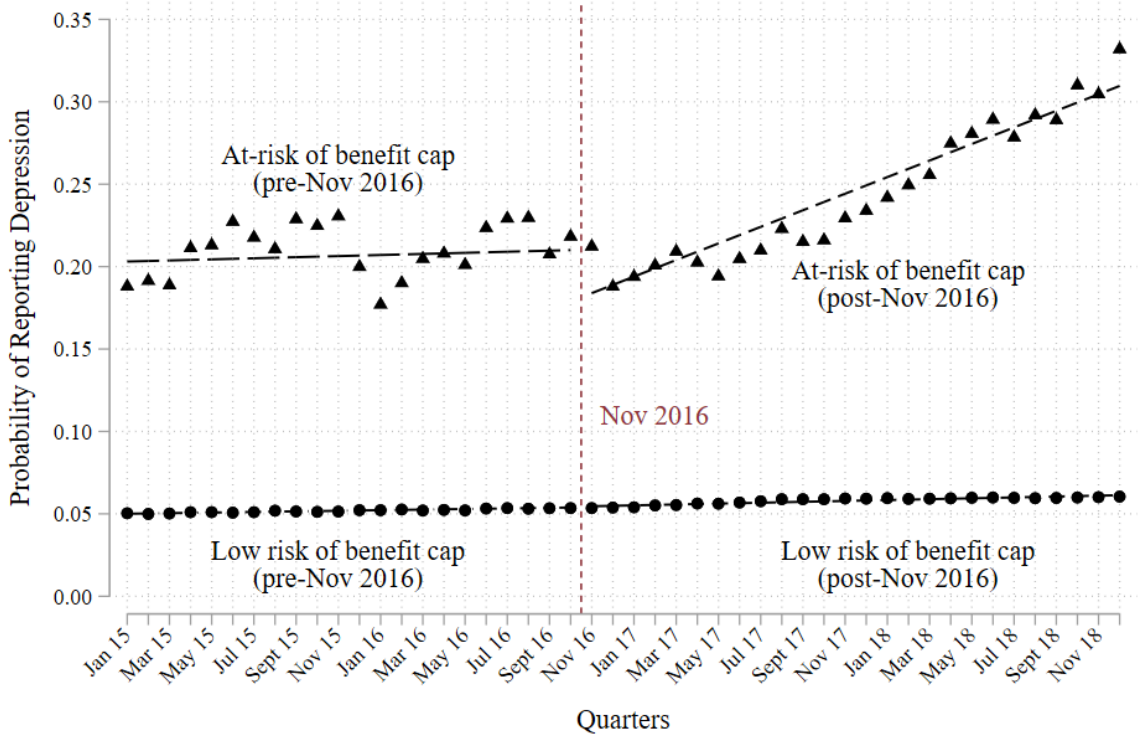

Notes: Data comes from the Labour Force Survey. Each dot represents the 3-month moving average of the probability of reporting depression. The lines of best fit are extrapolated from the Interrupted Time Series Analysis reported in Supplementary analysis 3.

Sensitivity tests

Our first sensitivity test is to conduct a falsification test. We would not expect the benefit cap to affect other non-mental health related outcomes over this time-period and so test whether we find a similar increase in other health problems in our difference-in-differences models. Supplementary analysis 5 shows that we find no association between those at risk of being capped and other non-mental health outcomes after the cap was lowered, suggesting our findings are not driven by compositional shifts unaccounted for by our variables nor by some other spurious trend.

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 on those that are more similar to the capped group than the unrestricted comparison group used in the main analysis. We focus on three contrasts, comparing families at risk of being capped with others who are not subject to the cap: (1) those who own their home or who are purchasing it through a mortgage (i.e., other large families), (2) those who are not receiving any benefits but who are renters and have the same household structure (i.e., other larger families in rented accommodation),

and (3) those who are receiving disability benefits and therefore exempt. Each model is in the same direction and the difference-in-differences estimates are of approximately the same size (see Supplementary analysis 6 for more details).

Third, we also explore a number of other restrictions to our data to examine whether the association is stable across various sub-groups. We re-estimate our main models among (1) only lone parents, (2) only families with more than 2 children, and (3) lone parents with more than 2 children. In each case, we find that being at high risk of being capped increases the prevalence of mental ill health (Supplementary analysis 7).

Finally 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 mental health in a variety of ways, described in detail in Supplementary analysis 2. Here, we report briefly the results of using a cut-off threshold in the probability of being capped. We assume the capped individuals are those who have an estimated probability of being affected by the benefit cap greater than 0.1 (for details of how this optimal cut-off was chosen, see Supplementary analysis 8). We conduct a difference-in-difference analysis similar to that presented in figure 1; the results, shown in in figure 5, are very similar in that they indicate a statistically significant increase in the probability of mental health problems for people at risk of being capped, after the introduction of the lower cap. The identified effect is somewhat larger than in figure 1, which would be expected if our FRS model were able to more precisely identify those at risk of being capped, and suggests that, if anything, our main estimates are conservative. We also run a series of additional models to test the consistency of our results and these are reported in Supplementary analysis 8. All the results from the FRS model are consistent with our findings with the simpler risk model from the LFS.

Figure 5 Probability of reporting mental ill health increased more among those at-risk of being capped than everyone else after the benefit cap was introduced, using the Family Resources Survey to identify the at-risk group

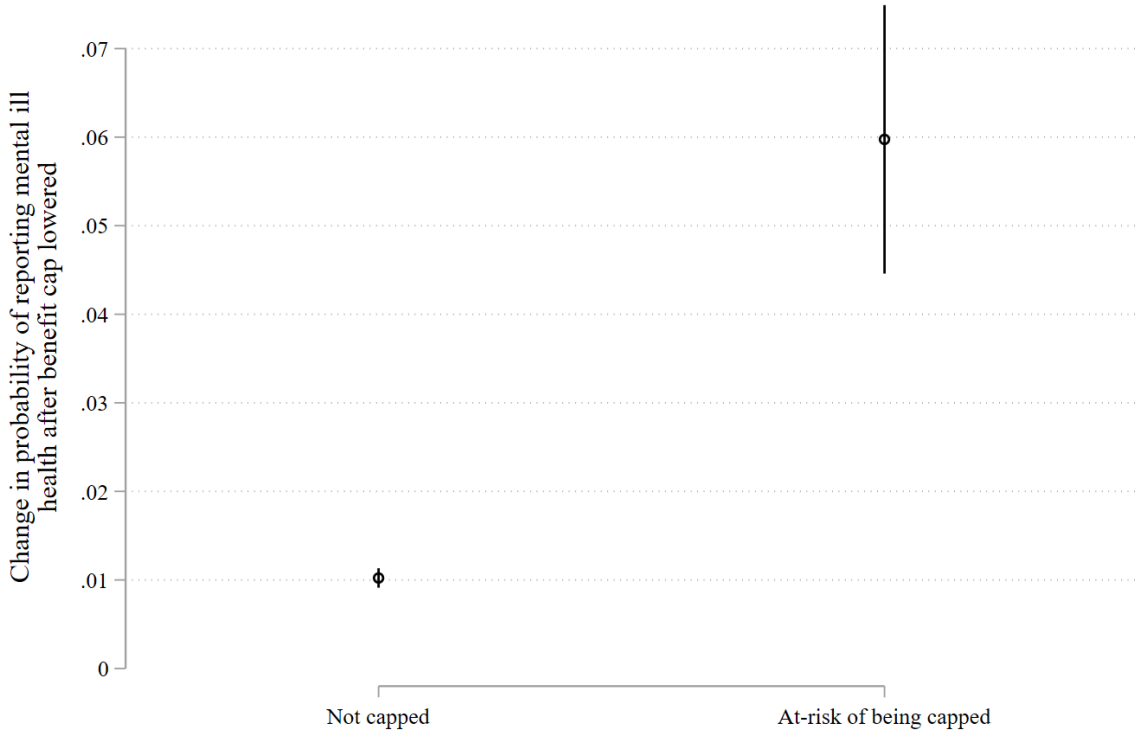

Notes: Figure based on results from Supplementary analysis 8. Data comes from the Labour Force Survey and the Family Resources Survey. Vertical lines represent 95% confidence intervals.

4. Discussion

We exploit the timing of the introduction of the more restrictive benefit cap in the UK in 2016 to examine the mental health effects of this reduction in total income on well-being. Two key findings emerge from our analysis. First, we find that the policy increased the risk of experiencing depressive-like symptoms among those affected by the cap. Second, we find that these negative effects on mental health emerge over a number of months. By the end of our study period, the risk of experiencing mental ill health among those at-risk of being capped had increased by around 10 percentage points, a relative increase of around 50%. To put this into perspective, in November 2019 there were ~76,000 households being capped (DWP, 2020b). Our estimates suggests that at least 16,000 people (~21%) in these households would have been living with depression-like symptoms if the benefit cap had remained unchanged, with around 6,600 additional people (an additional 9%) experiencing depressive-like symptoms as a result of the lowering of the cap. This is a conservative estimate because

this number only includes those being capped since 2016, and not those who were exposed to the cap before it was lowered. It may also underestimate the spill-over effects onto other family members (such as partners) whose mental health may also be affected by the impact of the cap but which are less reliably captured in the survey.

One broader implication of our results concerns the mental health effects of income shocks brought about by welfare reform. While the impact of income on physical health is still contested (Gunasekara et al., 2011), there is growing evidence using quasi-experimental designs that increases in income can lead to reductions in depression and anxiety, especially in low-income families (Cooper and Stewart, 2015, 2017). There is less evidence on the impact of *cuts* rather than increases in social security. Our study provides support for the hypothesis that reductions in income can be harmful to mental health in the short-term.

But our results are also significant for what they tell us about the nature of this particular reform. The cap most affected lone parents whose total income was close to median earnings for in-work families. The explicit assumption underlying the policy was that this provided a reasonable and adequate 'minimum income' that would still allow "people… to take responsibility for their decisions in the light of what they can afford" (Freud, 2013). Our results provide evidence that capped income is not in fact adequate. In practice, the families affected by the cap are very often living in relative poverty simply because of their household composition, and in addition face high housing costs. That is, while many believed that those capped received a non-trivial sum of money from the government in absolute terms, these families were still finding it difficult to make ends meet because of high spending needs. In this respect, the concerns of some charities and religious leaders have been borne out (CPAG, 2012; Evening Standard, 2010). Pushing capped families even further into poverty has negatively affected well-being.

If the cap is harming the mental health of parents (and particularly mothers) then this reform will have cascading effects on children's outcomes too (Wickham et al., 2017). Providing more income to low-income families has been found to improve both cognitive and social and behavioural development in children (Dahl and Lochner, 2012; Khanam and Nghiem, 2016; Manley et al., 2015). At the more extreme end, reductions in welfare payments have been linked to increased child maltreatment (Cancian et al., 2013). Parental well-being, and particularly maternal anxiety and depression, appear to be a key part of the mechanism linking poverty to children's outcomes (Cooper, 2017; Garfinkel et al., 2016). Thus implementing the benefit cap may deliver short-run savings and may even induce some families into work, but there may be long-run consequences on the life chances of the children in these families.

Our findings also have implications for the ability of the cap to successfully incentivise labour market activity, one of the policy's central goals. Observational evidence indicates that those affected by the cap were initially pursuing employment with more energy as a result of the policy, and that capped individuals did indeed move into work at a higher rate than similar individuals who were not capped (DWP, 2014). Our results do not contradict these earlier findings but instead suggest that the cap may have had heterogeneous effects. The cap may increase job search activity and even re-employment for some of those affected, while at the same time resulting in a non-trivial number of people experiencing poorer mental health. We cannot perform a formal cost-benefit analysis in part because we cannot estimate the employment effects in our data while employment can be good for health, not all jobs improve well-being compared to being unemployed (Broom et al., 2006; Butterworth et al., 2013); and so some of those who are experiencing poorer mental health may now be in work. However, we can still put these different effects into dialogue with each other. Estimates of the re-employment effects range from 3.5 percentage points to 4.7 percentage points while the increase in mental ill health is, on average, around 2.6 percentage points but may be as high as 9 percentage points by the end of the period. It is difficult to draw strong conclusions from this comparison but the mental health effects are non-trivial compared to the employment effects, especially given that many of those experiencing poorer mental health will still be out of work and so the effect of the cap could be to push them even further away from the labour market (García-Gómez et al., 2010). Indeed, some of those now in work may exit more quickly because of their health while others may be less likely to get work again in the near future.

These findings have implications for other countries beyond the United Kingdom. In countries like the United States, where social security has already adopted some of the principles underlying this new logic of social security, (fairly) hard limits on social security spending have already been implemented in the form of block grants to states, which limit the amount that can be spent on support for low-income families. These limits are not targeted at specific groups, however, and have often been criticised as a result. A benefit cap like the one implemented in the UK may be a more palatable (and potentially even popular) way to achieve the same goal. Policy transfer between these contexts is not new: Britain adopted some of the principles underlying Clinton's welfare reforms, and these in turn borrowed from British debates going back to the implementation of the 'poor laws' in Britain in the nineteenth century (Somers and Block, 2005). Our results suggest other countries should be cautious about adopting the basic structure of the benefit cap. The social harms created by this system should give policymakers pause, especially in countries where the rate of mental ill health among low-income families is already high.

There are, of course, important limitations to our analysis. First, our self-reported measure of mental health is not a clinical diagnostic tool and while the health measure we use is highly correlated with more formal measures of mental health there is still a possibility of measurement error in our dependent variable. This is unlikely to materially affect our results, of course, unless the policy change simultaneously affected how people responded to this question (Reeves et al., 2016). Second, while we have used multiple control groups to test the robustness of our results, the absence of a true experiment means all of our control groups are less than ideal (Sekhon, 2009). Third, more work is needed to understand exactly how and why implementing a more restrictive cap harmed mental health, and this will likely require qualitative research that can trace the experiential aspects of the cap. Finally, neither of our data sets allow us to perfectly identify those who are capped. We address this limitation by adopting an intention-to-treat approach, which is likely more conservative in this instance (Gupta, 2011), and by conducting a series of sensitivity analyses which explore different ways of identifying those at risk of being capped. These reveal that our main estimates may in fact be conservative and that the impact of the benefit cap could be even larger.

The benefit cap reminds us of the unintended consequences of policy and illustrates how they can exacerbate inequalities. The cap not only increases the risk of mental ill health but it does so among lone parents (usually women) who live in high-rent areas. In this respect, our results reinforce other work which shows how lone parents have been particularly badly hit by recent welfare reforms and subsequent economic hardship, for example being disproportionately likely to be sanctioned (have their welfare payments temporarily stopped) and to be evicted (Desmond, 2012; Reeves and Loopstra, 2017). One of the troubling aspects of the rise of this new logic of social security is how often the policies which flow from it seem to penalise lone parents and their children (Gregg et al., 2009; Knijn et al., 2007). In this instance, our analysis suggests attempts to promote labour market activation and to reduce costs through the benefit cap have had the adverse consequence of damaging the mental health of those exposed to this reform.

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

References

Broom, D.H. et al. (2006) 'The Lesser Evil: Bad Jobs or Unemployment? A Survey of Mid-Aged Australians', *Social Science & Medicine* 63(3): 575–86.

Butterworth, P., L.S. Leach, S. McManus, and S.A. Stansfeld (2013) 'Common Mental Disorders, Unemployment and Psychosocial Job Quality: Is a Poor Job Better than No Job at All?', *Psychological Medicine* 43(8): 1763–72.

Campbell, M. (2020) *Capping Motherhood: An Equality-Based Analysis of the UK Benefit Cap Cases* (SSRN Scholarly Paper No. ID 3587480).

Cancian, M., M.-Y. Yang, and K.S. Slack (2013) 'The Effect of Additional Child Support Income on the Risk of Child Maltreatment', *Social Service Review* 87(3): 417–37.

Cooper, K. (2017, September) 'Poverty and Parenting in the UK'. PhD, The London School of Economics and Political Science (LSE).

Cooper, K. and K. Stewart (2015) *Does money in adulthood affect adult outcomes?*

Cooper, K. and K. Stewart (2017) 'Does Money Affect Children's Outcomes? An Update'.

CPAG (2012, January 23) 'Benefit Cap Myths Exposed', *CPAG*. https://cpag.org.uk/news-blogs/news-listings/benefit-cap-myths-exposed (accessed 9 September 2020).

Dahl, G.B. and L. Lochner (2012) 'The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit', *American Economic Review* 102(5): 1927–56.

Desmond, M. (2012) 'Eviction and the Reproduction of Urban Poverty', *American Journal of Sociology* 118: 88–133.

Dunning, T. (2012) *Natural Experiments in the Social Sciences*, Strategies for Social Inquiry. Cambridge: Cambridge University Press.

DWP (2014) *The benefit cap: a review of the first year* (No. Cm 8985).

DWP (2020a) *Benefit cap: number of households capped to May 2020*.

DWP (2020b) *Benefit Cap: Data to November 2019*.

Eikemo, T.A. and C. Bambra (2008) 'The Welfare State: A Glossary for Public Health', *Journal of Epidemiology & Community Health* 62(1): 3–6.

Evening Standard (2010, November 9) 'Bishop of Southwark: Vulnerable at Risk from Housing Benefit Cap'. http://www.standard.co.uk/news/bishop-of-southwark-vulnerable-at-risk-from-housing-benefit-cap-6533954.html (accessed 9 September 2020).

Fenton-Glynn, C. (2015) 'Austerity and the Benefit Cap: In Whose Best Interests?', *Journal of Social Welfare and Family Law* 37(4): 467–9.

Finlay, S., B. Marshall, and L. Hill (2013) *The Benefit Cap: Public perceptions and pre-implementation effects* (No. 850).

Freud, D. (2013) *HL Deb (13 September 2011) Vol 730, Col 739*.

García-Gómez, P., A.M. Jones, and N. Rice (2010) 'Health Effects on Labour Market Exits and Entries', *Labour Economics* 17(1): 62–76.
Garfinkel, I., S.S. McLanahan, and C. Wimer (eds) (2016) *Children of the Great Recession*. New York: Russell Sage Foundation.
Gregg, P., S. Harkness, and S. Smith (2009) 'Welfare Reform and Lone Parents in the Uk', *Economic Journal* 119(535): F38–65.
Grogger, J. (2002) 'The Behavioral Effects of Welfare Time Limits', *American Economic Review* 92(2): 385–9.
Gunasekara, F.I., K. Carter, and T. Blakely (2011) 'Change in Income and Change in Self-Rated Health: Systematic Review of Studies Using Repeated Measures to Control for Confounding Bias', *Social Science and Medicine* 72(2): 193–201.
Gupta, S.K. (2011) 'Intention-to-Treat Concept: A Review', *Perspectives in Clinical Research* 2(3): 109–12.
Hills, J. (2011) 'The Changing Architecture of the UK Welfare State', *Oxford Review of Economic Policy* 27(4): 589–607.
Hills, J. (2014) *Good Times, Bad Times: The Welfare Myth of Them and Us*. Bristol: Policy Press.
Hollingsworth, K. (2015) 'Judging Children's Rights and the Benefit Cap: R (SG and Others) v Secretary of State for Work and Pensions Case Commentary', *Child and Family Law Quarterly* 27(4): 445–66.
Hussain, M.A., M. Ejrnæs, and J.E. Larsen (2020) 'Are Benefit Reductions an Effective Activation Strategy? The Case of the Lowest Benefit Recipients in Denmark', *Journal of Social Policy*: 1–19.
James, G., D. Witten, T. Hastie, and R. Tibshirani (2013) *An Introduction to Statistical Learning: With Applications in R*. Springer Science & Business Media.
Jones, A.M. and J. Wildman (2008) 'Health, Income and Relative Deprivation: Evidence from the BHPS', *Journal of Health Economics* 27(2): 308–24.
Kaiser Family Foundation (2017, November 27) 'Facilitating Access to Mental Health Services: A Look at Medicaid, Private Insurance, and the Uninsured', *KFF*. https://www.kff.org/medicaid/fact-sheet/facilitating-access-to-mental-health-services-a-look-at-medicaid-private-insurance-and-the-uninsured/ (accessed 8 July 2020).
Kaur, K., R. Mosley, T. Nutley, and J. Thompson (2014) *Benefit Cap: Analysis of outcomes of capped claimants* (No. 11).
Khanam, R. and S. Nghiem (2016) 'Family Income and Child Cognitive and Noncognitive Development in Australia: Does Money Matter?', *Demography* 53(3): 597–621.
Knijn, T., C. Martin, and J. Millar (2007) 'Activation as a Common Framework for Social Policies towards Lone Parents', *Social Policy and Administration* 41(6): 638–52.
Lupton, R., J. Hills, K. Stewart, T. Burchardt, and P. Vizard (2016) *Social Policy in a Cold Climate: Policy, Poverty and Inequality in England*. Policy Press.

Mackenbach, J.P. et al. (2005) 'The Shape of the Relationship between Income and Self-Assessed Health: An International Study', *International Journal of Epidemiology* 34(2): 286–93.
Manley, J., L. Fernald, and P. Gertler (2015) 'Wealthy, Healthy and Wise: Does Money Compensate for Being Born into Difficult Conditions?', *Applied Economics Letters* 22(2): 121–6.
Mulholland, H. (2010, October 27) 'David Cameron Refuses to Back down over Housing Benefit Cap', *The Guardian*.
ONS (2019) 'Families and Households in the UK'. https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/families (accessed 18 September 2020).
Osborne, G. (2010) *Spending Review 2010: Presented to Parliament by the Chancellor of the Exchequer by Command of Her Majesty*.
Pierson, P. (1994) *Dismantling the Welfare State? : Reagan, Thatcher, and the Politics of Retrenchment*. Cambridge: Cambridge University Press.
Reeves, A., A. Clair, M. McKee, and D. Stuckler (2016) 'Reductions in the United Kingdom's Government Housing Benefit and Symptoms of Depression in Low-Income Households', *American Journal of Epidemiology* 184(6): 421–9.
Reeves, A. and R. Loopstra (2017) '"Set up to Fail"? How Welfare Conditionality Undermines Citizenship for Vulnerable Groups', *Social Policy and Society* 16(2): 327–38.
Scruggs, L. and J.P. Allan (2006) 'The Material Consequences of Welfare States - Benefit Generosity and Absolute Poverty in 16 OECD Countries', *Comparative Political Studies* 39(7): 880–904.
Sekhon, J.S. (2009) 'Opiates for the Matches: Matching Methods for Causal Inference', *Annual Review of Political Science* 12(1): 487–508.
Somers, M.R. and F. Block (2005) 'From Poverty to Perversity: Ideas, Markets, and Institutions over 200 Years of Welfare Debate', *American Sociological Review* 70(2): 260–87.
Taylor-Gooby, P. and E. Taylor (2015) 'Benefits and Welfare: Long-Term Trends or Short-Term Reactions?', *British Social Attitudes*. London: NatCen.
Timmins, N. (2017) *The Five Giants: A Biography of the Welfare State* (Revised edition edition.). London: HarperCollins.
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).
Tucker, J. (2019) *Universal Credit: What needs to change to reduce child poverty and make it fit for families*.
Wickham, S., M. Whitehead, D. Taylor-Robinson, and B. Barr (2017) 'The Effect of a Transition into Poverty on Child and Maternal Mental Health: A Longitudinal Analysis of the UK Millennium Cohort Study', *The Lancet Public Health* 2(3): e141–8.

Work and Pensions Committee (2019) *The benefit cap* (Twenty-Fourth Report of Session 2017–19).
Ziliak, J.P. (2015) 'Temporary Assistance for Needy Families', *Economics of Means-Tested Transfer Programs in the United States, Volume 1*: 303–93.

Supplementary analysis

Supplementary analysis 1: Composition of those who are at-risk of being capped before and after the reform

Supplementary analysis 2: Method for modelling predicted probabilities of being capped in the FRS and estimating these for LFS respondents

Supplementary analysis 3: Impact of the benefit cap by region

Supplementary analysis 4: Interrupted Time Series Analysis of the impact of the benefit cap on mental health

Supplementary analysis 5: The introduction of the benefit cap did not increase the prevalence of other health outcomes (non-mental health)

Supplementary analysis 6: Exploring whether those at risk of being capped experience worse mental health than those who are exempt

Supplementary analysis 7: The introduction of the benefit cap and the prevalence of mental ill health among lone parents, larger households (more than 2 children), and lone parent households with more than 2 children.

Supplementary analysis 8: Estimating the effects of benefit cap using information from the LFS and the FRS

Supplementary analysis 1: Composition of those who are at-risk of being capped before and after the reform

		Among those at risk of being capped			Compared to whole population
Variable	Administrative data on Capped Households	Before Benefit cap made more restrictive	After Benefit Cap made more restrictive	Test of difference in means: After – Before (p-value)	Difference-in-difference (p-value)
Gender (female=1)	92%[1]	94.77%	95.46%	0.572	0.549
Age	33.01[2]	32.86	33.54	0.226	0.790
Ethnicity (White British = 1)	--	81.81%	81.34%	0.576	0.613
London	26.14%[3]	15.22%	12.78%	0.001	0.006
Health problems (not depression)	--	25.87%	26.04%	0.888	0.757
Lone parent	71.8%[3]	97.12%	96.54%	0.575	0.147
Number of children	3.07[2]	2.09	2.13	0.165	0.015
Children present	93%[3]	98.99%	99.12%	0.593	0.918
Single	79.4%[3]	72.50%	74.23%	0.156	0.123
Unemployed (ILO definition)	--	23.17%	17.21%	<0.001	<0.001
Economically inactive	--	72.92%	74.62%	0.028	0.030

Notes: Economically inactive people are those who are not in work and who are not unemployed according to the ILO's definition (which includes people who are not in work but who are actively seeking employment).

1 - This figures refers to single claimants only and apply to those capped through their Housing Benefit in November 2019.

2 – This figure refers to those capped through their Housing Benefit in November 2019.

3 – This figure refers to those capped through Universal Credit and Housing Benefit in November 2019.

Supplementary analysis 2: Method for modelling predicted probabilities of being capped in the FRS and estimating these for LFS respondents

The main limitation of the Labour Force Survey (LFS) is that it does not contain information on the income received from government cash transfers. The Family Resources Survey (FRS), by contrast, is a household survey that collects detailed data from respondents about their income from all sources. The FRS does not, however, have a measure of mental health. We therefore attempt to combine the information contained in both the LFS and the FRS to check our main results from the LFS.

For the purposes of UK government cash transfers, income is assessed at the 'benefit unit' level which is what we follow in this analysis. A benefit unit consists of one adult or two adults in a relationship, plus any dependent children (defined as under 16 years of age, or under 19 years of age and in full-time education) – also known in official demographic statistics as a 'family'. The cash transfer income of a benefit unit is capped if their income from all benefits that are included in the cap exceed the specified level.

We pool data from the 2013/14, 2014/15 and 2015/16 iterations of the cross-sectional FRS which include interviews that took place between April 2013 and March 2016 i.e. in the period after the initial benefit cap was introduced but before the more restrictive cap was implemented in November 2016. We pool data in this way to increase the number of cases identified as being affected by the cap, given the small proportion (approximately 0.24%) of UK benefit units affected[2] (ONS, 2019). For each benefit unit in the sample we sum the total income received from the cap-affected benefits, and can thus identify benefit units with benefit income that is above the November 2016 cap threshold. The FRS data also allows us to identify whether benefit units are exempt from being affected by the cap i.e. whether they are claiming any of the disability-related 'exemption' benefits, in sufficient hours paid employment or had worked within the past year. We are thus able to identify benefit units who would have been affected by the more restrictive cap, had it been in place before November 2016.

Using this information we develop a predictive model for whether a benefit unit is affected by the benefit cap by using variables that are both common across the LFS and FRS datasets and are related to the risk of being affected by the benefit cap. These common variables include economic status, housing tenure, number of dependent children, household composition, ethnic group, age, and government region. We develop the predictive model for benefit units who have a non-zero probability of being affected

[2] Authors' calculation from published official statistics. There were 29.0 million benefit units in Great Britain in mid-2017 (ONS, 2019) of which 68,900 were subject to the benefit cap, or 0.24%.

by the benefit cap (those renting their home, with dependent children, not in paid employment). We use a logistic regression for our predictive model, using 10-fold cross-validated measures of prediction error to avoid overfitting to the FRS dataset (James et al., 2013).

The final step in our predictive model is to apply this predictive logistic regression model to the LFS data for 2017 to 2018 which gives us, for each LFS case, a predicted probability of being affected by the more restrictive post-2016 benefit cap. The key assumption is that the characteristics of households who received benefits above the benefit cap level before the cap came into place is the same as those households who were actually capped when the policy came into being. This seems like a reasonable assumption given the limited behavioural change that the cap appeared to induce (Kaur et al. 2014).

How well did this procedure work?
The number of benefit units affected by the benefit cap, according to our calculations from the FRS, was broadly in line with that reported by national statistics. In the 2015/16 FRS sample 55 of a total of 22,540 observations were identified as having relevant cash benefit incomes above the 2016 cap threshold. Using survey weights this equates to an estimate of 64,700 benefit units affected (95% CI: 46,500-83,000), which is consistent with the 70,000 reported by official government statistics in April 2017 (see figure 1).

The results of the logistic regression model for predicting whether a benefit unit is at risk of being affected by the benefit cap are shown in table 2.1. The results are as expected: benefit units in private rented rather than social rented housing are at greater risk of the benefit cap due to the higher rents they face, as are those living in the higher cost London and South East England regions. Benefit units with a larger number of children and those headed by someone from a minority ethnic group are also at greater risk.

The cross-validated prediction error for this model was 0.08 (i.e. 8% of cases were incorrectly classified by the model). Alternative models adding interaction terms between the prediction variables were also explored but did not significantly alter the prediction error; therefore we chose to use this relatively parsimonious prediction model.

Table 2.1: Predictive model for being capped in the Family Resources Survey

	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	-5.918	0.8378	-7.064	1.614e-12
Rented from housing association	-0.5112	0.2576	-1.984	0.0472
Rented privately unfurnished	0.6891	0.2137	3.224	0.001262
Rented privately furnished	0.6812	0.4464	1.526	0.127

Couple, two children	0.5205	0.6684	0.7787	0.4361
Couple, three children	2.227	0.6211	3.586	0.0003361
Couple, four or more children	1.9	0.6614	2.873	0.004072
Lone parent, one child	-0.8839	0.7015	-1.26	0.2077
Lone parent, two children	0.6494	0.5961	1.089	0.276
Lone parent, three children	2.827	0.5769	4.9	9.584e-07
Lone parent, four or more children	3.903	0.5938	6.573	4.918e-11
Number of children under 5 years old	0.5012	0.117	4.285	1.824e-05
Aged 25-34 years	0.7073	0.4432	1.596	0.1105
Aged 35-44 years	0.8904	0.4733	1.881	0.05994
Aged 45-54 years	1.297	0.5422	2.392	0.01675
Aged 55-64 years	0.4396	1.188	0.37	0.7113
North West	0.02636	0.5022	0.05248	0.9581
Yorks & Humber	0.1206	0.5313	0.2269	0.8205
East Midlands	0.144	0.5322	0.2705	0.7867
West Midlands	0.1473	0.5105	0.2885	0.773
East of England	0.192	0.5504	0.3489	0.7271
London	1.393	0.4763	2.924	0.003454
South East	1.055	0.4947	2.134	0.03288
South West	-0.7002	0.7552	-0.9273	0.3538
Wales	0.3162	0.609	0.5192	0.6036
Scotland	-0.05271	0.5103	-0.1033	0.9177
Northern Ireland	-0.001862	0.4756	-0.003915	0.9969
Mixed / Multiple ethnic groups	-0.4972	0.8423	-0.5903	0.555
Asian / Asian British	0.7774	0.3711	2.095	0.03617
Black African / Caribbean / British	-0.2775	0.3434	-0.808	0.4191
Other ethnic group	1.473	0.4653	3.166	0.001546

Figure 2.2 shows the results of the predictive model for the 2,355 FRS cases with non-zero probability of being affected by the benefit cap. The predicted probabilities display a reasonable degree of differentiation between those cases that are affected by the benefit cap and those that are not. However there is a significant number of outlying cases that the predictive model identifies as having a high probability of being affected by the cap, whereas the FRS benefit claim data shows they are not in fact affected by the cap i.e. false positive predictions.

Figure 2.2: Predicted probability of being affected by benefit cap, by FRS benefit cap status

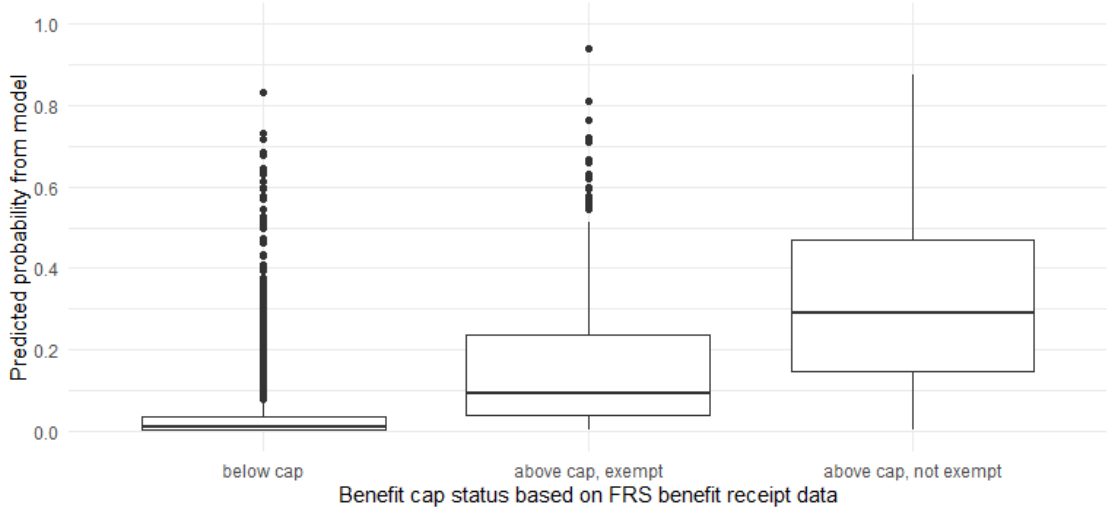

Figure 2.3 shows the predicted probabilities calculated by applying the predictive model to the 20,280 LFS observations that have non-zero probability of being affected by the benefit cap. Consistent with what we would expect, most cases have a very small probability of being affected by the benefit cap - the median probability is 0.02, the third quartile 0.09 and the maximum is 0.95. Once we transpose our model into the LFS, the model reveals that there is a small number of people with a high probability of being subject to the benefit cap.

Figure 2.3: Predicted probability of being affected by benefit cap for LFS cases

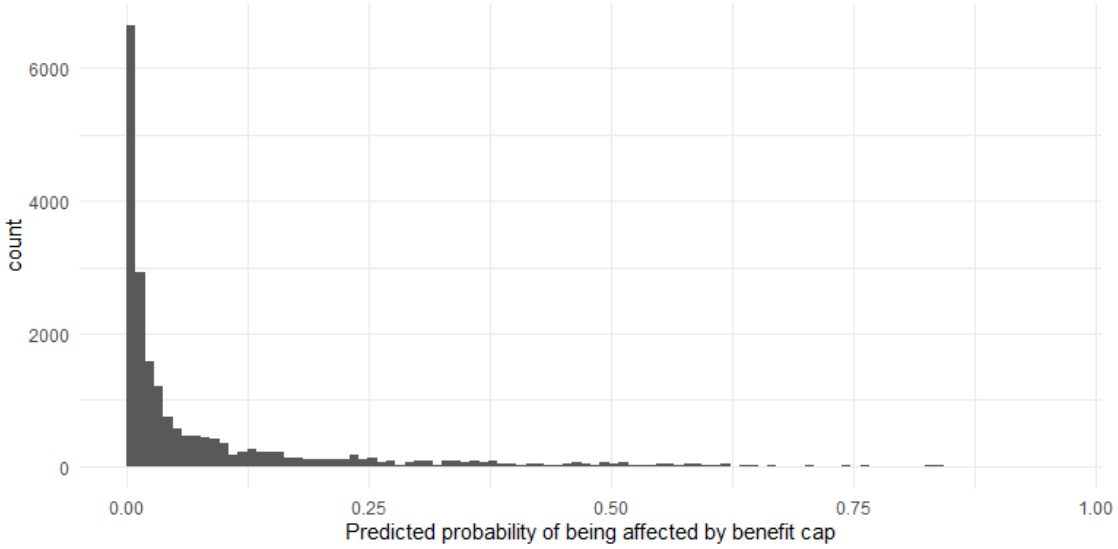

We show formal results from this approach in Supplementary analysis 7.

Supplementary analysis 3: Impact of the benefit cap by region

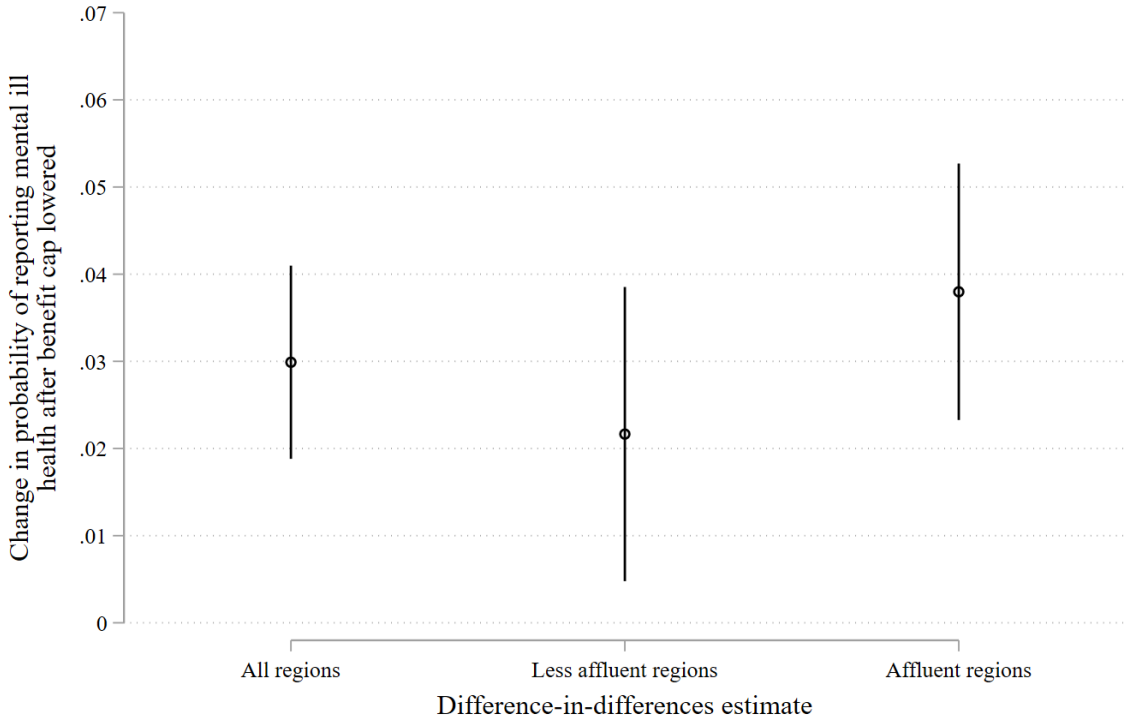

Supplementary analysis 4: Interrupted Time Series Analysis of the impact of the benefit cap on mental health

Below we report some of the parameters from our interrupted time series analysis, as visualised in figure 4. We find a small positive increase in the probability of experiencing depression in the control group prior to intervention. But we find no difference between this slope for the treated group, suggesting that although the levels of depression are different that the trends are parallel. Following the reform, there is a slight increase in the slope for the control group, suggesting that the probability of depression increased slowly after the benefit cap was introduced. However, we now observe a large divergence between these trends. After the benefit cap was lowered, the probability of depression increases much faster than it did before the reform for those at risk of being capped and, crucially, this slope is much steeper than the slope for the control group.

Outcome: Probability of reporting depression	Point estimate	p-value	95% CI
Pre-intervention slope for control group	0.00018	<0.001	0.00013 to 0.00023
Difference in pre-intervention slopes between treated and control group	0.00015	0.851	-0.0014 to 0.0017
Post-intervention slope for control group	0.000088	0.016	0.000017 to 0.00016
Difference in post-intervention slopes between treated and control group	0.0046	<0.001	0.050 to 0.051

Notes: Other parameters are estimated in the full-model but these are the essential point estimates to show the differences in the slopes before and after the trend.

Supplementary analysis 5: The introduction of the benefit cap did not increase the prevalence of other health outcomes (non-mental health)

	Probability of reporting other health problems		
	(1)	(2)	(3)
Difference-in-differences: Capped individuals after the reform	-0.0011 (0.012)	-0.0060 (0.012)	-0.0068 (0.012)
Change over time for the non-capped individuals	0.0047** (0.0009 5)	0.0050** (0.0009 1)	0.0051** (0.0009 1)
Difference between capped and non-capped individuals at baseline	0.010 (0.0078)	-0.059** (0.0075)	-0.059** (0.0075)
Constant (probability of depression among non-capped individuals before cap lowered)	0.25** (0.0006 6)	-0.22** (0.0080)	-0.22** (0.0080)
Adjusted for covariates		Y	Y
Restrict to those who have never had mental health problem			Y
Number of individuals	832300	832278	830584

Notes: Standard errors are reported in parentheses. Number of individuals is lower because all those with mental health issues have been excluded.

Supplementary analysis 6: Exploring whether those at risk of being capped experience worse mental health than those who are exempt

We focus on three exemptions.

1. We compare those at risk of being capped with those who are otherwise similar but who are not claiming social security benefits. This helps us see what is happening to renters with similar household compositions (namely lone parents and couples with 3 or more children).
2. Next we compare those at risk of being capped with those who own their home or who are currently buying it through a mortgage. This helps us see what is happening to people with a similar household structure and who are claiming benefits but who are not receiving housing benefit (the main driver of the benefit cap).
3. Finally, we compare those at risk of being capped with those who have a disability. This contrast again using this exemption to trace whether those who are capped experience a different trend in their well-being after the reform.

In all cases, we observe positive difference-in-differences estimates, suggesting that those at-risk of being capped are more likely to experience worse mental health after the more restrictive cap was implemented compared to those who were not exposed or who were exempt.

A. At risk of being capped compared to those claiming no benefits.

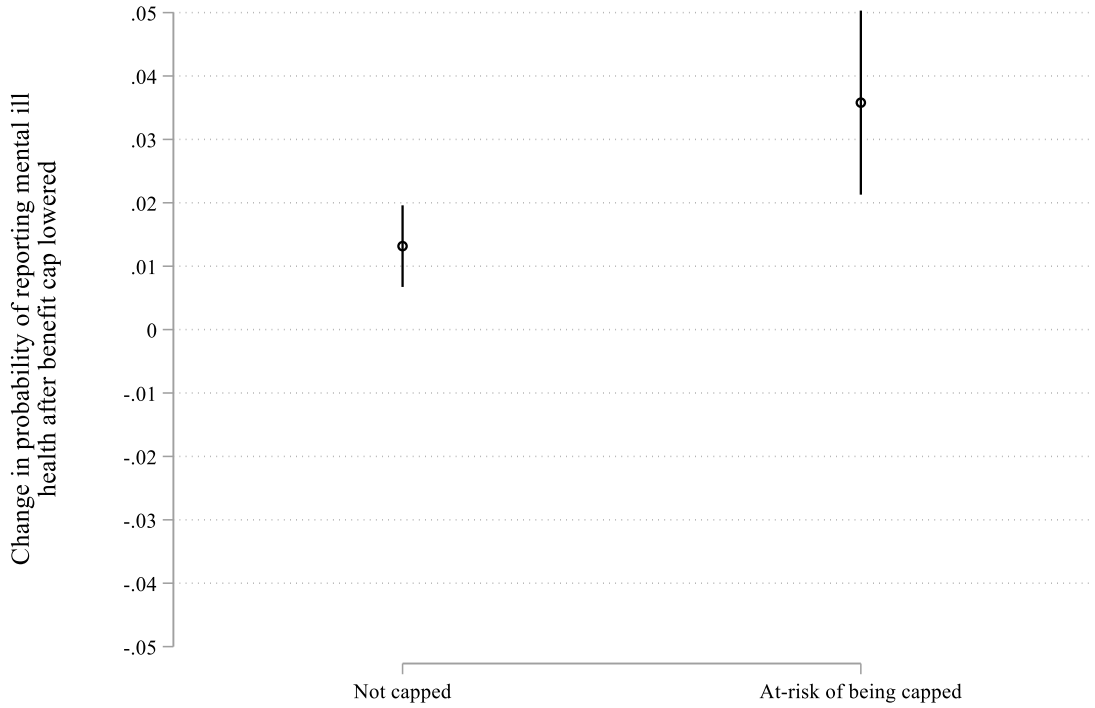

B. At risk of being capped compared to those with a mortgage or who own their homes

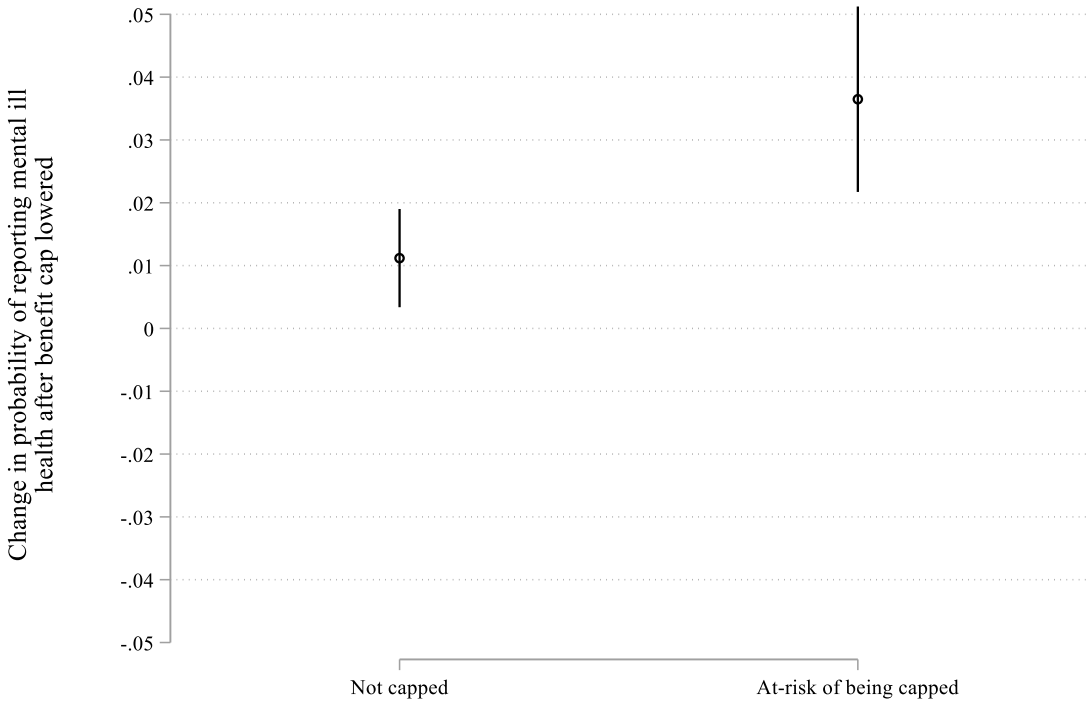

C. At risk of being capped compared to those who are disabled.

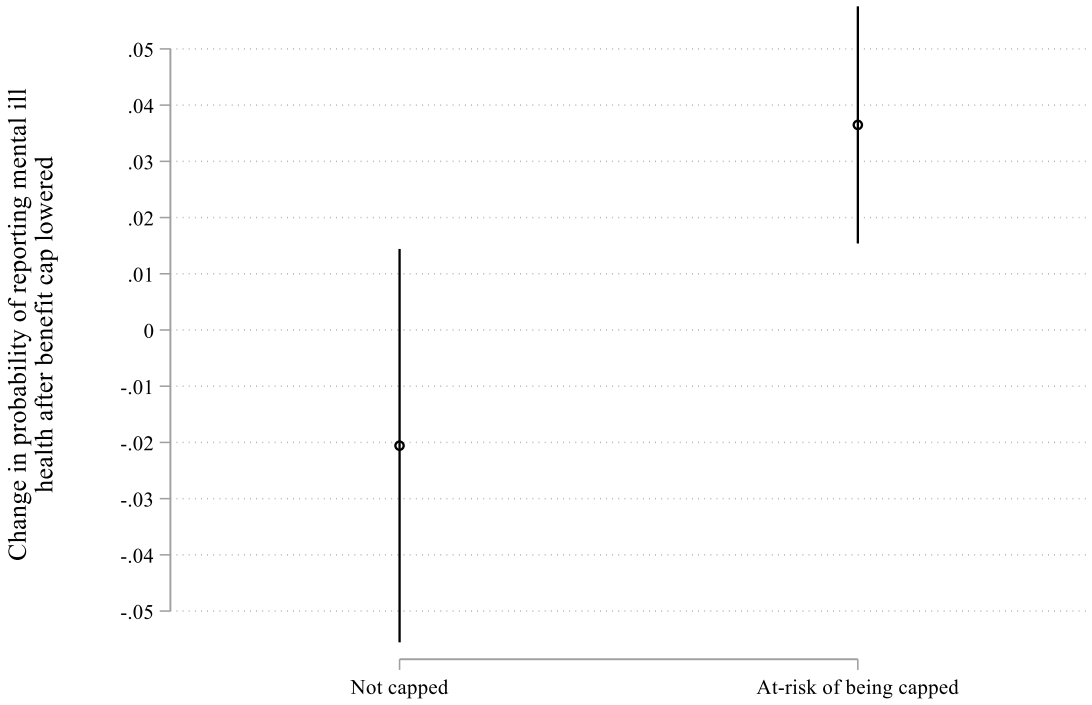

Supplementary analysis 7: The introduction of the benefit cap and the prevalence of mental ill health among lone parents, larger households (more than 2 children), and lone parent households with more than 2 children.

	Probability of reporting mental health problems		
	Lone parents	Households with more than 2 children	Lone parent households with more than 2 children
	(1)	(2)	(3)
Difference-in-differences: Capped individuals compared to uncapped individuals after the reform	0.018*	0.048**	0.035*
	(0.0079)	(0.011)	(0.016)
Change over time for the non-capped individuals	0.014**	0.012**	0.026**
	(0.0020)	(0.0018)	(0.0072)
Difference between capped and non-capped individuals at baseline	-0.0033	0.028**	-0.026*
	(0.0053)	(0.0070)	(0.012)
Constant (probability of depression among non-capped individuals before cap lowered)	-0.12**	-0.085**	-0.15**
	(0.0062)	(0.0066)	(0.024)
Controls for covariates	Y	Y	Y
Number of individuals	104056	68867	9615

Notes: Standard errors are reported in parentheses.

Supplementary analysis 8: Estimating the effects of benefit cap using information from the LFS and the FRS

In Supplementary analysis 2, we described the methods we used to derive the predicted probability of being capped for each individual in the LFS using information available to us in the FRS. Once we have the predicted probability of being subject to the benefit cap, we can then use different approaches to identifying the effect of the benefit cap.

Binary classifier approach
Our first approach is to convert the predicted probability into a 0/1 classifier to identify LFS cases at high risk of being affected by the benefit cap. To identify the optimal classifier we create a series of 100 binary classifiers by using probability cutoffs ranging from 0.01 to 1 in increments of 0.01. In the first iteration of these classifiers (cutoff 0.01), those benefit units with predicted probability greater than 0.01 would be classified as being at risk of being affected by the benefit cap. In the FRS data we compare this classification to the 'true' classification derived from the FRS benefit income calculations, thus determining whether each classification is a true positive or false positive. We choose an optimal cutoff probability by plotting a 'ROC curve' of true positive rate (tpr) vs false positive rate (fpr), and identifying the point at which $tpr = (1 - fpr)$. The true positive rate is the proportion of benefit cap cases that are correctly identified as such; the false positive rate is the proportion of non benefit cap cases that are incorrectly identified as benefit cap cases.

We use this optimal cutoff probability to assign each case in the LFS data to a binary classification. This allows us to, in a similar fashion to our simple LFS classifier, calculate a simple difference-in-difference estimate of the policy impact and plot the prevalence of reported mental health problems over time for the at risk group compared to the not-at-risk group. In addition we model the evolution of mental health problems over time using a linear regression model of the form

$$\text{Health}_{i,t} = \alpha + \beta_1 \text{Capped}_{i,t} + \beta_2 \text{time} + \beta_3 \text{Capped}_{i,t}\text{time} + \beta_4 \text{Policy}_t \text{Uncapped}_{i,t}\text{time} + \beta_5 \text{Policy}_t \text{Capped}_{i,t}\text{time}$$

In this model the coefficient of interest is β_5, the average additional quarterly prevalence increase after the reform was introduced, for capped individuals.

Figure 8.1 shows the ROC curve that plots the true positive rate tpr against the false positive rate fpr for 101 values of the cutoff probability for classifying a case as being affected by the benefit cap, varying from a cutoff probability of 0 to 1 in increments of 0.01. The optimal classifier where $tpr = (1 - fpr)$ can be read from the results table as a cutoff probability of 0.1, at which point tpr is 0.84 and fpr is 0.16. This is a binary classifier in

which all LFS cases with predicted probability of being affected by the benefit cap greater than 0.1 are classified as benefit cap cases, those equal to or below the cutoff are classified as non benefit cap cases.

Figure 8.1: ROC curve for binary classifiers

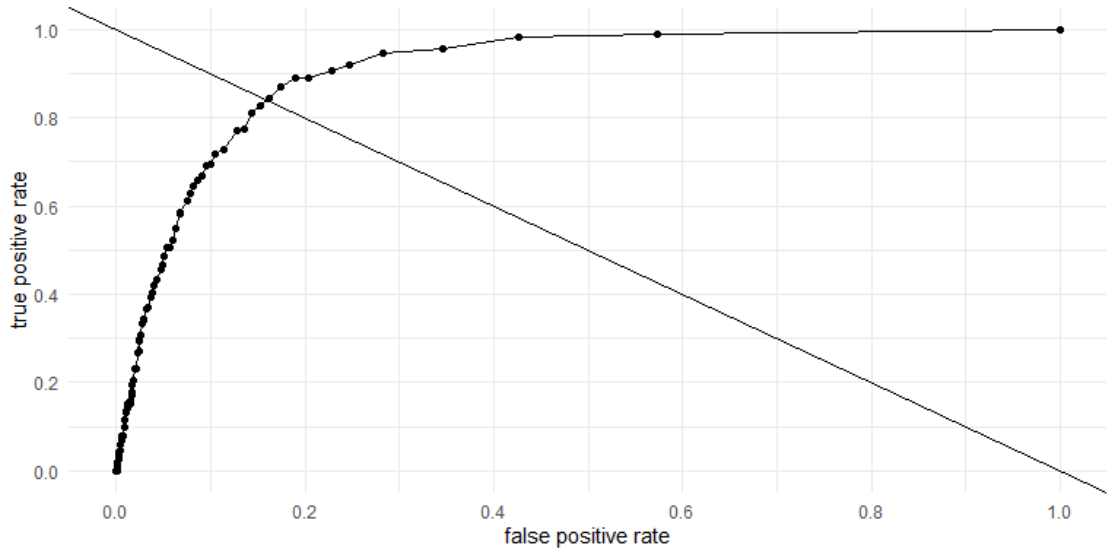

In an analogous fashion to the model in figure 4 (in the main paper) we estimate an unadjusted difference in difference model for the change in mental illness prevalence after the benefit cap was introduced for people at risk of being capped compared to those not at risk of being capped. The results shown in figure 4 (in the main paper) display a similar pattern to those for our simpler benefit cap risk indicator, although the effect size is larger than the estimate from our original model (β = 0.060, (95% CI: 0.045-0.075). If anything, then, our main estimates are likely to be conservative.

Figure 8.2 below replicates the analysis of figure 2 earlier in our paper. Again we find that the level of mental ill health before the reform was relatively stable amongst those at risk of being capped, with a steady rise following the lowering of the benefit cap.

Figure 8.2: Introduction of the benefit cap and the prevalence of mental ill health among those who are at-risk of being capped and those who are not, by quarter - FRS-based binary indicator

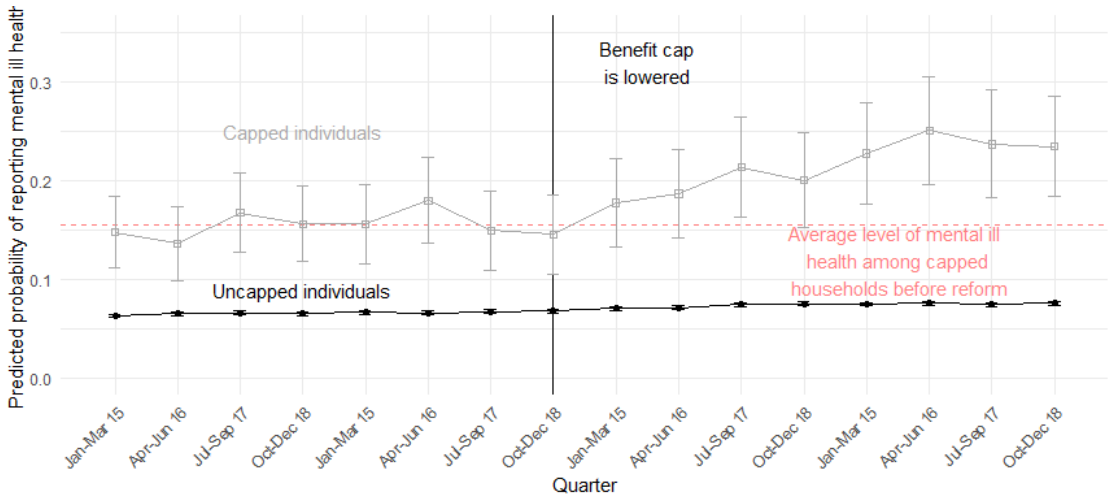

Table 8.3 shows the results of the linear regression model that fits the prevalence of self-reported mental ill health as a function of time, with separate slopes fitted for the capped and uncapped groups, before and after the reform. The only slope in the model that is significantly different from zero is the additional quarterly prevalence increase after the reform for capped individuals (last row in the table).

Table 8.3: Linear regression of mental ill health time series

	Estimate	Std. Error	t value	Pr(>|t|)
Prevalence among uncapped individuals in first quarter	0.06455	0.006268	10.3	1.146e-10
Difference between capped and uncapped individuals in first quarter	0.08201	0.008864	9.251	1.045e-09
Average quarterly prevalence increase for uncapped individuals, period before reform	0.0004924	0.001527	0.3225	0.7496
Average additional quarterly prevalence increase for capped individuals, period before reform	0.001524	0.002159	0.7058	0.4866
Average additional quarterly prevalence increase after reform, uncapped individuals	0.0003366	0.001178	0.2857	0.7774
Average additional quarterly prevalence increase after reform, capped individuals	0.004171	0.001178	3.54	0.001531

Examining the effect of the benefit cap using all the information from the predicted probabilities

A drawback of the 'binary classifer' approach is that it discards information about the probability of a case being affected by the benefit cap. For example, with a binary classifer cutoff probability of 0.1, a case with predicted probability 0.2 is treated the same as a case with predicted probability 0.5. Our second approach is to use the full information from the

predicted probabilities by using them as a measure of treatment intensity. We estimate a model of the form

$$\text{Health}_{i,t} = \alpha + \gamma_1 \text{AtRiskofCapping}_{i,t} + \gamma_2 \text{time} + \gamma_3 \text{CappingProbability}_{i,t} + \gamma_4 \text{CappingProbability}_{i,t} \text{time} + \gamma_5 \text{Policy}_t \text{CappingProbability}_{i,t} \text{time}$$

In this model the coefficient of interest is γ_5, the additional marginal quarterly increase in prevalence of mental ill health for a unit increase in the probability of being affected by the benefit cap, after the reform was introduced. The coefficient γ_4 is used to test whether there was any such relationship before the reform was introduced.

Table 8.4 shows the results from the linear regression of mental ill health prevalence on time and the probability of being affected by the benefit cap. The key result is in the final row of the table, which shows a positive and statistically significant relationship between the probability of being affected by the benefit cap and mental ill health prevalence, in the period after the benefit cap reform was introduced. By contrast, there is no statistically significant relationship between capping probability and increase in mental ill health in the period before the reform was introduced.

Table 8.4: Linear regression of increase of mental ill health prevalence on probability of being affected by the benefit cap

	Estimate	Std. Error	t value	Pr(>|t|)
Prevalence among uncapped individuals in first quarter	0.06688	0.0002436	274.6	0
Average quarterly increase in mental ill health prevalence for all individuals	0.0009734	5.203e-05	18.71	4.246e-78
Difference in first quarter between uncapped individuals and those with non-zero probability of being capped	0.2084	0.002098	99.32	0
Marginal increase in prevalence of mental ill health from increase of 0.1 in probability of being affected by the benefit cap	-0.03867	0.002322	-16.65	2.843e-62
Marginal quarterly increase in prevalence of mental ill health from increase of 0.1 in probability of being affected by the benefit cap - before reform was introduced	-0.0009691	0.0005508	-1.759	0.07849
Additional marginal quarterly increase in prevalence of mental ill health from increase of 0.1 in probability of being affected by the benefit cap - after reform was introduced	0.003756	0.000953	3.941	8.122e-05

Figure 8.5 displays the predicted evolution of mental ill health prevalence over time in the period before and after the benefit cap reform was introduced, according to the fitted model in table 8.4. It shows the evolution of mental ill health for individuals with zero probability of being capped against those with a probability of 0.28 of being capped. This is chosen for comparison with the binary indicator approach above, as the mean predicted probability of being capped for those individuals above the binary cutoff of 0.1. As can be seen, the model predicts an increase in mental ill-health prevalence for those at risk of being subject to the benefit cap, in a similar relationship to that shown in the binary indicator model of figure 8.2.

Figure 8.5: Introduction of the benefit cap and the prevalence of mental ill health among those who are at-risk of being capped and those who are not, by quarter - FRS-based binary indicator

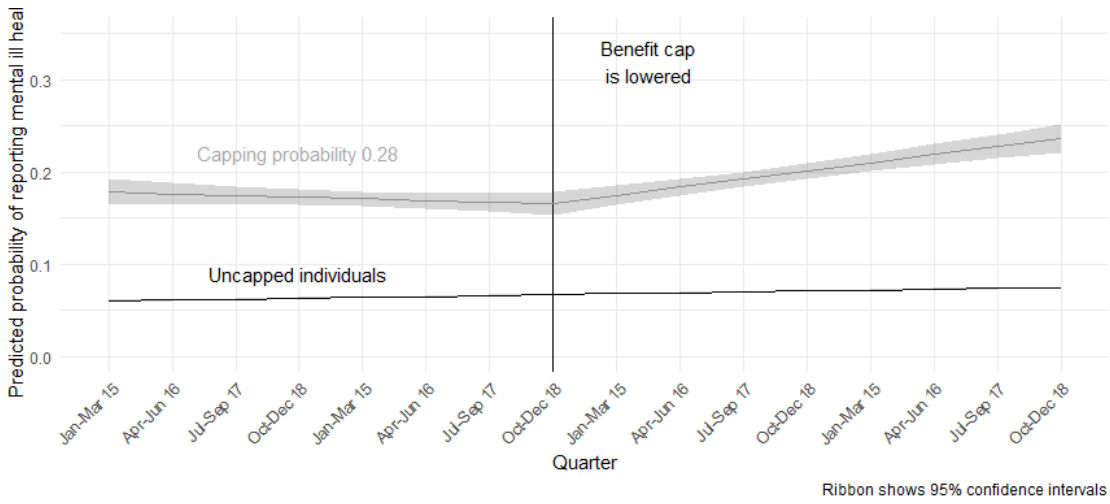

Exploring differences between those at risk of being capped in the FRS and those at risk of being capped in the LFS

These two different approaches to specifying who is and who is not at risk of being capped identify different groups of people. The LFS-based approach identifies 6824 households at risk and the FRS-based approach identifies around 4684. There is some overlap between these two approaches, that is, around 1498 people are found in both specifications. But this means that each measure is identifying some individuals that are not captured by the other. Here we explore these groups and then consider whether our estimates are different depending on the combination we pay most attention to.

First, we explore descriptive differences. As noted above, our LFS-based measure has some similarities to administrative data in terms of gender, age, and the presence of children in the household. But it was less accurate

in terms of the number of people based in London, the number of children in the household and whether households were lone parent. When we compare the LFS-based measure and the FRS-based measure, we find that the FRS-based measure corrects for some but not all of these biases. The FRS-based measure is more accurate when it comes to proportion of capped households based in London and the number of children in the household. But the FRS-based measure was less accurate in terms of the number of women, the number of single households, and the number of lone parents, compared to the administrative data. Neither is perfect, therefore, but both have strengths.

Next, we explore the difference-in-differences estimates using these two measures together. We use those who are categorized as not at risk of being capped in both measures as our baseline (or control group). We then compare this group to: 1) those only identified as being at risk in the LFS-based measure (the main one used in the paper), 2) those only identified as being at risk in the FRS-based measure (the sensitivity analysis), and 3) those identified as being at risk in both measures.

What these results show us is that the LFS-based sample potentially has a smaller difference-in-differences estimate than the FRS-based estimates, which are both slightly higher, albeit that the indicator with both FRS and LFS identified individuals has wide confidence intervals which overlap with the LFS only measure (p = 0.19). In short, whichever way we cut the data,

see a clear pattern: those at risk of being capped (defined in various ways) consistently see

Figure 8.5: Exploring the difference between FRS-based and the LFS-based measures

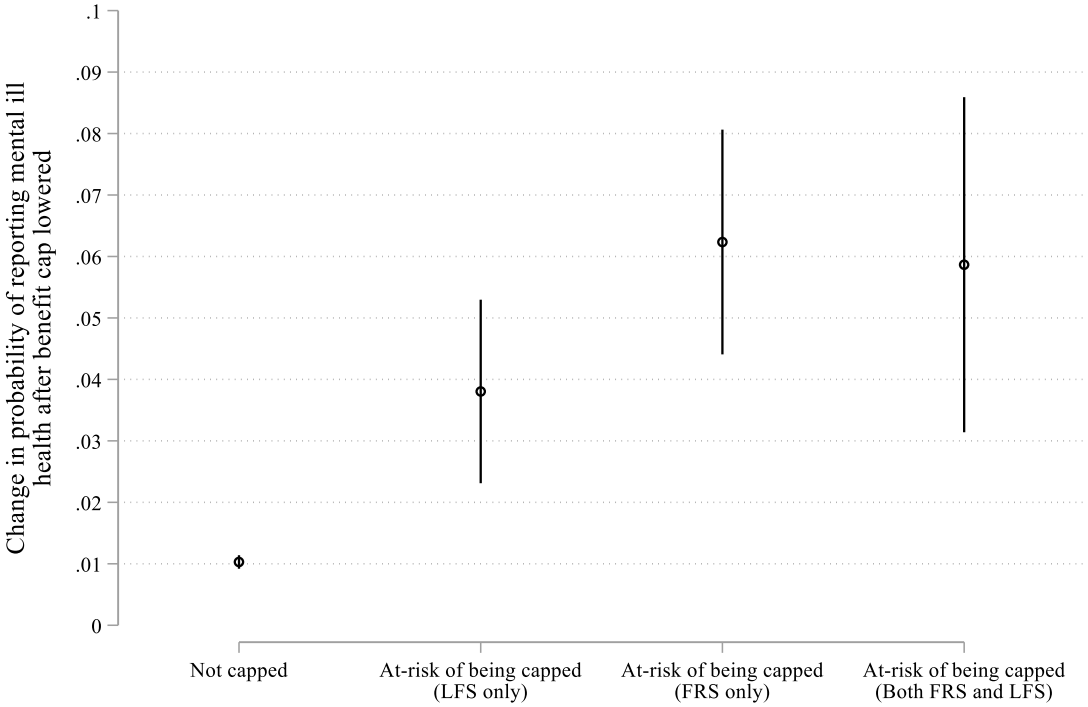

Viewed together, these results reinforce the primary message from our main results. We find a consistent and pronounced increase in the risk of mental ill health after the introduction of the benefit cap, among those who are risk of being affected by it.