Zero-hours Contracts in a Frictional Labor Market*

Juan J. Dolado†
Universidad Carlos III de Madrid
CEPR and IZA

Etienne Lalé‡
Université du Québec à Montréal
CIRANO and IZA

Hélène Turon§
University of Bristol
and IZA

Abstract

We propose a model to evaluate the U.K.'s zero-hours contract (ZHC) – a contract that exempts employers from the requirement to provide any minimum working hours, and allows workers to decline any workload. We find quantitatively mixed welfare effects of ZHCs. On one hand they unlock job creation among firms that face highly volatile business conditions and increase labor force participation of individuals who prefer flexible work schedules. On the other hand, the use of ZHCs by less volatile firms, where jobs are otherwise viable under regular contracts, reduces welfare and likely explains negative employee reactions to this contract.

JEL codes: E24, J22, J23, J63, L84.

Keywords: Zero-hours contracts; Working hours; Gig economy; Flexibility.

*We are grateful to Amanda Gosling, Andrea Ichino and Felix Wellschmied for their detailed comments, as well as to participants at the IZA workshop on *Heterogeneity and the Labor Market*, the workshop on *Alternative Work Arrangements* in Oxford, the 25th SMYE conference, the 2021 Simposio de Análisis Económico, and seminars at Birmingham, Cemfi, Durham, EUI, Kent, QMUL, UAB and UC3M for helpful suggestions. Dolado, Lalé and Turon acknowledge financial support from the Spanish Ministry of Science and Innovation (grant 2021/00366/001-Plan Estatal), Canada's Social Sciences and Humanities Research Council (Insight development grant 430-2019-00967) and European University Institute (Fernand Braudel Fellowship), respectively.

†Address: Departamento de Economia, Universidad Carlos III de Madrid, Calle Madrid 126, 28903 Getafe (Madrid), Spain – Email: dolado@eco.uc3m.es.

‡Address: Department of Economics, Université du Québec à Montréal, C.P. 8888, Succ. centre ville, Montréal (QC) H3C 3P8, Canada – Email: lale.etienne@uqam.ca.

§Address: Department of Economics, University of Bristol, Priory Road Complex, Priory Road, Bristol BS8 1TU, United Kingdom – Email: helene.turon-lacarrieu@bristol.ac.uk.

1 Introduction

Zero-hours contracts (ZHCs) have spread in the United Kingdom (U.K.) during the past decade, particularly in the low-pay segment of the labor market. These contracts, which exempt employers from the requirement to provide any minimum working hours while allowing employees to decline any workload, have become the focus of a heated controversy in the British media and political arena (see Adams et al. [2015] and Adams et al. [2018]). On the one hand, employers, and some employees, point to the benefits of having flexible labor contracts in the face of fluctuating demand conditions.[1] On the other hand, trade unions and other commentators have raised fierce concerns about the potential exploitation of workers, especially for those employed by online platforms that typically enjoy significant monopsonistic power (see Dube et al. [2020]). This controversy has ramifications for the design of labor laws and regulation of employment relations. Indeed, the legal status of workers on ZHCs lies between the categories of "employee" and "self-employed", making it unclear what benefits and welfare programs these workers should be entitled to. Similar debates exist in other OECD countries, where on-call contracts have been put into focus by the emergence of the gig economy.[2]

Despite this ongoing important debate, research on the equilibrium and welfare effects of on-call contracts is only at an early stage. Part of the reason for this is that detailed data on these labor contracts has only recently become available; see Abraham et al. [2021] for a discussion of the many challenges of measuring segments of the labor market that make more intensive use of these contracts. In the U.K. where the number of workers under ZHCs is coming close to the one million mark (Figure 1), these workers typically account for a small share of the sample of the country's labor force survey, which limits the scope of analyses that can be conducted solely using these data. One alternative to make up for these data shortcomings is to conduct specific questionnaire surveys (Katz and Krueger [2019a], Boeri et al. [2020]) or field experiments (Mas and Pallais [2017]) to gather evidence about the effects of these contracts. Another approach consists in combining a structural model with the empirical information provided by the labor force survey's data. This is the approach we undertake in this paper.

We seek to understand the pros and cons of ZHCs based on a frictional labor-market model where firms and workers are both heterogeneous in their valuation of ZHCs compared to regular contracts. Firms are characterized by more or less volatility in their revenue function, typically capturing the fluctuating demand conditions that affect low-pay service sectors where ZHCs

[1] An eloquent illustration was the offer made in 2016 by McDonald's to 115,000 of its U.K. employees to switch to regular contracts with a minimum number of guaranteed hours every week. The offer took place in the wake of complaints by staff in McDonald's restaurants that they were struggling to get access to loans and mortgages as a result of not having guaranteed employment each week. However, McDonald's reported that about 80% of these workers chose to remain on flexible ZHCs; see this article in The Guardian for details.

[2] If anything, the COVID-19 pandemic has reinforced this state of affairs. Many workers deemed "essential" during the pandemic were in jobs with on-call contracts, thus cumulating the higher risk of exposure to the virus to the disadvantages of precarious labor contracts. The case of Uber drivers is emblematic of these dynamics. The issue of whether Uber drivers are "workers" as opposed to "independent, third-party contractors" of Uber had been raised long before the pandemic. On 19 February 2021, the U.K. Supreme Court unanimously upheld a ruling that they are workers of the company, with rights to be paid at least the national minimum wage, to holiday pay and other benefits. See this article in The Washington Post and related discussion in this blogpost of the Economic Policy Institute.

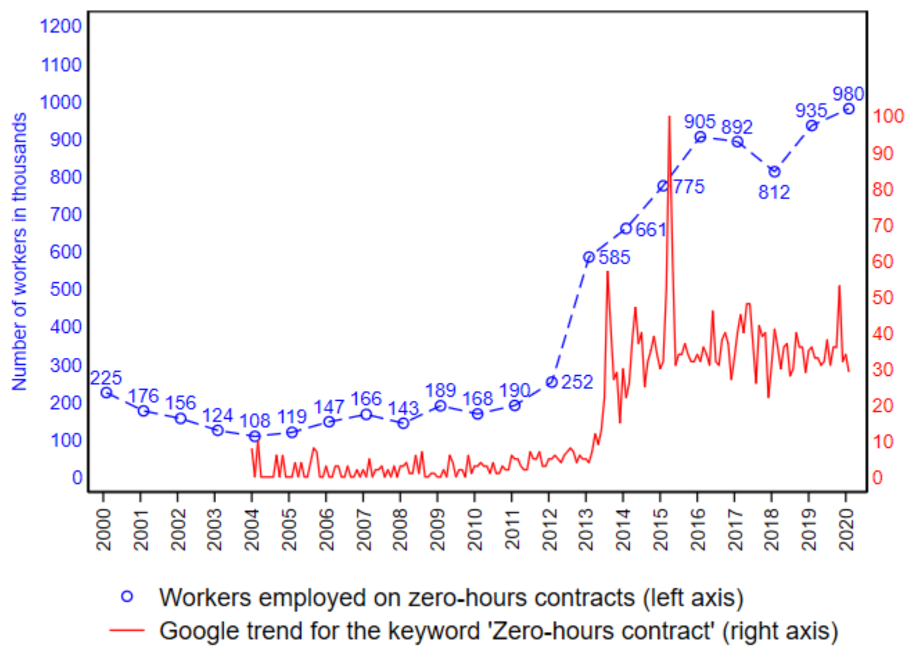

Figure 1: Workers on ZHCs and Google trend for the keyword "Zero-hours contract"

Notes: Left axis (dashed line): Office for National Statistics, calculations based on data from the Labour Force Survey weighted to official population projections. Data is annual from 2000 to 2013, semi-annual from 2014 to 2019 and quarterly in 2020. From 2014 to 2019, the number reported is the average of the two data points. For 2020, the number reported is the value for the 1st quarter. Data is not seasonally adjusted. Right axis (solid line): Google Trends, search for the keyword "Zero-hours contract" in the United Kingdom. Data is available monthly from January 2004 onwards.

are commonly used. Workers differ from each other with respect to their utility from working short hours, and since short hours are more prevalent in ZHCs, this creates a preference for or against these contracts. Thus, trade-offs between contract types arise for different agents, leading to sorting patterns of workers across firms that choose to offer either ZHCs or regular contracts. The model emphasizes three channels through which ZHCs affect the labor market. First, a job-creation effect, as firms endowed with more volatile business conditions can enter the market and/or are able to post more vacancies using these flexible contracts. Second, a substitution effect, whereby some jobs that would be otherwise viable under regular contracts (typically part-time jobs with fixed hours) become advertised as ZHCs. Third, a participation effect, as workers who prefer flexible work schedules join the labor market to take advantage of ZHCs. The relative importance of each of these three channels in shaping the equilibrium and welfare effects of ZHCs is ultimately a quantitative question. We address it by calibrating our model to data and policies from the U.K. Motivated by the U.K. debates on the consequences of the spread of ZHCs (Figure 1), we then use our model to analyze a ban on ZHCs.[3]

[3]The dashed line in Figure 1 shows that at the beginning of 2020 there are almost a million worker under ZHCs, which represents 3% of the U.K. labor force. The spread of ZHCs has received considerable policy and media attention, which is somewhat evidenced by the solid line in Figure 1 reporting the Google trend for the keyword "Zero-hours contract". Ever since the expansion of ZHCs in the 2010s, there have been protests calling for a ban on ZHCs in the U.K.; see, among many others, this article in Politics Home summarizing the views of Labour, the Lib Dems and the Conservatives on ZHCs.

Our main quantitative findings are as follows. A ban on ZHC leads to an increase of the unemployment rate in the low-pay labor market between 2.0 and 2.7 percentage points (p.p.). At the same time, it makes the employment rate of the low-pay labor market drop by between 4.8 and 5.4 p.p. The difference between the impact on unemployment and that on the employment rate comes from the labor force participation effect of ZHCs. We find that there are 3-4 percent of workers in this segment of the labor market who would prefer not to participate if there were no flexible contracts, such as ZHCs, that provide access to a shorter work schedule. The ranges of values for the unemployment and employment effects come from the so-called substitution effect of ZHCs. Suppose for instance that all ZHC jobs in the market are offered by firms that would otherwise find it profitable to hire workers under regular contracts. This minimizes the impact of a ZHC ban on the unemployment rate. Still, according to our model, employment falls by almost 5 p.p. The unemployment/employment impact is on the other hand maximized when all ZHCs jobs are provided by firms that would not be able to create jobs if flexible contracts were made illegal. In sum, even when substitution effects are largest, we find that a ban on ZHCs decreases job creation substantially. This is because those firms that find it profitable to substitute ZHCs for regular contracts conditional on filling a vacancy also find it more profitable to post more vacancies when ZHCs come into operation, as well as because all firms create more jobs when labor force participation is higher.

Next, to get the full picture of the equilibrium consequences of a ban on ZHCs, we study its impact on accession rates to regular employment. In a frictional labor market setting, there are two factors that counteract the decline in the aggregate job-finding rate driven by lower job creation. First, workers who were previously employed in ZHCs become more effective at searching for regular jobs – indeed, we estimate that their on-the-job search efficiency in ZHCs is lower than when unemployed. Second, vacancies that advertise regular contracts are more likely to reach these workers once ZHC vacancies are no longer "diluting" their chances of doing so. As a result, although unemployment increases in response to this policy reform, we find that regular employment expands and that the average duration out of regular employment decreases. The broader lesson from these experiments is that jobs that can serve as a stepping stone towards regular employment, such as ZHCs, may also imply that labor market trajectories are more unstable on average.

Given these results, the welfare consequences of a ban on ZHCs are not obvious. In our model, these consequences are well-defined for workers who participate in the labor market even in the absence of ZHCs. We find that the gain from avoiding transitions through ZHCs is offset by the increase in unemployment, meaning that the reform generates welfare losses among these workers. Depending on their valuation of short hours, and on the strength of the substitution effects, the welfare impact ranges between -0.9 and -1.1 percent of foregone consumption. While these effects suggest that ZHCs are wholly beneficial for the low-pay labor market, our model also rationalizes concerns on how hours flexibility is shared between firms and workers – specifically, that workers bear the costs of unpredictable hours schedules under ZHCs. In the (partial equilibrium) experiment where we transplant workers who always participate in the labor market into an environment where their ZHCs have been replaced by regular contracts, we find welfare gains that range between 0.2 and 0.5 percent in consumption equivalent variation. In other words, although the substitution effect is offset by the other labor market impacts of ZHCs, the effect is not negligible in absolute terms. This may explain the seemingly paradoxical responses to the spread of ZHCs described in the opening paragraph. In a general equilibrium setting, other forces (participation and job creation) come into play and, as it turns out, overturn these negative consequences of ZHCs.

Finally, a by-product of our analysis is to derive model-based predictions about workers' willingness to pay (MWP) for shorter work shifts. Indeed, in our model we uncover regions of the parameter space where the marginal utility of working short hours leads to a preference for ZHCs over regular employment contracts over unemployment, regions where the marginal utility leads to a preference for ZHCs over unemployment over regular employment, and so on. Consider those workers who would always rank employment over non-employment, and conditional on working they would prefer a regular employment contract over a ZHC – we call them "more attached workers". These workers do not value short hours very much: they would be willing to give up at most £4.7 of consumption (per week) to avoid working one hour beyond their work availability. On the other hand, the "least attached workers", who participate in the labor force only to work on ZHCs but would not accept to work on regular contracts, have a much higher valuation of short hours. That is, the lower utility bound indicates that they would give up at least £10.9 of consumption (per week) to avoid working one hour beyond their work availability. This MWP is 45 percent higher than the hourly minimum wage (£7.5 per hour) that workers can earn in this segment of the labor market.

The paper contributes to a growing body of research on understanding alternative work arrangements, much of which is motivated by the advent of the online gig economy.[4] In addition to monitoring the trends in the development of alternative work arrangements (e.g., Abraham et al. [2019, 2021], Katz and Krueger [2019b], Collins et al. [2019], Boeri et al. [2020]), the focus of this research is on understanding the reasons why firms use non-standard employment, and how this affects a broad set of outcomes of workers; e.g., Bloom et al. [2015], Mas and Pallais [2017], Koustas [2018], Katz and Krueger [2019a], Angelici and Profeta [2020], among many others. Our contribution is to address these questions through the lens of a structural model. This approach allows us to model heterogeneous agents on both sides of the market and simulate their behavior in response to the changing flexibility of labor contracts as well as the impact on their welfare. Our analysis speaks to this research by producing estimates of the valuation of flexibility that can be compared to the MWPs that have been estimated empirically. In addition, the structure afforded by the model enables us to investigate the relation between ZHCs and labor market policies, thereby complementing the empirical analysis of U.K. policies presented in Datta et al. [2019]. The U.K. is an interesting case because of the growth of ZHCs that took place recently, but similar contracts exist in Australia, Canada, Finland and Ireland, suggesting that these policy analyses can be of broader relevance.[5]

[4] However, alternatives to regular, fixed hours employment contracts are clearly not a new feature of labor markets. Other flexible work arrangements, such as those labeled "reservist", "on call", and "if and when" contracts (see Dickens [1997]) date back to the 19th century where workers hired under piece-rate contracts were not guaranteed any amount of fixed work on a daily or weekly basis, e.g. in industries involving dock labor. In this respect, the lessons of our analysis extend beyond the recent experience of the U.K. labor market.

[5] Differences between the U.K.'s ZHC and the ZHCs in place in Australia, Canada, Finland and Ireland

To our knowledge few attempts have been made to analyze flexible hours contracts in a structural model. The one most closely related to ours is Scarfe [2019]. The author develops a frictional labor-market model to analyze the coexistence of "casual work" (akin to ZHCs in our analysis) and regular employment, and calibrates the model on data for Australia where casual workers account for about 10% of the labor force. Her model and ours differ along several important dimensions. Foremost, we emphasize *ex ante* heterogeneity as a key source of variation to understand firms and workers' ranking of different labor contracts. In Scarfe [2019], agents are homogeneous *ex ante* and the choice of contracts depends much on "luck", namely the stochastic draw of match productivity at the time of meeting between firms and workers. Our assumptions are guided by our focus on linking the pros and cons of ZHCs to individuals' heterogeneous valuations of flexibility. Scarfe [2019]'s approach is rather guided by the objective of making her model tractable to theoretically analyze casual jobs. For this reason, we view the two studies as being complementary to each other.[6] The other related paper in this vein of the literature is Frazier [2018]. The author considers a directed search model of hours and wages to analyze the effect of regulations restricting variation in working hours. The main finding is that search frictions generate imperfect sorting between workers and firms where the key trade-off is between the wage level and hours flexibility. Sorting plays an important role in our analysis as well, since it introduces *ex ante* rankings between labor contracts which is absent from Frazier [2018]'s analysis. This ranking, we argue, is key to evaluate the pros and cons of ZHCs.

The rest of the paper is organized as follows. To set the stage for our analysis, in Sections 2 and 3 we describe the regulatory framework of ZHCs in the U.K. and present a set of stylized facts about these contracts. In Section 4, we present our model of the equilibrium and welfare effects of ZHCs. In Section 5, we explain the model's calibration, and in particular how it is used to draw inferences about firms' and workers' heterogeneous valuation of ZHCs. We analyze the labor market effects of a ban on ZHCs in Section 6. Finally, Section 7 concludes. An Appendix gathers further details about the model and robustness checks of the results.

2 Regulatory framework

This section reviews the legal status of individuals on ZHCs in the U.K., as well as their entitlement to welfare under these contracts.

relate to the legal status of zero-hours workers (see Section 2) and the levels of regulation of these contracts. Similar on-call contracts also exist in Scandinavian countries and in Cyprus and Malta. Likewise, they are also used, albeit subject to much heavier regulations, in Germany, Italy and in the Netherlands. They are either explicitly forbidden or not used in the remaining countries of the European Union. In the United States, on call working arrangements are also growing in importance. Despite the absence of federal regulation, several states operate 'show-up pay' laws, where employers are required to pay workers for a minimum number of hours, if they have been called to work, though coverage varies across these states and a number of exemptions exist.

[6] In terms of quantitative results, Scarfe [2019] finds that a ban on casual work (akin to a ban on ZHCs in our model) leads to an almost doubling of the unemployment rate (i.e., a rise from 7% to 13%). This large unemployment effect is related to her assumption of *ex ante* homogeneity among firms: since *any* entering firm is susceptible of becoming an employer with casual workers, the ban has a large negative effect on firms' entry. On the other hand, in our model we find that most firms' heterogeneous types is such that they are indifferent to the existence of ZHCs, in the sense that they would always find it more profitable to post regular contracts.

Workers' rights. As will be explained below, ZHCs typically give staff a "worker" employment status, which lies between the traditional categories of "employee" and "self-employed". This intermediate status confers such individuals with the following employment rights:[7]

- Right not to be discriminated against under the Equality Act 2010; right to receive pro-rata holiday pay and other working time rights (Working Time Regulations 1998); and right to receive Statutory Sick Pay (as long as they have met the Lower Earnings Limit);
- Automatic enrollment for pensions;
- Protection from unlawful deductions from wages;
- Right to receive the hourly National Minimum Wage or National Living Wage.[8]

However, in contrast to employees, ZHC workers are not entitled to redundancy pay when dismissed but have more rights than the self-employed who are only entitled for protection for health and safety on a client´s premises and against discrimination. It is noteworthy that exclusivity clauses in ZHCs, which stop a zero-hour worker from taking on another job, were banned in May 2015. Employers cannot enforce the clause, and since January 2016, workers have been able to claim compensation at an employment court if they are punished or dismissed for looking for work elsewhere.

Whereas in the U.K. workers under ZHCs are not obliged to provide any minimum working hours, in this respect it is noteworthy that in Ireland individuals are contractually obliged to be available for work if called by employers. By contrast, "self-employed" individuals have no employment rights besides certain discrimination rights. At the other end of the spectrum "employees" have the whole range of employment rights, such as paid maternity leave, including unfair dismissal and redundancy and family rights.

The distinction between the status of "worker" and that of "employee" has been subject to court litigation recently. A well-known case is whether companies like Uber or Deliveroo should hire under employment contracts or freelance work. To the extent that some of these firms use contractors rather than employees, they do not fall into the above definition of ZHCs. The most important difference between these two categories of workers is that employers must offer "employees" work in exchange for pay, and "employees" are required to do the work, whereas "workers" can turn work down, depending on their time *availability*.

Moreover, whether an individual is considered to be an "employee" or a "worker" depends not just on what it is the offered contract, but also on what happens day to day. While a contract might stipulate that there is no obligation to work, if the individual is "punished" for

[7] There is, however, some controversy between trade unions and employers' associations about whether individuals on ZHCs are "workers" or "employees". While the Trade Unions Congress considers that most of them are "workers", the Chartered Institute of Personnel and Development [CPID, 2013] reports that two-thirds of employers (interviewed in a survey carried out by the institute) classify zero-hours individuals on ZHCs as "employees".

[8] In the U.K. there have been several minimum wages in place during the recent period. From April 2016 there are three minimum wage rates for young workers (those aged 16-17, 18-20, and apprentices), another minimum wage rate for young adults (those aged 20-24), and finally the new National Living Wage (NLW) for individuals aged 25 and above. The NLW was raised in April 2019 to £8.21 an hour.

not accepting all the offered hours offered, or consistently work a set number of hours, then a tribunal might decide that the worker is actually an “employee”.

Entitlement to welfare. Since workers under ZHCs are often low-wage earners, they are entitled to means-tested benefits and tax credits. In the past, the benefits one could claim hinged on whether the individual worked more than 16 hours in a week, as in the case of the Income Support program or Jobseekers’ Allowance (JSA). When working 16 hours a week or more, individuals could also claim the Working Tax Credit, Child Benefit and Housing Benefit if they needed help with the rent and had savings less than £16,000. However, the Universal Credit (UC) in 2013 replaced all of these income support schemes with a taper rate of 65 percent implemented from a typical monthly work allowance of (net of taxes) £490 for single workers. The UC taper rate was reduced to 63 percent in 2018, and a further reduction to 55 percent is expected to come into force at the end of 2021.

3 Stylized facts about ZHCs

Next, we present some stylized facts about ZHCs. These facts are useful not only for context purposes but also to guide the development of the model presented in Section 4.

We use data from the U.K.’s household labor force survey (LFS). This survey covers a large number of individuals, which is important for our purposes since, despite their growth, ZHCs still represent a fairly small portion of the labor market. The U.K.’s LFS has a modest longitudinal dimension as it follows individuals over five quarters, with one fifth of the sample being renewed every quarter. Among these five interviews, respondents are asked twice (one semester apart) whether they hold a ZHC. We pool eight quarterly waves of the longitudinal LFS, up to the onset of the COVID-19 pandemic. Specifically, the survey responses that we analyze cover the period from September 2018 to March 2020. One motivation for focusing on this time period is that Figure 1 suggests that the U.K. labor market was stable during those years along the dimensions we analyze.

As our focus is on the segment of the labor market where jobs pay around the minimum wage, we restrict our data analysis to individuals who either work in a low-paying occupation in at least one quarter or are ‘not employed’ in any of five quarters while not being recorded as ‘inactive’ throughout the five quarters.[9] We define low-paying occupations as those classified as either ‘Administrative and secretarial’, ‘Caring, leisure and other service’, ‘Sales and customer service’, ‘Process, plant and machine’, or ‘Elementary’ occupations. We exclude all retirees and those under age 16 in the first quarterly interview. This leaves us with 9,342 individuals aged 16 to 69 belonging to the low-pay segment of the labor market.

The unemployment rate in the first interview is 11.2 percent. Every other quarter, the LFS includes a question asking respondents whether they hold a ZHC. In particular, 356 individuals,

[9] This sample restriction does not exclude from our sample some long-term unemployed who are looking for a high-paying occupation job. Consequently our approach might overstate the unemployment rate of the low-pay segment of the U.K. labor market.

i.e. 3.8 percent of our sample, reports being employed under a ZHC in the first interview where this question is included. This represents 4.4 percent of all employees in the low-pay segment of the market. Likewise, 431 individuals, or 4.6 percent of our sample (5.2 percent of all employees), report having a ZHC in either the first or the second interviews. This is about 1.5 times the share of ZHC workers in the overall labor market. As expected, all individuals on ZHCs report being employed.

Table 1: Characteristics of people by labor contracts

(a) **Gender**	N	Z	R
Men	44.0	43.5	39.6
Women	56.0	56.5	60.4
(b) **Age**	N	Z	R
16 to 19 years	21.2	16.6	3.0
20 to 24 years	9.3	11.5	4.9
25 to 29 years	5.8	7.3	6.2
30 to 34 years	6.8	6.2	7.7
35 to 39 years	6.3	3.9	8.7
40 to 44 years	5.7	5.6	9.8
45 to 49 years	7.4	9.6	11.9
50 to 54 years	8.7	8.7	14.9
55 to 59 years	10.6	12.4	15.9
60 to 64 years	12.5	11.0	12.6
65 to 69 years	5.7	7.3	4.6
(c) **Education**	N	Z	R
Degree or equivalent	16.3	21.9	18.0
Higher education	6.7	8.7	8.9
GCE A level or equivalent	20.4	21.9	23.8
GCSE grades A*-C or equivalent	31.6	25.8	27.8
Other qualification	9.2	10.7	9.8
No qualification	13.9	8.2	9.2
No answer or don't know	1.8	2.8	2.7
(d) **Industry**		Z	R
Agriculture, forestry and fishing		1.1	0.9
Mining and quarrying		0.0	0.3
Manufacturing		3.9	7.9
Electricity, gas, AC supply		0.0	0.5
Water supply, sewerage, waste		0.0	1.0
Construction		0.8	3.7
Wholesale, retail, repair of vehicle		8.7	17.7
Transport and storage		8.7	7.9
Accommodation and food services		19.9	4.0
Information and communication		0.8	1.2

Financial and insurance activities	0.6	2.9
Real estate activities	0.3	0.7
Prof., scientific, technical activities	2.5	5.1
Admin. and support services	6.7	5.5
Public admin. and defense	2.0	7.2
Education	8.4	10.8
Health and social work	20.5	15.5
Arts, entertainment and recreation	6.7	1.9
Other service activities	3.4	3.2
Households as employers	0.8	0.3
Extraterritorial organizations	0.0	0.1
No answer	3.9	1.5

Note: Authors' calculations based on data from the Labour Force Survey. N: Unemployed, Z: Employed in a zero-hours contract, R: Employed not in a zero-hours contract. All table entries are expressed in percent.

In Table 1 we start out our description of ZHCs by comparing the cross-sectional characteristics of workers in these contracts (denoted in short as Z) with workers in regular contracts (R) and workers without employment (N).[10] First, we note little difference in terms of gender: female employees account for 56.5 percent of ZHC employment vs. 60.4 percent of regular employment, and this difference is not statistically significant. Differences with respect to age, on the other hand, are large. The (not reported) mean age of workers under ZHCs is 40.8 vs. 46.3 years old for employees in regular contracts. In order to get a more precise view of the incidence of ZHCs over the life cycle, Panel (b) of Table 1 displays the shares of each 5-year age bands in Z and R employment and in unemployment N. This shows an increased prevalence of Z contracts at both ends of the working life. Indeed, relative to both R and N individuals, the group of Z workers displays a much lower share of workers aged 30 to 44 years old. Relative to R workers, the age distribution of workers in Z contracts is largely skewed towards younger workers; relative to N individuals, the share of workers aged 45 and above is slightly higher. The latter fact could be related to the demand for flexibility among older workers to smooth out the transition to retirement.[11]

Panel (c) of Table 1 displays the distribution of educational attainment among employed and unemployed workers. As in Datta et al. [2019], these distributions only exhibit modest differences. 21.9 percent of ZHC workers hold a degree or equivalent vs. 18 percent of employees in regular contracts, and 19 percent of ZHC workers hold no or "other" qualifications, which is the same fraction as for those in regular employment. Hence, two facts stand out from these findings: (i) a substantial fraction of employment in "low occupations" are highly qualified,

[10] 'Regular contracts' is somewhat of a misnomer because this category lumps together all employed workers not on ZHCs, which could include workers on part-time employment contracts, temporary jobs, etc. Thus, differences between Z and R workers might be dampened by the fact that workers employed in those 'regular' but precarious employment contracts are likely to resemble workers on ZHCs.

[11] Ameriks et al. [2020] document, based on survey data for the United States, that there is a large potential for increasing labor force participation among older workers through the use of jobs with shorter work schedules.

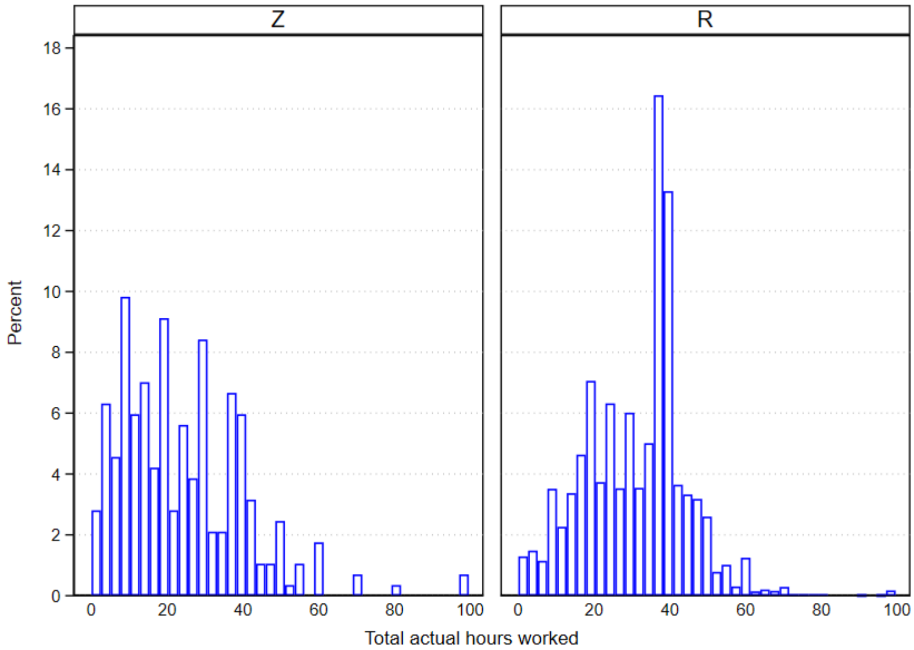

Figure 2: Distribution of actual hours worked

Notes: Authors' calculations based on data from the Labour Force Survey. Total actual hours worked exclude holidays. Z: employed in a zero-hour contract, R: employed not in a zero-hour contract.

corresponding perhaps to young college students and elderly college graduates who use ZHCs to complement their earnings;[12] and (ii) the different contract types in this segment of the labor market are filled with very similar workers in terms of education. Next, Panel (d) of Table 1 shows the breakdown of employment in either contract type by industry. We note that industries that are over-represented in ZHC employment are 'Accommodation and food services', 'Health and social work' and 'Arts, entertainment and recreation'. Conversely, under-represented industries in ZHC employment are 'Manufacturing', 'Construction', 'Wholesale, retail, repair of vehicle' and 'Public administration and defense'. The relation between ZHCs and 'Accommodation and food services' and 'Health and social work' is in line with what one would expect and has been highlighted by the COVID-19 pandemic.

In our dataset, the only measure of hours that does not suffer from a large fraction of missing data is "total actual hours in the main job". Since there is no information about whether the respondent is on holiday during the LFS's reference week (the time frame used to measure actual hours worked), we assume that individuals reporting hours in the lowest decile of the distribution of hours worked are on holiday. Figure 2 reports the cross-sectional distributions of hours among individuals at work with either type of contract. As can be inspected, ZHC

[12] In general, the relationship between education and the probability of being employed in a flexible job is a complex one. On the one hand, workers with more educational attainment who benefit from more bargaining power are more likely to obtain more schedule and location flexibility, and at the same time employers face a lower cost of providing flexibility in high-skilled jobs. On the other hand, less educated workers are more likely to be employed in jobs which require working overtime, having night shifts, or receiving shorter advance notice about their schedules. Mas and Pallais [2020] find a statistically significant, positive relationship between higher education and the probability of employment in some form of flexible work arrangement; see Section 1.2 of their paper and the insightful discussion therein.

workers spend on average fewer hours on the job, and the cross-sectional standard deviation of these hours is higher than in regular contracts. To inspect further the source of these differences, we use the longitudinal dimension of the LFS to compute the mean and standard deviation of hours worked for each individual over the five quarterly interviews in which they are included in the survey. Table 2 displays the average of the individual-level means and individual-level standard deviations obtained for those who are continuously employed under ZHCs and under regular contracts, respectively.[13] As expected, these figures suggest that ZHC jobs offer fewer hours of work and a more volatile hours schedule than regular jobs. The LFS also includes questions about an individual's willingness to work more hours. 16.6 percent of individuals in ZHCs indicate that they would like to work more hours vs. 10.1 percent in regular contracts. In a similar vein, 18.2 percent of ZHC workers report looking for another (or additional) job whereas only 5.0 percent of employees in regular jobs do so. This suggests that a higher fraction of workers in ZHCs are in underemployment, but that a majority of them are satisfied with both their job or their hours, to the extent that they are not looking to change either.

Table 2: Mean and standard deviation of actual hours worked

Continuous employment in:	***Z***	***R***
Mean	19.4	28.1
Standard deviation	7.8	7.2

Notes: Authors' calculations based on data from the Labour Force Survey. Total actual hours worked exclude holidays. Z: Employed in a zero-hours contract, R: Employed not in a zero-hours contract.

We now turn to different measures of labor market mobility across contract types. These measures will play an important role in informing the model's calibration in Section 5. Table 3 shows the distribution of job tenure with the current employer among workers in either type of contract. Rather surprisingly, nearly half of ZHC workers report job tenures longer than 2 years, which contrasts with the popular image of ZHCs as precarious employment contracts.[14] At any rate, job tenures are on average shorter in ZHCs than in regular contracts, but probably less so than expected: 9.2 percent of ZHC workers were recruited in the last 3 months vs. 3.4 percent of employees in regular contracts, and 30.3 percent of ZHC workers have been with their current employer for less than a year, vs. 14.3 percent of employees in regular contracts. For the purposes of our model, we also look at the duration of spells of unemployment. This distribution, reported in Table C1 of the Appendix, indicates that over half of the unemployed have been without jobs for less than 6 months. At the same time, a substantial 21 percent have exceeded two years without jobs.

Finally, Table 4 displays the transition matrix between the three labor market states under consideration, namely, unemployment N, employment under a zero-hour contract Z and employment under a regular contract R.[15] Several interesting observations emerge from this

[13]As before, we assume that the observations in the first decile of the hours distribution correspond to holidays and excluded them from these calculations.

[14]12.3 percent of ZHC workers report job tenures longer than 10 years. Besides reporting error, another factor that could explain this figure is that their contracts might have been reclassified as ZHCs.

[15]Note that these transitions occur over one semester since we only observe the response to the question:

Table 3: Distribution of job tenure by labor contracts

Length of time with current employer:	***Z***	***R***
Less than 3 months	10.7	4.2
3 to 6 months	8.3	4.0
6 to 12 months	11.3	6.1
1 to 2 years	19.6	9.7
2 to 5 years	26.0	19.3
5 to 10 years	12.2	15.5
10 to 20 years	9.8	26.0
More than 20 years	2.5	15.2

Notes: Authors' calculations based on data from the Labour Force Survey. Z: Employed in a zero-hours contract, R: Employed not in a zero-hours contract. All table entries are expressed in percent.

matrix. First, 11 percent of exits from unemployment are to ZHCs. Second, the transition rate to unemployment is almost 50 percent larger in ZHCs than in regular contracts (6.2 vs. 4.4 percent). Third, the job-to-job transitions involving a change of contract type are predominantly from ZHCs to regular contract. These are key empirical findings to inform the design and calibration of our model. In particular, in Section 5 we will exploit the information that is contained in the distribution of job tenure and unemployment duration beyond that conveyed by the transition matrix shown in Table 4.

Table 4: Transition rates between employment status and labor contracts

	To:	***N***	***Z***	***R***
From:	***N***	62.2	4.2	33.5
	Z	6.2	87.3	6.5
	R	4.4	0.5	95.2

Notes: Author's calculations based on data from the Labour Force Survey. N: Unemployed, Z: Employed in a zero-hours contract, R: Employed not in a zero-hours contract. All table entries are expressed in percent.

4 The model

We propose a model that analyzes the incentives of heterogeneous firms and heterogeneous workers to post/accept zero-hours (Z) and regular (R) contracts. In addition, the model provides an understanding of the equilibrium stocks and flows between employment in either contract type and non-employment.

'Do you hold a zero-hour contract?' every other quarter.

4.1 Economic environment

Time is discrete and runs forever. Agents discount the future at a common rate ρ. Given our focus on the low-pay segment of the labor market, it is assumed that all jobs pay the national minimum hourly wage, denoted as w. We think of this segment of the labor market as being subjected to rapidly changing business conditions. Typically, we consider a time period to be two weeks and assume that business conditions in a given fortnight are independent of those in the preceding fortnight. These shifting business conditions motivate the use of flexible labor contracts, such as ZHCs. In line with the regulatory framework of ZHCs, we consider a "regime" where the firm chooses the number of hours worked and allows the worker to choose the timing of work. Later on in the analysis we consider another regime where workers in ZHCs can decline any workload beyond a threshold of hours worked.

Workers' and firms' types. Workers and firms are heterogeneous. Workers' type is denoted as i, and in an equilibrium $i = 1, \ldots, 4$ corresponds to the possible ranking of the opportunities that are available to workers. Letting N denote a worker's asset value of being not employed, and W_Z and W_R denote workers' asset values of being employed under a Z contract and R contract, respectively, workers' type i are labelled as follows:

$$\begin{cases} N < W_Z < W_R & \text{for type 1} \\ W_R < N < W_Z & \text{for type 2} \\ W_Z < N < W_R & \text{for type 3} \\ N < W_R < W_Z & \text{for type 4} \end{cases} \tag{1}$$

(ignoring types that would rank non-employment higher than any form of employment). Hence, type-1 and type-4 workers are the ones most attached to the labor market as non-employment is their least preferred option. By contrast, types 2 and 3 are less attached in the sense that they would prefer non-employment over employment in a certain labor contract. In particular, if Z contracts were to be banned, type-2 workers would prefer to remain inactive in equilibrium (that is to say if the asset values in the equilibrium without Z contracts are such that $W_R < N$ for type-2 workers). We do not model the labor force participation rate in the equilibrium with Z contracts. Yet, through type-2 workers, we can assess by how much this rate would *change* in response to a ban on ZHCs, and thereby allow the model to capture the so-called *participation* effect of these contracts.

Firms can be of one of three types $j = c, s, r$, depending on their ranking of the asset values V_Z and V_R of advertising a vacant position as either a Z contract or R contract. In equilibrium, firms' type j are labeled as follows:

$$\begin{cases} V_R < 0 < V_Z & \text{for type } c \\ 0 < V_R < V_Z & \text{for type } s \\ 0 < V_Z < V_R & \text{for type } r. \end{cases} \tag{2}$$

That is, when Z contracts come into operation, type-c firms would be the ones that predominantly *create* these jobs (hence label c) while, when abolished, they are the ones which would abstain from creating any jobs. Thus, type-c firms enable us to capture the *job creation* effect associated to Z contracts. Type-s firms, on the other hand, advertise jobs as Z contracts even though they would profitably hire workers under R contracts, and they would *switch* to R contracts if Z contracts were banned. Hence, this type of firms allows us to capture the so-called *substitution* effect. Finally, firms of type r would *remain* offering R contracts irrespective of whether Z contracts are legal or banned. Notice that although a type-r firm always posts R contracts, the expected discounted value of its profits will in general depend on the presence of Z contracts in the labor market. More generally, whether Z contracts are available or not will affect equilibrium conditions, which in turn matters for the calculations of the asset values and implies that all firm types may adjust their job creation decisions in response to a ban on Z contracts.

We now turn to the economic environment that gives rise to the coexistence of these heterogeneous worker and firm types.

Workers' preferences. Workers derive utility from consumption and leisure. In non-employment, they receive unemployment benefits b, while in employment they are paid at the minimum wage w and work for h hours. When earning labor income wh, workers may retain part of their unemployment compensation depending on the taper rate τ of welfare benefits, so that earned labor income is:

$$\text{inc}\,(h) = \max\left\{wh, b + (1-\tau)\,wh\right\} \tag{3}$$

Workers are available to work for a hours at no utility cost. Hours worked beyond a generates a disutility, which is assumed to be linear with respect to the extra hours of work and measured by: $\alpha_i \max\{h-a, 0\}$. Observed that hours worked h are not subject to any indivisibility, i.e. if employed in a Z contract the worker is able to choose the timing of her work shifts, so that the disutility depends "only" on the $h-a$ excess hours. α_i, measuring the marginal utility cost of working excess hours, is specific to each worker type i.[16] In an equilibrium, this parameter drives workers' heterogeneous valuations of Z contracts. We assume that earned labor income in (3) is the only source of income for workers and we rule out saving/borrowing, so that workers consume all their labor income. Utility from consumption is derived according to a CRRA function.[17] The intra-period utility function is thus given by:

$$u^i\,(h, a) = \frac{\text{inc}\,(h)^{1-\eta} - 1}{1-\eta} - \alpha_i \max\left\{h-a, 0\right\}. \tag{4}$$

[16]Throughout the text, we describe workers as having a preference for short hours. In Equation (4), we model this by introducing a disutility from working long hours. This formulation gives rise to a marginal willingness to pay to avoid working long hours (which we analyze in Section 5.4). We find the latter more intuitive to interpret because it can be compared to the rate of overtime pay.

[17]We use a CRRA function for the calculation of welfare effects in Section 6.

There is a unit continuum of workers. The shares of each worker type i are exogenously given and are denoted as ω_i.

Production technology. Firms face shocks to their revenues which reflect fluctuations in demand conditions and/or production shocks. We entertain the idea that profits depend on actual hours worked, h, as well as on hours $\widetilde{h}$ measuring the number of working hours that would exactly meet the demand that a firm faces at a given point in time. Deviations between h and $\widetilde{h}$ are costly due to, e.g., reputation costs (if the firm is not able to produce enough to satisfy the demands of its consumer), marketing expenses (if the firm needs to expand marketing to sell extra units of output), etc. The firms' instantaneous profit function is:[18]

$$\pi\left(h, \widetilde{h}\right) = (p - w)\, h - \frac{\phi}{2}\left(h - \widetilde{h}\right)^2 \tag{5}$$

$\widetilde{h}$ is stochastic and is drawn from a distribution H_j, where j denotes firm-specific types. These stochastic draws are assumed to be independent across periods. This simplifying assumption avoids carrying hours $\widetilde{h}$ as a state variable, and is intended to capture rapidly changing business conditions, which in turn motivates the use of flexible contracts in the low-pay labor market. In equilibrium, the properties of H_j – its mean and variance – yield a ranking of flexible and regular contracts consistent with the taxonomy of firm types in (2).

Actual hours worked, h, depend on the type of labor contract operated by the firm. $h = \widetilde{h}$ under a Z contract, so that this contract enables firms to fully offset the quadratic losses in (5) by effectively making workers work irregular hours. Under a R contract, on the other hand, h is constant over time and is set fixed equal to an exogenous level of hours $\bar{h}$. This feature of R contracts, together with the assumption of quadratic losses, implies that firm's expected profits are negatively related to the variance of $\widetilde{h}$.

Search frictions. Job seekers and vacant positions are brought together via a random search process. We assume that workers search off and on the job provided that the incentives to do so are strictly positive. Thus, type-1 workers, who accept all job offers when not employed, search on the job when employed under a Z contract since they would prefer to hold a R contract. Likewise, type-4 workers employed under a R contract search on the job in order to find a Z contract. Type-2 and type-3 workers, on the other hand, do not find it profitable to search on the job. Letting $x > 0$ denote on-the-job search intensity, in equilibrium it must be that $x_i = x$ for $i \in \{1, 4\}$ and $x_i = 0$ for $i \in \{2, 3\}$. While searching on the job, we assume that workers switch to a different firm if and only if the benefit of doing so is strictly positive, i.e. we rule out job-to-job transitions within the same contract type. Another implication of this setting, which is key to capture ZHCs's participation effect, is that type-2 workers would not search *at*

[18] The environment faced by firms can also be described in the following way. A firm receives orders $\widetilde{q}$ and produces output $q = ph$, where p is the hourly labor productivity. It incurs a convex reputation costs $\underline{d}\,(q - \widetilde{q})^2$ when it produces less than its orders. If the firm wants to sell more than its orders, i.e. if $q > \widetilde{q}$, it needs to spend on marketing to sell the new orders, and the cost is given by $\overline{d}\,(q - \widetilde{q})^2$. For simplicity, we assume that $\underline{d} = \overline{d} = d$. The instantaneous profit of firms is then $q - wh - d\,(q - \widetilde{q})^2$, which boils down to $\pi(h, \widetilde{h})$ in Equation (5) after defining $\widetilde{h} = \tilde{q}/p$ and $\phi = 2dp^2$.

all if Z contracts were to be banned, since their strictly preferred option would be to remain without employment.

On the other side of the market, firms attract workers (either not employed or those employed and searching on the job) by posting vacancies at a per-period cost $\kappa > 0$. Here we make a key assumption: firms must advertise the type of contract that they offer upon posting a vacancy.[19] Hence the trade off faced by firms is as follows. Suppose a firm posts a Z contract, and there are many workers in the labor market who dislike Z contracts. While the firm expects higher per-period profit once the position will be filled, it faces a low acceptance rate of its posted vacancy and/or a high probability that the worker will eventually leave to another firm. This is a standard trade off between maximizing profit flows vs. attracting more or retaining workers over longer employment spells. The converse occurs for firms that choose to post a R contract, if there are many job seekers who prefer the stable work schedule afforded by R contracts.[20]

At the aggregate level, the number of contacts per unit of time depends on market tightness θ, the ratio between the number of vacant positions and job seekers (weighted by the search intensity). The contact rate for non-employed workers is $\lambda(\theta)$, where $\lambda(.)$ is an increasing, concave function of θ. For employed workers of type i, the contact rate with vacant positions is $x_i\lambda(\theta)$. The job-filling rate, i.e. the probability that a randomly chosen vacant job meets a randomly chosen job seeker, is $\lambda(\theta)/\theta$, which is a decreasing and convex function of θ. All jobs are destroyed with a per period probability δ. When this happens, the firm must leave the market. Conditional on not being hit by the δ shock, a firm's position becomes vacant if the worker leaves to another firm. In this event since the job is not destroyed, the firm remains in the market and re-advertises its now vacant position.[21]

In order to enter the market, firms must incur a one-off, business creation cost $K > 0$. The measure of active firms is endogenous and is pinned down by an equilibrium free entry condition (details follow).

4.2 Bellman equations

We now turn to the expression of the asset values of workers and firms. These values depend on exogenous parameters as well as on equilibrium market tightness θ and the equilibrium cross-sectional distribution of agents. We denote as $e_{i,j}$ the measure of job matches between workers of type $i = 1, \ldots, 4$ and firms of type $j = c, s, r$; n_i the measure of unemployed workers of type i; v_j the measure of vacancies of firms of type j; $n = \sum_i n_i$ the aggregate measure of unemployed workers; and $v = \sum_j v_j$ the aggregate measure of vacancies. To define the relevant probabilities used by agents to form expectations, we also define $e_j = \sum_i e_{i,j}$ the measure of jobs at firms

[19] It is also assumed that firms are fully committed to the posted contract. Thus we rule out the possibility of advertising a job as a R contract and eventually offering a Z contract upon meeting, and vice versa. Commitment to the posted terms of trade is a common assumption in wage posting models; see Rogerson et al. [2005].

[20] Note that the compensating differential channel – lower wages in exchange of higher flexibility – is ruled out by the assumption that all jobs in this segment of the labor market pay the statutory minimum wage w. This assumption seems realistic: for example the modal hourly wage in Datta et al. [2019]'s survey on ZHCs is £8, with a large proportion of individuals being paid close to the NLW rate of £7.50.

[21] This is unlike the textbook frictional labor-market model which does not distinguish between job destruction and the exit of firms from the market.

of type $j = c, s, r$; $e_i = \sum_j e_{i,j}$ the measure of employed workers of type $i = 1, \ldots, 4$; e_k the measure of jobs with contract $k = Z, R$; and v_k the measure of vacant positions advertised as $k = Z, R$.

For workers of type i, we denote as N^i the value of non-employment and W_k^i the value of being employed in a contract $k = Z, R$. These asset values solve:

$$N^i = u_N + \frac{1}{1+\rho} \left[(1 - \lambda(\theta)) N^i + \lambda(\theta) \sum_{k'} \frac{v_{k'}}{v} \max\left\{N^i, W_{k'}^i\right\} \right], \tag{6}$$

and

$$W_k^i = u_k^i + \frac{1}{1+\rho} \left[\delta N^i + (1-\delta) \left((1 - x_i \lambda(\theta)) W_k^i + x_i \lambda(\theta) \sum_{k'} \frac{v_{k'}}{v} \max\left\{W_k^i, W_{k'}^i\right\} \right) \right]. \tag{7}$$

The "max" operator in Equations (6) and (7) captures the decision to accept or reject an offer $k' = Z, R$ from either non-employment or employment. The variable x_i, which is a policy function, must be consistent with the worker's own search decisions in the sense that $x_i = x > 0$ if and only if there exists a contract k' such that $W_{k'}^i > W_k^i$. The probability that the offer is a labor contract $k' = Z, R$ is $v_{k'}/v$. Lastly, in Equation (6) the flow value of non-employment is $u_N = u^i(0, a)$, which is the same across all worker types i. Furthermore, the flow value of employment u_k^i in Equation (7) differs across worker types and depends on the equilibrium mix of firm types:

$$u_Z^i = \frac{e_c}{e_Z} \int u^i\left(\tilde{h}, a\right) dG_c\left(\tilde{h}\right) + \frac{e_s}{e_Z} \int u^i\left(\tilde{h}, a\right) dG_s\left(\tilde{h}\right) \quad \text{and} \quad u_R^i = u^i\left(\bar{h}, a\right) \tag{8}$$

There is an obvious simplification in (8): we rule out workers' learning about the specific type of their own employer so as to compute the expected flow value of employment in a Z contract directly based on the equilibrium mix of type-c and type-s firms among all Z employers. We think learning is an interesting issue on its own but is unlikely to be of first order importance for the quantification of the effects of ZHCs.

Next, for firms of type j, V^j denotes the asset value of holding a vacant position advertised as a contract $k = Z, R$, and $J_{i,k}^j$ is the asset value of filling this position with a type-i worker. Firms' asset values of advertising a vacant position is given by:

$$V_k^j = -\kappa + \frac{1}{1+\rho} \left[V_k^j + \frac{\lambda(\theta)}{\theta} \sum_i \frac{n_i \mathbb{1}_{\{W_k^i > N^i\}} + \sum_{j'} x_i e_{i,j'} \mathbb{1}_{\{W_k^i > W_{k(j')}^i\}}}{n + \sum_{i'} x_{i'} e_{i'}} \left(J_{i,k}^j - V_k^j\right) \right], \tag{9}$$

where $\mathbb{1}_{\{.\}}$ is the indicator function. Conditional on meeting a worker, which occurs at rate $\lambda(\theta)/\theta$ for vacant jobs, the probability that she will accept contract k depends on her current labor market status and preferred employment contract. For filled jobs, the asset values solve:

$$J_{i,k}^j = \pi_k^j + \frac{1-\delta}{1+\rho} \left[V_k^j + \left(1 - x_i \lambda(\theta) \sum_{k'} \frac{v_{k'}}{v} \mathbb{1}_{\{W_{k'}^i > W_k^i\}} \right) \left(J_{i,k}^j - V_k^j\right) \right]. \tag{10}$$

In this equation, the probability that the job remains filled depends on the equilibrium offers from other firms (through $v_{k'}/v$) and on workers' preferences over those offers. The flow values of employing a worker under contract k are:

$$\pi_Z^j = \int \pi\left(\widetilde{h}, \widetilde{h}\right) dG_j\left(\widetilde{h}\right) \quad \text{and} \quad \pi_R^j = \int \pi\left(\bar{h}, \widetilde{h}\right) dG_j\left(\widetilde{h}\right). \tag{11}$$

Finally, observe that in Equation (10) the continuation value is multiplied by $1-\delta$. This is because with probability δ the firm exits the market and its asset value becomes zero under the equilibrium free entry condition.

4.3 Free entry

Firms enter the market until the expect value of doing so is exhausted. They pay a business creation cost K to enter the market and then draw their type j from a distribution $(\gamma_j)_{j=c,s,r}$. In an equilibrium, firms' vacancy posting decision must be consistent with the taxonomy of types in (2): type-c and type-s firms post Z contracts, so that their asset value after paying K becomes V_Z^c and V_Z^s, respectively, while for type-r firms the asset value of a vacancy after market entry is V_R^r. Thus the free entry condition reads:

$$K = \gamma_c.V_Z^c + \gamma_s.V_Z^s + \gamma_r.V_R^r. \tag{12}$$

Market tightness θ, which is the ratio between v and $n + \sum_i x_i e_i$, adjusts to satisfy the above condition.

4.4 Steady-state equilibrium

We are in a position to define a steady-state equilibrium of this economy. A steady-state equilibrium is a list of asset values N^i, W_k^i, V_k^j, $J_{i,k}^j$; a stationary distribution of job matches $e_{i,j}$, non-employed workers n_i and vacancies v_j; and labor market tightness θ such that:

1. Given the measures $e_{i,j}$, n_i, v_j, and market tightness θ , the asset values N^i, W_k^i, V_k^j, $J_{i,k}^j$ solve the Bellman equations (6), (7), (9), (10);

2. Given N^i, W_k^i, worker types satisfy the rankings presented in (1); given V_k^j, $J_{i,k}^j$ firm types satisfy the rankings presented in (2);

3. Given V_k^j, where $k = Z, R$ is the contract offered by type-j firms, market tightness θ solves the free entry condition in Equation (12);

4. Given market tightness θ, the measures $e_{i,j}$, n_i, v_j are time-invariant with respect to the law of motion described in Appendix A.

Condition 2 of the above definition is key to understand the workings of the model. Suppose that this condition holds. Given market tightness θ, the stationary distribution of job matches $e_{i,j}$, non-employed workers n_i and vacancies v_j can be found by iterating on the equilibrium stock-flow equations. Next, the fact that agents' rankings are satisfied allows us to simplify all the "max" operators in Equations (6), (7), (9), (10), because we know workers' preferred options. We can thus solve for workers' and firms' asset values in one step since the Bellman equations define a set of autonomous linear equations. Finally, we must verify that the free entry condition is met, and more importantly that the rankings are consistent with agents' types. What makes the computation challenging is that we cannot impose *ex ante* that the rankings be satisfied.

These observations are also useful to understand the out-of-steady-state dynamics of the model. Consider what would happen along a transition path of the economy. The trajectory of the cross-sectional distribution would directly impact the calculation of the asset values (recall that the distribution matters for the u_k^i's), and more importantly the rankings of both firms and workers might change along the transition path, making the computation intractable. For these reasons, in Section 6 we confine ourselves to making steady-state comparisons, with a view of focusing on the long-run labor market impacts of ZHCs.

5 Calibration and inference on workers' and firms' types

In this section, we calibrate the model and present a first set of results about workers' and firms' types. Our calibration method proceeds in two steps: first, we calibrate the model parameters that are directly related to labor turnover moments; and second, we calibrate the parameters that determine expected firms' profits and workers' utility functions.

5.1 Parameters set externally

The model period is set to be two weeks. We choose $\rho = 0.0015$ to yield an annual discount rate of 4 percent. We set w to the 2017 U.K. national minimum wage for workers aged 25 and over, namely $w = 7.50£$ per hour. The other parameter that we choose using external information relates to the elasticity of the job creation process. We use a standard Cobb-Douglas matching function to determine the number of contacts per unit of time:

$$m(s, v) = M v^{\psi} s^{1-\psi}, \tag{13}$$

where v denotes the number of vacancies and s the number of job seekers weighted by their search intensity. We think of ψ as being an important parameter for the job creation effects predicted by the model. Given this, we estimate ψ using U.K. data for the low-pay segment of the labor market. Our procedure, presented in Appendix B, controls for occupation fixed effects as well as time variation in matching efficiency, and yields estimates of ψ between 0.60 and 0.70. These figures are higher than the value of 0.5 that is commonly used in the literature, but certainly not inconsistent with the bias-corrected estimate of ψ from Borowczyk-Martins et al. [2013]. We use the mid-point of our empirical estimates and set ψ to 0.65.

5.2 First-step calibration parameters

We first calibrate M, θ, δ, x, the ω_i's and γ_r to match data moments on job and worker turnover. Transitions out of unemployment enable us to pin down the job-finding rates, and hence determine either M or θ. Whether workers transit to Z or R jobs then depends on the mix of contracts among posted vacancies, which is in turn informative about γ_r. The more crucial part concerns transitions out of Z and R jobs as they depend on δ, x, as well as on the ω_i's. Our approach relies on close examination of the distribution of job tenure/duration in a given labor market state, which conveys information on the heterogeneity of exit rates, and of the transition matrix across N, Z and R (Table 4). The reason why the distribution of job tenure conveys additional information is the following. If the job distribution data are well matched by an exponential survival function, then we can fit it well by a homogeneous exit rate. Otherwise, we need at least two worker types to fit the data satisfactorily. It turns out that, given the data at hand, we only detect the presence of type-1 and type-2 workers. This result is intuitive. According to Table 4, there are virtually no transition from R to Z jobs. This rules out the presence of type-4 workers, i.e. $\omega_4 = 0$. Consistent with this, the job tenure distribution of R contracts suggests a homogeneous exit rate. The model can replicate this pattern by having only type-1 workers, only type-3 workers, or a combination of them (as they both have exit rate δ from R contracts). On the other hand, the job-tenure distribution of Z contracts can only be matched by having a combination of type-1 and type-2 workers in these jobs. Thus $\omega_1 > 0$. Finally, to make it profitable for type-r firms to post R contracts, we need to have sufficiently many workers *among job seekers* who are willing to accept R contracts. Only type-1 workers would satisfy this criterion, given that they search on the job while employed in Z contracts. This pushes up ω_1, and since $\omega_1 + \omega_2 + \omega_3 = 1$, the share ω_3 decreases towards 0. In practice, we set $\omega_3 = \omega_4 = 0$, so that in the calibration we are left with searching for ω_1 to match the data and set $\omega_2 = 1 - \omega_1$. Last, we need another data moment to separately identify M and θ. For this, we use U.K. estimates of the job-filling rate (the probability of filling a vacancy) from Kuhn et al. [2021]. They estimate that the monthly job-filling rate for the U.K. is between 0.35 to 0.38. At the bi-weekly frequency, this yields a target for $\lambda(\theta)/\theta$ of 0.20. We calibrate M, θ, δ, x, ω_1, γ_r jointly to minimize the distance to the distributions of job tenure in Z and R jobs, the transition matrix in Table 4, and the targeted job-filling rate. Notice that at this point we obtain the sum $\gamma_c + \gamma_s$, but not γ_c and γ_s separately from each other.

Figure 3 compares the distributions of job tenure in the LFS data with the model-generated ones. For completeness, the bottom panel of this figure reports the distribution of unemployment duration, although it is not directly targeted by our calibration. The key observation is that the distribution of job tenure in Z contracts can only be matched with heterogeneous workers within this labor market state. Some of these workers (namely type-1 workers) quit at a higher rate, which explains why we find more workers at short tenure (less than 6 months) than at slightly longer tenure (6 to 12 months). Other workers (type-2) remain much longer in these jobs, which explains why we find larger shares of workers at job tenure longer than 2 and 5 years. This phenomenon is much less pronounced in the data in what concerns R contracts. As a result, we obtain a satisfactory fit with homogeneous workers within these jobs. Panel

(b) of Table 6 reporting the parameter values that come out of this first calibration step shows that, in order to match the data, we need a large share of type-1 workers and a large probability of drawing type-r for firms upon entry: ω_1 is 0.97 and γ_r is 0.95. These numbers are consistent with the fact, described in Section 3, that ZHCs make up for a small share of employment, even in the low-pay segment of the labor market.

Table C2 in the appendix performs a similar comparison exercise for the quarterly transition matrix across N, Z and R. It shows that the model overfits transitions from N to Z and slightly understates those from Z to R. On the other hand, the model matches perfectly all the outflow rates from R jobs.

Table 5: Description of baseline equilibrium

		Model	**Data**
n	Unemployment rate	9.2	10.1
$e_Z/(e_Z+e_R)$	Employment share of ZHCs	6.5	7.2
$v_Z/(v_Z+v_R)$	Vacancy share of ZHCs	19.4	–
$e_{1,Z}/e_1$	Share of employed type-1 workers in ZHCs	4.8	–
$e_{1,Z}/e_Z$	Share of filled ZHCs employing type-1 workers	66.8	–

Notes: The table reports model-predicted moments for the baseline steady-state equilibrium, and moments computed from the LFS data based on the ergodic distribution of the transition matrix across N, Z, R (Table 4). All table entries are expressed in percent.

Table 5 presents several moments describing the equilibrium allocation that results from the first step of the calibration. We compare the first two of these moments to the ergodic distribution of the empirical transition matrix across N, Z and R shown in Table 4. The model's baseline unemployment rate is 9.2 percent, which is fairly similar to the value implied by the empirical transitions in and out of unemployment (10.1 percent). The model's ZHC share of employment is lower than the value implied by the empirical transition matrix (7.2 percent), but note that the latter is high compared to that of respondents in the LFS who report being employed under a ZHC (5.2 percent of employed workers during either the first or the second interviews; see Section 3). The other moments reported in Table 5 illustrate the sorting patterns predicted by the model. First, while ZHCs make up for a small share of employment (6.5 percent), they account for almost one-fifth of all vacant positions. This is due to higher worker turnover on these contracts, implying that they get readvertised more frequently compared to R contracts. Notice that, in a random search environment, this implies that ZHCs exert a negative effect on R vacancies by making it more difficult for these vacancies to contact workers who prefer regular employment. We will return to this point in Section 6. Second, since most vacant positions are R contracts, and because of on-the-job search, there are fewer than 5 percent of employed type-1 worker who hold a Z contract. Meanwhile, since type-1 workers make up for a very large share of the labor force, they account for two thirds of all ZHC jobs. Thus, type-1 workers are key to sustain an equilibrium with ZHC jobs.

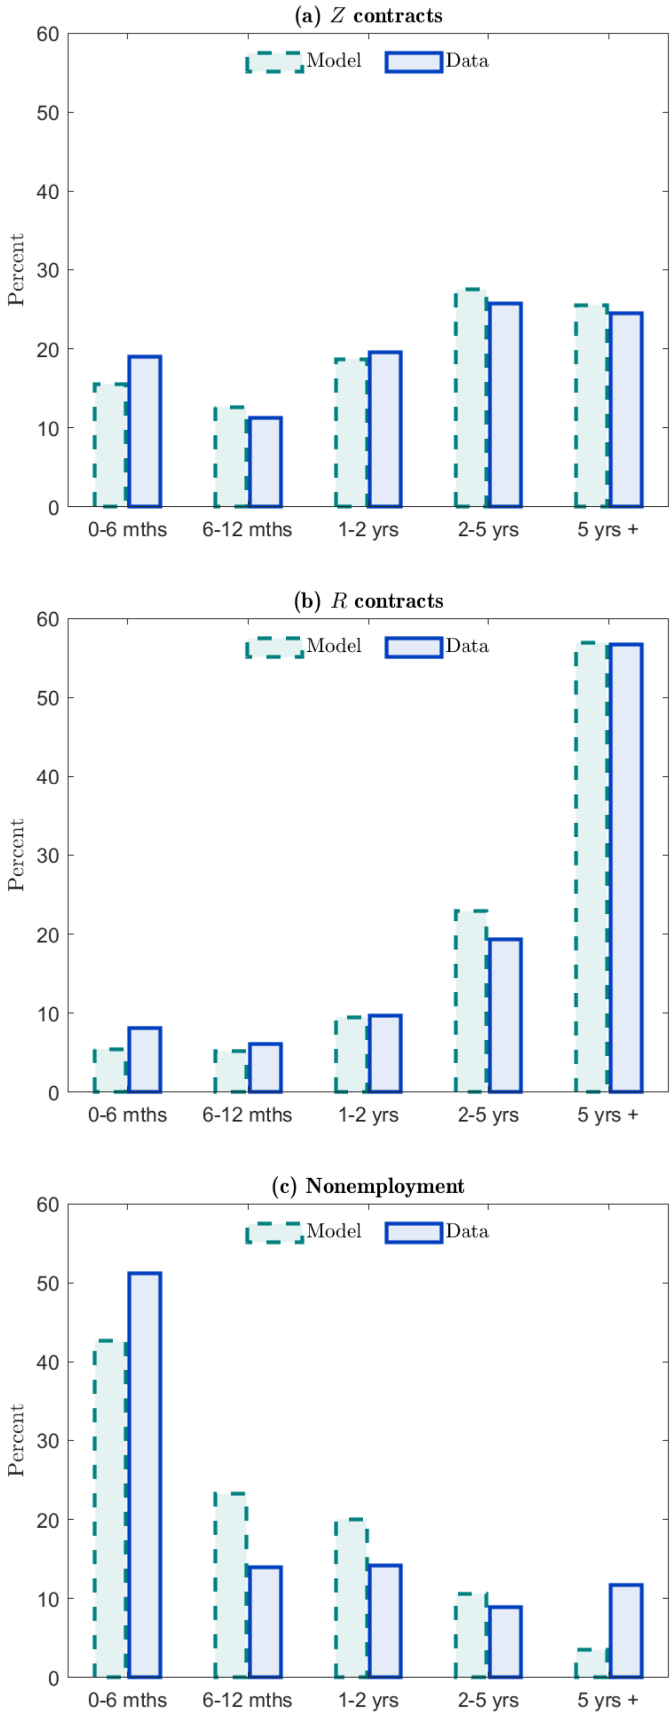

Figure 3: Model fit: Job tenure and distribution of unemployment duration

Notes: Dashed bars: Model's predicted moments; Solid bars: authors' calculations based on data from the Labour Force Survey. Panel (a) reports the distribution of job tenure in Z contracts. Panel (b) reports the distribution of job tenure in R contracts. Panel (c) reports the distribution of the duration of unemployment spells.

5.3 Second-step calibration parameters

Given the first-step calibration parameters, we now examine parameters that relate to firms' hours worked and production costs. We must calibrate the following parameters: p, $\bar{h}$, ϕ, κ, K, and the stochastic distributions $H_j(.)$ for each type j. We assume that the $H_j(.)$'s are Beta distributions over the interval [0, 50], where 50 refers to the upper bound on weekly hours worked. The reason we use Beta distributions is that they are flexible in terms of generating bell-shaped distributions, bi-modal distributions, etc. and the parametrisation relies on the choice of two moments only, such as the mean and standard deviation of the distribution. Thus, in addition to the above parameters, we must calibrate the means μ_c, μ_s, μ_r and standard deviations σ_c, σ_s, σ_r of these distributions.

We set $p = 8.25$, thus effectively assuming that the marginal productivity of workers in the low-pay segment is 10 percent higher than the statutory minimum wage, w. We choose $\bar{h} = \mu_r$, which is a normalization in the sense that the gap between $\bar{h}$ and $\tilde{h}$ in the profit function (5) is scaled by ϕ. From Table 2, we set $\mu_r = 28$, and choose $\mu_c = \mu_s = 18$ to capture the 10 hours gap between the two contract types.[22] At this point, we are left with the task of choosing ϕ, κ, K, and the standard deviations of hours, σ_c, σ_s, σ_r. There is no direct empirical counterpart for ϕ, but note that any combination of parameter values for ϕ and κ pins down the value of K through the free-entry condition (12). Thus we set up calibration targets for κ and K, and then search for the value of ϕ that allows the model to match these targets. For κ, we follow Elsby and Michaels [2013] and target an expected cost of vacancy posting (i.e. $\kappa\theta/\lambda(\theta)$) that amounts to 14 percent of average quarterly labor earnings. We find that $\kappa = 38$£ (per week) achieves this objective.[23] As regards the startup costs of creating a business, K, estimates for the U.K. suggests that it averages around £22,500.[24] 95 percent of businesses in the U.K. have between 1 and 10 employees (House of Commons [2021]). Assuming that the average business has 5 workers, and since each firm in our model has only one job, this suggests a target for K at around £4,500. We come very close to this target by setting $\phi = 0.16$; we obtain $K = 4,376$£. The calibrated ϕ implies that the cost of deviating from $\tilde{h}$ by 5 hours in a given week reduces firms' accounting profits ($(p - w)h$) by 10 percent.

The calibration of ϕ, κ, K is of course not independent of σ_c, σ_s, σ_r. For these parameters, however, we must choose them in a way that delivers rankings of the different contracts (i.e. V_Z and V_R) consistent with the taxonomy of firm types in (2). As will be shown in the next section, we find that $\sigma_c = 6$, $\sigma_s = 3$, $\sigma_r = 2$ achieve this objective. There is some arbitrariness

[22] The assumption that $\mu_c = \mu_s$ is partly motivated by the fact that our model precludes job-to-job turnover within Z employment. Our reasoning in favor of this is that, although workers can easily update beliefs about the distribution of hours in their current job by observing the number of working hours every period, getting an accurate estimate of the hours variance would be much harder. Hence, ZHC workers would not be able to ascertain whether they work for a type-c firm and, as a result, accept outside offers of another Z contract from firms which could be either of the c or the s types.

[23] Although Elsby and Michaels [2013]'s figure for vacancy posting cost is computed out of U.S. data, it seems well in line with numbers for the U.K. See, for example, https://theundercoverrecruiter.com/true-costs-hiring-uk/: they estimate that the advertising cost using social media and job sites is between £200 and £400 per new hire. Our calibration yields $\kappa\theta/\lambda(\theta)$ equal to £362.

[24] See https://www.capalona.co.uk/blog/how-much-does-it-cost-to-start-a-business/ and https://www.telegraph.co.uk/business/sme-home/start-up-costs/. We are not aware of any official statistics, but the evidence gathered from the Internet puts the value of K between £20,000 and £30,000.

in this choice, which seems unavoidable given the lack of additional empirical information on hours worked in ZHC jobs. Our choice is motivated by the observation that hours worked are more volatile in ZHC jobs, as shown in Table 2, and that the rankings of V_Z and V_R imply that σ_c must be larger than σ_s if type-c and type-s firms have the same mean of hours (i.e. $\mu_c = \mu_s$). We conduct an extensive robustness analysis of the sensitivity of our results to changes in σ_c and σ_s (under the constraint that (2) continues to hold) and find that our conclusions are not sensitive to these choices. As reported in Appendix D, we show more generally that our results are robust to changes to the means μ_c, μ_s, μ_r and standard deviations σ_c, σ_s, σ_r, provided that ϕ, κ, K are jointly recalibrated to match their targets. Figure C1 in the appendix shows the distribution $H_j\,(.)$'s of the baseline calibration.

Table 6: Parameter values

(a) Parameters set externally		
ρ	Discount rate of 4 percent per annum	0.0015
ψ	Elasticity of job-filling rate w.r.t. tightness (Appendix B)	0.65
w	Minimum hourly wage in £ (U.K. policies)	7.50
(b) First stage calibration parameters		
M	Matching function elasticity	0.1278
θ	Labor market tightness	0.2418
δ	Job destruction probability	0.0047
x	On-the-job search efficiency	0.3524
ω_1	Share of type-1 workers	0.9689
γ_r	Probability of type-r firms upon entry	0.9498
(c) Second stage calibration parameters		
p	Productivity of hours worked	8.25
(μ_c, μ_s, μ_r)	Average of weekly hours by firm type (Table 2)	$(18, 18, 28)$
$(\sigma_c, \sigma_s, \sigma_r)$	St. dev. of weekly hours by firm type (Table 2)	$(6, 3, 2)$
ϕ	Marginal cost of deviating from targeted hours	0.16
κ	Flow cost of vacancy posting, in £ per week	38.0
K	Startup cost of new businesses, in £1,000	4.38

Notes: Panel (a) reports parameters that are calibrated outside the model. Panel (b) reports parameters calibrated using data moments on job and worker turnover. Panel (c) reports parameters calibrated using information on hours worked and labor costs, under the assumption that $\gamma_c = \gamma_s$. The model period is set to be two weeks. Consequently, weekly hours worked and the flow vacancy posting cost are multiplied by 2 before plugging into agents' profit and utility functions.

Table 6 reports the outcomes of the calibration. Notice that, through Equation (12), the parameters reported in panel (c) can only be calibrated by choosing values for γ_c or γ_s. On the other hand, parameters reported in panel (b) of the Table depend only on $\gamma_r = 1 - \gamma_c - \gamma_s$ which is determined through the first step of the model's calibration. In Section 6, we will cover the whole spectrum of values for γ_c and γ_s such that $\gamma_c/(\gamma_c + \gamma_s)$ varies between 0 and 1. In all instances, we always recalibrate the parameters in the bottom panel of Table 6, but in general we find that they vary little with the assumption we make about γ_c or γ_s. This is intuitive, since γ_r is close to 1. For the sake of exposition, in Table 6 and in the results shown

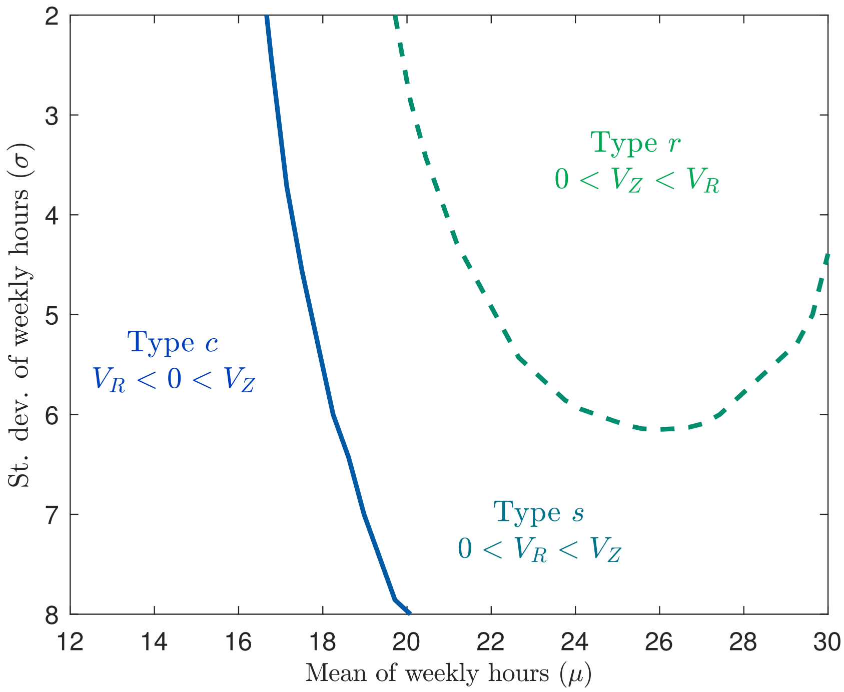

Figure 4: Firms' types across (some) regions of the parameter space

Notes: The figure shows firms' rankings of the asset values of advertising Z jobs (V_Z) and R jobs (V_R), as a function of firms' mean μ_j and standard deviation σ_j of the demand of weekly hours worked.

in the next section, we use $\gamma_c = \gamma_s$.

5.4 Heterogeneous workers' and firms' types

We begin with the analysis of firms' types. Figure 4 depicts the different regions of the parameter space μ_j and σ_j that lead to different preferences over ZHCs and regular contracts. For instance, the top-right corner corresponds to firm types r, for which the mean of $\tilde{h}$ is quite high and its standard deviation relatively low. This seems plausible in the sense that firms that would always prefer to post R contracts might be facing less volatility in consumer's demand. When either the mean of $\tilde{h}$ is lower or the standard deviation is higher, the firm type is s, as depicted in the Figure. These firms post a ZHC, but would be viable under a regular contract. In the area on the left part of the figure, where the mean of hours is lowest, firms are of type c, i.e. they are only viable when posting Z contracts. These type definitions are consistent with intuition, in that the higher and more stable hours are, the more likely the firm is to be of type r. Also, observe that by choosing $\mu_c = \mu_s = 18$ hours, our choice set for σ_c and σ_s is quite limited, as we must choose them along the vertical line $\mu = 18$ below and above the curve that separates out type-c and type-s firms. Appendix D shows that our results remain broadly similar when we allow μ_c and μ_s to differ from each other.

Next, we turn to the analysis of worker types. Figure 5 plots regions of the parameter space where on the horizontal axis different values of available working hours a are considered, and on the vertical axis the utility from short hours is varying. Specifically, the numbers shown on the vertical axis are the marginal willingness to pay (MWP; α_i over the marginal utility of consumption) to avoid working over and above available hours a. To calibrate the model, we

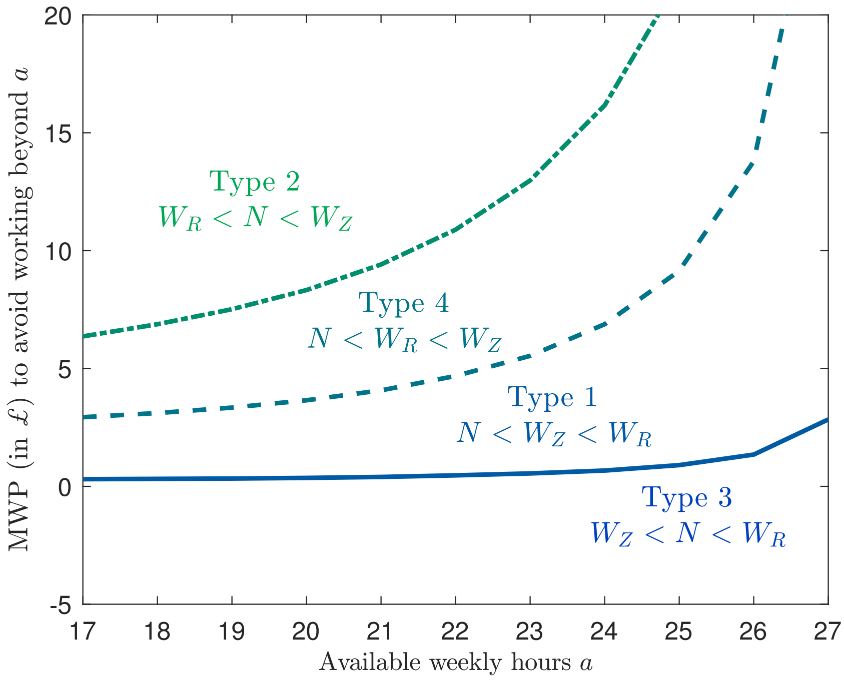

Figure 5: Workers' types across (some) regions of the parameter space

Notes: The figure shows workers' rankings of the asset values of working a Z contract (W_Z), R contract (W_R), and the value of not working (N), as a function of workers' available weekly hours a and disutility of work, where the latter is expressed as workers' marginal willingness to pay to avoid working hours beyond a (i.e. α_i over the marginal utility of consumption, and consumption is set equal to workers' average weekly labor income).

choose available hours a to be 22 weekly hours, which we interpret as four days of 5 hours of daily work plus half an hour of travel-to-work per day. According to Figure 5, type-1 workers are individuals who do not value short hours very much, especially at $a = 22$ hours per week. The type-1 worker with the lowest α_i would give up £0.5 of consumption (per week) to avoid working one hour beyond a, while at the highest α_i the worker would give up 'only' £4.6 of her consumption. Quite sensibly, type-2 workers have a higher valuation of short hours. They would be willing to give up *at least* £10.9, or 1.45 times the minimum wage, to avoid working one hour over their available hours a. Notice that the MWPs are increasing functions of a: the higher available working hours a, the lower the marginal utility of consumption that the worker gets from working 1 hour beyond a hours, and therefore the higher her willingness to pay to avoid working this extra hour. For instance at $a = 24$ hours per week, the upper bound for type-1 workers is a MWP of £6.9 (thus close to the minimum wage), and the lower bound for type-2 workers is £16.2. According to Figure 5, type-4 workers have preferences for short hours between those of type-1 and type-2 workers, and type-3 workers are typically workers who dislike short hours.

Table 7 presents the parameter values that we choose to pin down workers' utility flows. The coefficient of relative risk aversion η is set to 2.0, which is a standard value in the literature. As just mentioned, we set available hours per week a to 22 hours.[25] According to the OECD, the replacement ratio of UI benefits for those previously employed at the minimum wage, living

[25] We show in Appendix D that our results are robust to deviations of a above and below the value used to run the main experiments.

Table 7: Parameter values for welfare assessment

η	Relative risk aversion coefficient	2.0
b	Unemployment benefits in £ per week (U.K. policies)	148.5
τ	Taper rate (U.K. policies)	0.63
a	Available hours per week	22.0

Notes: The table reports the parameter values that are used for the assessment of the welfare effects of ZHCs. The model period is set to be two weeks. Consequently, weekly hours worked and the flow vacancy posting cost are multiplied by 2 before plugging into agents' profit and utility functions.

in a couple with two children and a partner earning the average wage, is 80 percent.[26], [27] We set b to £148.5 per week, which exactly match 80 percent of average weekly earnings wh in equilibrium. The taper rate τ is set to 63 percent in line with U.K. policies. Notice a key feature of our model: instead of choosing specific values for α_i, we only need to assume that they fall in the ranges implied by Figure 5. In the quantification of the welfare effects, we will report results evaluated over these ranges of values.

6 Policy experiments

This section contains the main discussion of the pros and cons of ZHCs. We simulate the impact of several counterfactual policies, and use the results as a basis for discussing policy issues related to the development of ZHCs and similar labor contracts.

6.1 Impact of a minimum wage rise

We begin with a simple comparative static analysis of the effects of a minimum wage rise according to our model. The motivation for this exercise comes from Datta et al. [2019] who estimate that the 2016 rise in minimum wage was accompanied by an increase in the use of ZHCs in low wage sectors in the U.K.

The arrows in Figure 6 show the effects of a 3 percent increase in the statutory minimum wage, w. Most importantly, we observe that the region depicting the combinations of mean and standard deviation of hours for which a firm is of type r shrinks when the minimum wage increases. This implies that firms' use of ZHCs would increase following an increase in the minimum wage. We also see that the composition of ZHC employers would change: firms of type c would become more numerous, that is to say firms that can be viable only through the use of ZHCs. We cannot, however, quantify these effects since we do not have firm data that would allow us to estimate the mass of firms that are present in the different regions of Figure

[26] See `https://stats.oecd.org/Index.aspx?DataSetCode=NRR`. In the OECD's net replacement rate calculator, we set Family type to "Couple with 2 children - partner's earnings: AW" and Previous in-work earnings to "Minimum wage".

[27] As reported in Appendix D, a lower value of unemployment benefits amplifies the effects of a policy reform, such as a ban on ZHCs, that leads to more frequent spells of unemployment.

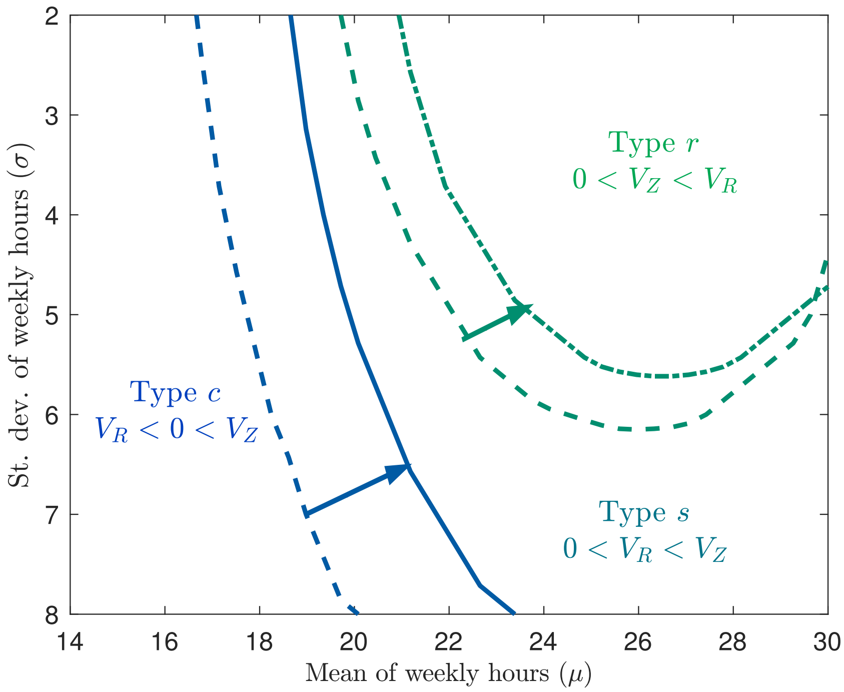

Figure 6: Effect of a rise in the statutory minimum wage on firms' types

Notes: The figure shows firms' rankings of the asset values of advertising Z jobs (V_Z) and R jobs (V_R), as a function of firms' mean μ_j and standard deviation σ_j of the demand of weekly hours worked. The horizontal arrows shows the effects of a rise in the statutory minimum wage on the thresholds that define each type of firms.

6. Yet, since the qualitative story is consistent with Datta et al. [2019]'s minimum wage results, we view it as a useful "over-identifying" test of the model.

6.2 Ban on ZHCs

Since this proposal is at the heart of the media debate on ZHCs, we simulate a counterfactual policy consisting of a ban on these contracts. We carry out simulations under varying assumptions regarding the fraction of type-c and type-s firms among firms that choose to offer ZHCs in the baseline equilibrium. Results are reported in Table 8.

Consider first the results reported in the first column next to 'Baseline'. In this simulation, we assume that $\gamma_c/(\gamma_c+\gamma_s)$ is 0, so that all ZHCs jobs are offered by type-s firms. The ban on ZHCs leads to an increase of the unemployment rate by 2 percentage points (p.p.). The main driver is that type-s firms face lower expected profits following a ban on ZHCs (they can no longer substitute ZHCs for regular contracts conditional on filling a vacancy), and as a result there are fewer entries of firms, making the unemployment rate increase. At the same time, the employment rate drops by a larger amount: it decreases by 4.8 p.p. The difference is accounted for by the so-called participation effect: type-2 workers drop from the labor force following a ban on ZHCs. The magnitude of the effect seems plausible, and indirectly this validates the calibration's outcome that assigns a small value to ω_2. Since all workers who remain active in the market are willing to accept a job offer when they receive one, the duration of posted R vacancies decreases (by about 2.5 weeks). Fewer resources are devoted to vacancy posting, but the overall impact of the ban on output is negative, due to lower firm entry in equilibrium.

In the other columns of the table, we analyze the effect of a ZHC ban where some ZHCs are offered by type-c firms, making $\gamma_c/(\gamma_c+\gamma_s)$ increase gradually all the way to $\gamma_c/(\gamma_c+\gamma_s)=1$. By doing so, we assume a comparatively larger role for the job creation effect of ZHCs. In the rightmost column, a ban on ZHC causes the unemployment rate to increase by 2.7 p.p. and the employment rate to drop by 5.4 p.p. These are admittedly large effects, but bear in mind that our analysis is confined to the low-pay segment of the labor market. According to the model, when the job creation effect is largest, a ban on ZHC leads to a 4 percent decrease in net output in this sector of the economy.

Accession to regular employment. In the equilibrium without ZHCs, even though the unemployment rate is higher, regular employment (R) is also higher. Quantitatively the difference is small – R employment increases by 1.5 percent –, but understanding the sources of this difference helps to illuminate the mechanisms affecting labor reallocation following a ban on ZHCs. Our model offers a simple accounting for this purpose. Let $\widetilde{e}_R$, $\widetilde{\theta}$ and $\widetilde{n}_1$ denote respectively regular employment, market tightness and the measure of unemployed type-1 workers in the absence of ZHCs. We can write the ratio between regular employment without and with ZHCs, i.e. $\widetilde{e}_R/e_R$ as the product of three components:[28]

$$\frac{\widetilde{e}_R}{e_R}=\underbrace{\frac{\lambda(\widetilde{\theta})}{\lambda(\theta)}}_{\text{job creation}}\times\underbrace{\frac{1}{v_R/v}}_{\text{vacancy competition}}\times\underbrace{\frac{\widetilde{n}_1}{(1-\delta)\,xe_{1,Z}+n_1}}_{\text{search efficiency}}. \tag{14}$$

The first component measures the effect of job creation: there are fewer firms entering the market after a ban on ZHCs, which reduces labor market tightness and hence regular employment through a lower aggregate job-finding rate. In the experiment with $\gamma_c=\gamma_s$ for instance, we find that $\lambda(\widetilde{\theta})/\lambda(\theta)$ equals 71 percent, meaning that lower job creation would reduce regular employment by almost 30 percent *ceteris paribus*. The second component measures the effects of competition between different type of vacancies. In the equilibrium with ZHCs, R vacancies are diluted in the pool of vacancies, which makes it more difficult for firms to find workers who would accept the offered contract. Indeed, recall from Table 5 that although ZHCs make up a small share of employment, they account for almost 20 percent of all vacancies. The reduction in vacancy competition in isolation from the other effects would increase regular employment by 24 percent. Lastly, there is an increase in search efficiency units for R employment following a ban on ZHCs. Prior to the ban, type-1 workers who are employed on ZHCs are less effective in contacting R vacancies, given that on-the-job search efficiency x is well below 1. We find that this effect contributes an increase in R employment by 15 percent *ceteris paribus*.

[28]Recall that according to our calibration, R employment is fully accounted for by type-1 workers. Use Equation (A.5) that describes the law of motion of $e_{1,r}$ to compute the steady-state stock of R employment:

$$e_R=e_{1,r}=\frac{1}{\delta}\lambda\,(\theta)\,\frac{v_R}{v}\,(x\,(1-\delta)\,e_{1,Z}+n_1)\,,$$

since $v_r=v_R$. Moreover, in the equilibrium without ZHCs we have: $\widetilde{e}_R=\frac{1}{\delta}\lambda(\widetilde{\theta})\widetilde{n}_1$. We obtain Equation (14) by taking the ratio of these two equations.

Table 8: Equilibrium and welfare effects of a ban on Z contracts

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(a) Equilibrium (re)allocation						
Employment rate (in %)	90.8	86.0	85.9	85.7	85.6	85.4
		-4.79	*-4.93*	*-5.08*	*-5.23*	*-5.40*
Unemployment rate (in %)	9.2	11.2	11.4	11.5	11.7	11.8
		2.03	*2.17*	*2.33*	*2.48*	*2.65*
Duration of R vacancies (in weeks)	10.5	8.0	8.0	7.9	7.8	7.8
		-2.47	*-2.53*	*-2.59*	*-2.66*	*-2.73*
Net output (1 = baseline)	1.00	0.98	0.97	0.97	0.97	0.96
(b) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-0.93	-0.99	-1.06	-1.12	-1.15
At 25th percentile of α	0.00	-0.96	-1.02	-1.08	-1.13	-1.14
At 50th percentile of α	0.00	-0.99	-1.04	-1.09	-1.15	-1.13
At 75th percentile of α	0.00	-1.02	-1.07	-1.11	-1.16	-1.13
At 99th percentile of α	0.00	-1.05	-1.09	-1.13	-1.17	-1.12

Notes: Panel (a) of the table reports the equilibrium allocation effects of a ban on ZHCs. Panel (b) of the table reports the welfare consequences for workers who remain in the labor market after a ban on ZHCs. The first column called 'Baseline' describes the equilibrium with ZHCs. The other columns describe the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

What is the overall impact of the ban on time spent *out* of regular employment? After the policy reform, this duration is given by the duration of unemployment spells. Prior to the ban, this is the duration that type-1 workers spent in unemployment as well as in Z employment waiting to eventually transit to R employment. The difference between these durations, which we denote as $\triangle$, is readily measured in our model, since[29]

$$\triangle = \frac{\omega_1}{\lambda(\widetilde{\theta})\widetilde{n}_1} - \frac{\omega_1}{\lambda(\theta)\frac{v_R}{v}\left(x\left(1-\delta\right)e_{1,Z}+n_1\right)} \approx -7 \text{ weeks} \tag{15}$$

in the main experiments. Thus, even though the unemployment rate increases after a ban on ZHCs, type-1 workers spend on average more time in regular employment. That is, lacking a stepping stone towards regular employment, workers would face higher unemployment, and at the same time face more stable employment.[30] The implications on welfare are not obvious, however. In the equilibrium with ZHCs, when type-1 workers are not in regular employment, they spent only a fraction of that time in unemployment and spend the remainder in Z employment, which makes them better off compared to being unemployed. This difference must be factored in to compute the overall welfare effect. This is the issue we turn to in the next paragraphs.

Welfare effects. The lower panel of Table 8 provides an assessment of the welfare consequences of a ban on ZHCs on type-1 workers, who remain in the labor market after the policy reform.[31] In order to evaluate these consequences, we assume that these workers are uniformly distributed across the utility parameters α_i that correspond to this worker type (see Figure 5). The specific assumption of a uniform distribution is irrelevant for any of the equilibrium effects computed in the upper panel of Table 8. It only matters to make statements about the distributional consequences of a ZHC ban among these workers. The Table shows that a ban on ZHCs has a wholly negative impact on workers' welfare. Depending on their own valuation of short hours, and on the strength of the job creation effect, they suffer a welfare loss that amounts to between -0.9 and -1.1 percent of foregone consumption. These are sizable welfare losses, essentially driven by the fact that workers in this segment of the labor market would face a longer expected duration of unemployment following a ban on ZHC.

In order to understand better the welfare consequences of a ban on ZHCs, we conduct a partial equilibrium experiment to isolate the role of the substitution effect. We ask how type-1 workers would be impacted if their ZHCs were replaced by regular contracts, holding constant all the other features of the economy. The results of this experiment are reported in Table 9. Workers' welfare increases by between 0.2 and 0.5 percent in consumption equivalent vari-

[29]The average duration of spells out of regular employment (i.e. the worker might be unemployed or employed in a Z contract) is given by $\frac{\omega_1}{\delta e_R} - \frac{1}{\delta} = \frac{\omega_1}{\lambda(\theta)\frac{v_R}{v}(x(1-\delta)e_{1,Z}+n_1)} - \frac{1}{\delta}$. In the equilibrium without ZHC, the duration of these spells is $\frac{\omega_1}{\delta e_R} - \frac{1}{\delta} = \frac{\omega_1}{\delta \widetilde{e}_R} - \frac{1}{\delta} = \frac{\omega_1}{\lambda(\widetilde{\theta})\widetilde{n}_1} - \frac{1}{\delta}$.

[30]This illustrates well the importance of going beyond the measurement of unemployment duration to get a full picture of workers' trajectories in markets with atypical employment; see Güell et al. [2021] for a recent illustration.

[31]Type-2 workers drop from the labor force permanently. From the model's perspective, this means that their consumption becomes equal to b forever, but this seems an unreasonable assumption to compute their welfare after a ban on ZHCs. For this reason we focus on type-1 workers in our welfare analysis.

Table 9: Welfare effects of a ban on Z contracts: The role of substitution

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
Welfare (in % of CEV)						
At 1st percentile of α	0.00	0.54	0.54	0.53	0.53	0.52
At 25th percentile of α	0.00	0.45	0.45	0.45	0.44	0.44
At 50th percentile of α	0.00	0.37	0.36	0.36	0.36	0.35
At 75th percentile of α	0.00	0.28	0.28	0.27	0.27	0.27
At 99th percentile of α	0.00	0.20	0.19	0.19	0.19	0.18

Notes: The table reports the welfare consequences of a ban on ZHCs for workers who remain in the labor market, under the (partial equilibrium) assumption that the only effect of the ban is to replace ZHCs with regular contracts. Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

ations.[32] In our view, these numbers dovetail well with the negative reactions against ZHCs popularized in the British media and political arena. In equilibrium, however, the substitution effect is counteracted mostly by the job creation effect, so that at least in steady-state comparisons workers suffer from a ban on ZHCs.[33]

Sensitivity analysis. We conduct several robustness analyses of the results presented in Tables 8 and 9. As shown in Appendix D, they are robust to various changes to the baseline parameter values, once the other parameters are recalibrated following the procedure presented in Section 5. We also show that similar effects of a ban on ZHCs would arise after the reform of the Universal Credit's taper rate τ from 63 to 55 percent. In order to measure these effects, we assume that this reform would, in the first place, draw more type-2 workers (i.e., the less attached ones) into the labor force, given that its goal is to provide stronger work incentives to individuals at the margin of non-participation. Last, in Appendix D, we report larger welfare losses if the replacement ratio of UI benefits is lower than under the baseline experiments, which is likely relevant for the evaluation of the consequences of a ban on ZHCs for younger workers.

Alternative view on ZHCs. We examine another regime of the model that allows workers employed in ZHCs decline any workload beyond their available working time. Foremost, we view this scenario as a "stress test" of the model, in that it enables us to check the robustness of the calibration to the hypothesis that ZHCs are much less favorable to firms. Concretely, workers can turn down working hours that become a source of disutility, making expected

[32] Notice that the mix of type-c and type-s firms has a negligible impact on these figures, due to the fact that the utility flows derived from each contract type are similar to each other. Again, this lines up well with the assumption that workers under ZHCs do not attempt to learn the specific type of their own employer.

[33] Notice that in principle the labor force participation effect also matters for the welfare of workers who remain in the market after the policy reform. By dropping from the labor force, type-2 workers reduce congestion externalities on the workers' side of the market. But in quantitative terms this effect is dwarfed by the job creation effect of ZHCs.

Table 10: Equilibrium and welfare effects of a ban on Z contracts: Alternative view

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(a) Equilibrium (re)allocation						
Employment rate (in %)	90.8	86.4	86.2	86.0	85.8	85.6
		-4.41	*-4.60*	*-4.81*	*-5.02*	*-5.25*
Unemployment rate (in %)	9.2	10.8	11.0	11.2	11.5	11.7
		1.63	*1.83*	*2.04*	*2.27*	*2.50*
Duration of R vacancies (in weeks)	10.5	8.2	8.1	8.0	7.9	7.8
		-2.29	*-2.38*	*-2.48*	*-2.57*	*-2.67*
Net output (1 = baseline)	1.00	0.98	0.98	0.98	0.97	0.97
(b) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-1.31	-1.41	-1.53	-1.64	-1.75
At 25th percentile of α	0.00	-1.38	-1.49	-1.59	-1.70	-1.81
At 50th percentile of α	0.00	-1.45	-1.55	-1.65	-1.76	-1.86
At 75th percentile of α	0.00	-1.53	-1.62	-1.72	-1.82	-1.92
At 99th percentile of α	0.00	-1.59	-1.69	-1.78	-1.88	-1.97

Notes: Panel (a) of the table reports the equilibrium allocation effects of a ban on ZHCs. Panel (b) of the table reports the welfare consequences for workers who remain in the labor market after a ban on ZHCs. The first column called 'Baseline' describes the equilibrium with ZHCs. The other columns describe the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

utility and profit functions for Z contracts become:

$$u_Z^i = \frac{e_c}{e_Z} \int u^i \left(\min \left(\widetilde{h}, a \right), a \right) dG_c \left(\widetilde{h} \right) + \frac{e_s}{e_Z} \int u^i \left(\min \left(\widetilde{h}, a \right), a \right) dG_s \left(\widetilde{h} \right)$$
$$\text{and } \pi_Z^j = \int \pi \left(\min \left(\widetilde{h}, a \right), \widetilde{h} \right) dG_j \left(\widetilde{h} \right). \tag{16}$$

Figure C2 in the appendix shows that our choices of μ_c, μ_s, μ_r and σ_c, σ_s, σ_r remain consistent with the taxonomy of firms' types, despite important changes in how firms ranks Z and R contracts in different regions of the parameter space. On the other hand, in this scenario, the fact that some firms substitute ZHCs for regular contracts is actually beneficial to the other side of the market, since regular contracts entail working longer hours, which is valued negatively by some workers.

The results reported in Table 10 show, firstly, that the impact of a ZHC ban on the unemployment rate is dampened. We find increases in the unemployment rate between 1.6 and 2.5 p.p., and decreases in the employment rate between 4.4 and 5.3 p.p. Output is reduced by 2 to 3 percent. Comparing these results to Table 8, we observe that the mix of type-c and type-s employers who offer ZHCs matters slightly more for the equilibrium effects of a ban on ZHCs. Thus our baseline results likely represent an upper bound on the effects of ZHC on equilibrium allocations, given that in some circumstances ZHC workers might be given the option to decline additional working hours. Second, the negative welfare effects are larger, reflecting the assumption that all workers in ZHCs would be able to decline any workload when they see fit. In this scenario, the reduction in workers' welfare ranges from between 1.3 and 1.9 percent in consumption equivalent variations. In this respect, the baseline also represent an upper bound on the losses suffered by workers when firms substitute ZHCs with regular contracts.

6.3 Policy discussion

Turning back to the ongoing media and political debate about the pros and cons of ZHCs, we summarize here how our approach sheds light on the arguments often put forward. The debate focuses on the loss of workers' welfare due to income volatility and tends to assume two extreme opposite views, namely, that either (i) all those workers under Z contracts would be rehired under R contracts when banned, or (ii) all Z jobs would be destroyed under such a ban. Our simulations show that a fraction of ZHCs would survive as R contracts and that, for these matches, there would be a transfer of welfare from firms to workers. There is also a fraction of jobs and workers that only become operative when ZHCs are available. Consequently, their availability allows some additional job creation and higher labor force participation that would otherwise disappear. Our calibration results suggest that the substitution effect from R to Z contracts is of secondary importance relative to the latter effects. We find, however, that in absolute terms substitution effects are the source of important welfare losses.

In light of these findings, we see at least two policy recommendations that should be brought into debate regarding the availability of ZHCs:

(P1) ZHCs could be restricted to job matches where workers opt for Z when offered a choice of contract;

(P2) Access to ZHCs could be prioritized for workers employed in small firms rather than in large firms.

The reasoning behind (P2) is that the volatility of $\tilde{h}$ in a given job match is likely to be lower in large firms where orders can be spread over many employees. As a result, it is likely that substitution is a bigger driver of the use of ZHCs in large firms, whereas in small firms the job creation channel of ZHCs might be relatively more important.

Another direction for policy change would be to regulate the type of flexibility attached to ZHCs. In the above, for our baseline experiment we have assumed that the firm chooses the workload in terms of hours and the worker chooses the timing of work, and found important differences when we made different assumptions along this dimension of the model. Anecdotal evidence suggests that any combination of choices of hours and timing made by the worker or the firm exist within ZHCs. Given the costs incurred on both sides when the match is producing fewer/more hours than desired, a promising avenue for improving workers' welfare under these contracts would be to:

(P3) Recognize that the sharing of hours flexibility between workers and firms is often part of the incompleteness of employment contracts;

(P4) Take steps to regulate the sharing of hours flexibility between workers and firms.

In addition, an assumption in our model is that the fixed costs of employment for the firms and worker's rights are equal across the two contract types. This is an approximation, since there are some relevant differences in workers' rights as detailed in Section 2. In the media debate, ZHCs are often amalgamated with self-employment, as in the current news items on Uber drivers.[34] Despite being beyond the scope of this paper to draw general conclusions about the optimal regulation of the various forms of employment gathered under the label of 'gig economy', our analysis sheds light on some relevant issues in this respect. In particular, our findings point to the need to clarify the type of flexibility and rights attached to each employment relationship and, on the other, that it is key to foster gains from trade in some segments of the labor force where flexibility plays a big role both for firms and workers, without compromising the welfare of workers with low bargaining power.

Finally, it follows from the comparative statics analysis of our model that policy changes affecting labor market institutions, such as a rise of the statutory minimum wage, affect the attractiveness of ZHCs. The quantitative inference that can be drawn from the combination of model and data does not allow us to assess the magnitude of these effects. However, we find it important to highlight that these interactions should be accounted for when it comes to changing the regulation of ZHCs.

7 Conclusions

In this paper we provide a theoretical setup which helps discuss the effects of flexible contracts on labor market outcomes. In particular, we focus on zero-hours contracts (ZHCs) which have become increasingly popular in the U.K. after the Great Recession. Under these contracts, neither employers nor workers commit themselves to offer/accept any given number of working hours. Both may prefer these contracts due to their flexibility; however, workers who are more attached to the labor force, i.e. most workers, would prefer the income stream provided by regular contracts. This gives rise to higher labor turnover in ZHCs which, in the presence of vacancy-opening costs, may lead firms in some instances to replace them by regular contracts.

Our model, which is calibrated to the low-pay occupation segment of the U.K. labor market using LFS data, does a good job in replicating its main stylized facts. We identify worker types that value ZHCs differently because they have different marginal costs of working short vs. long hours schedules, possibly due to household care responsibilities. Similarly, firms do not value different contracts the same as they differ in the mean and variance of their demand of working hours, reflecting differences in the volatility of idiosyncratic business conditions. Accordingly, employers offering ZHCs are better able to satisfy their demand of working hours. Yet, in exchange, they face a higher separation rate because their employees are either of the 'labor-market attached' type (and are searching on the job for a regular contract) or instead belong to the 'less attached' type. This higher worker turnover rate under ZHCs than under regular contracts, and the mix of different worker types among job seekers, become the key determinants of the trade-off that firms face when choosing a contract type. We find that

[34] See this editorial in The Guardian.

simulations of a rise in the minimum wage increases firms' propensity to post ZHCs, which is supported by the empirical findings of Datta et al. [2019]. Foremost, our simulation of the impact of a ban of ZHCs suggests that a share of existing ZHCs would be replaced by regular contracts, while at the same time reducing labor force participation as well as job creation.

Among the issues discussed in the ongoing debate on ZHCs, there is some controversy about the nature of the relationship between some firms and their self-employed contractors, such as Uber drivers. Our findings suggest that most workers employed under ZHCs would prefer having a regular contract. At the same time, as estimated for the U.S. by Frazier [2018], there is not only a substantial willingness to pay for flexible working schedules in some segments of the (potential) labor force, but also productive opportunities in sectors facing highly volatile demand which may not be viable without the ability to adjust working hours at no cost. In light of these considerations, there is scope to improve the regulation of the gig economy. In particular, regulations should clarify the extent and sharing of flexibility of all employment relationships, and target the use of flexible contracts to segments of economic activity and the workforce where the existence of these casual contracts conditions the viability of firms' entry and participation of workers. Identifying such segments, which requires the availability of richer data on firms' profitability and workers' time use and preferences, is left for future research.

References

Zero-hours contracts: Myth and reality. *Chartered Institute of Personnel and Development*, November 2013. https://www.cipd.co.uk/Images/zero-hours-contracts_2013-myth-reality_tcm18-10710.pdf.

The Guardian view on the Uber drivers ruling: A challenge to government. *The Guardian*, February 2021. https://www.theguardian.com/commentisfree/2021/feb/19/the-guardian-view-on-the-uber-drivers-ruling-a-challenge-to-government (accessed July 30, 2021).

Katharine Abraham, John C Haltiwanger, Kristin Sandusky, and James R Spletzer. Measuring the gig economy: Current knowledge and open issues. In Carol A Corrado, Jonathan Haskel, Javier Miranda, and Daniel E Sichel, editors, *Measuring and Accounting for Innovation in the 21st Century*. National Bureau of Economic Research, 2021.

Katharine G Abraham, John Haltiwanger, Kristin Sandusky, and James Spletzer. The rise of the gig economy: Fact or fiction? *AEA Papers and Proceedings*, 109:357–61, 2019.

Karla Adam. Britain‘s Supreme Court rules Uber drivers are ‘workers' entitled to minimum wage and paid vacation. *The Washington Post*, February 2021. https://www.washingtonpost.com/world/europe/britain-uber-supreme-court-workers/2021/02/19/5f9a0b24-729f-11eb-8651-6d3091eac63f_story.html (accessed September 10, 2021).

Abi Adams, Mark Freedland, and Jeremias Prassl. «Zero-hours contracts» in the United Kingdom: Regulating casual work, or legitimating precarity? *Giornale di Diritto del Lavoro e di Relazioni Industriali*, 2015.

Abi Adams, Judith Freedman, and Jeremias Prassl. Rethinking legal taxonomies for the gig economy. *Oxford Review of Economic Policy*, 34(3):475–494, 2018.

John Ameriks, Joseph Briggs, Andrew Caplin, Minjoon Lee, Matthew D Shapiro, and Christopher Tonetti. Older Americans would work longer if jobs were flexible. *American Economic Journal: Macroeconomics*, 12(1):174–209, 2020.

Marta Angelici and Paola Profeta. Smart-working: Work flexibility without constraints. *CESifo Working Paper 8165*, 2020.

Nicholas Bloom, James Liang, John Roberts, and Zhichun Jenny Ying. Does working from home work? Evidence from a Chinese experiment. *The Quarterly Journal of Economics*, 130 (1):165–218, 2015.

Tito Boeri, Giulia Giupponi, Alan B Krueger, and Stephen Machin. Solo self-employment and alternative work arrangements: A cross-country perspective on the changing composition of jobs. *Journal of Economic Perspectives*, 34(1):170–95, 2020.

Daniel Borowczyk-Martins, Grégory Jolivet, and Fabien Postel-Vinay. Accounting for endogeneity in matching function estimation. *Review of Economic Dynamics*, 16(3):440–451, 2013.

Brett Collins, Andrew Garin, Emilie Jackson, Dmitri Koustas, and Mark Payne. Is gig work replacing traditional employment? Evidence from two decades of tax returns. *Unpublished paper, IRS SOI Joint Statistical Research Program*, March 2019.

Nikhil Datta, Giulia Giupponi, and Stephen Machin. Zero-hours contracts and labour market policy. *Economic Policy*, 34(99):369–427, 2019.

Linda Dickens. Exploring the atypical: Zero hours contracts. *Industrial Law Journal*, 26(3): 262–264, 1997.

Arindrajit Dube, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. Monopsony in online labor markets. *American Economic Review: Insights*, 2(1):33–46, 2020.

Michael WL Elsby and Ryan Michaels. Marginal jobs, heterogeneous firms, and unemployment flows. *American Economic Journal: Macroeconomics*, 5(1):1–48, 2013.

Nick Frazier. An equilibrium model of wage and hours determination: Labor market regulation in the retail sector. *Unpublished working paper, Rice University*, 2018. https://npfrazier.com/nfrazier-geneqm.pdf.

Maia Güell, Cristina Lafuente, Manuel Sánchez, and Hélène Turon. So different yet so alike: Micro and macro labour market outcomes in Germany and Spain. *SERIEs, Journal of the Spanish Economic Association*, pages 1–23, 2021.

Georgina Hutton and Matthew Ward. Business Statistics. *House of Commons Research Briefing*, December 2021. https://researchbriefings.files.parliament.uk/documents/SN06152/SN06152.pdf.

Lawrence F Katz and Alan B Krueger. The rise and nature of alternative work arrangements in the United States, 1995-2015. *Industrial & Labor Relations Review*, 72(2):382–416, 2019a.

Lawrence F Katz and Alan B Krueger. Understanding trends in alternative work arrangements in the United States. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 5 (5):132–146, 2019b.

Dmitri Koustas. Consumption insurance and multiple jobs: Evidence from rideshare drivers. *Unpublished working paper, University of Chicago*, 2018. http://dmitrikoustas.com/DKoustas-RideSmoothing-WP.pdf.

Moritz Kuhn, Iourii Manovskii, and Xincheng Qiu. The geography of job creation and job destruction. *NBER Working Paper 29399*, 2021.

Alexandre Mas and Amanda Pallais. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59, 2017.

Alexandre Mas and Amanda Pallais. Alternative work arrangements. *Annual Review of Economics*, 12:631–658, 2020.

Lawrence Mishel and Celine McNicholas. What we learned from the UK case rendering Uber drivers employees. *Economic Policy Institute*, Working Economics Blog, March 2021. https://www.epi.org/blog/what-we-learned-from-the-uk-case-rendering-uber-drivers-employees/ (accessed September 10, 2021).

Christina Patterson, Ayşegül Şahin, Giorgio Topa, and Giovanni L Violante. Working hard in the wrong place: A mismatch-based explanation to the uk productivity puzzle. *European Economic Review*, 84:42–56, 2016.

Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988, 2005.

Graham Ruddick. McDonald's offers fixed contracts to 115,000 UK zero-hours workers. *The Guardian*, April 2017. https://www.theguardian.com/business/2017/apr/25/mcdonalds-contracts-uk-zero-hours-workers (accessed July 30, 2021).

Rachel Scarfe. Flexibility or certainty? The aggregate effects of casual jobs on labour markets. *Edinburgh School of Economics Discussion Paper 294*, January 2019.

Alain Tolhurst. Record figures reveal almost one million workers on zero-hour contracts. *Politics Home*, February 2020. https://www.politicshome.com/news/article/record-figures-reveal-almost-one-million-workers-on-zerohour-contracts (accessed September 10, 2021).

Appendix

The appendix contains four sections. Section A presents the stock-flow equations that define the law of motion of the model. Section B presents the estimation of the vacancy-elasticity of the matching function based on U.K. data for the low-pay labor market. Section C collects additional figures and tables. Section D includes several robustness checks of the main results of the paper.

A Stock-flow equations

$e_{i,j}$ denotes the measure of job matches between workers of type $i = 1, \ldots, 4$ and firms of type $j = c, s, r$. n_i is the measure of non-employed workers of type i, and v_j the measure of vacancies of firms of type j. $v = \sum_j v_j$ is the aggregate measure of vacancies. Employment in firm types $j = c, s$ (i.e, the firms offering Z contracts) evolves according to:

$$e'_{1,j} = \left(1 - x\lambda(\theta)\frac{v_r}{v}\right)(1-\delta)\, e_{1,j} + \lambda(\theta)\frac{v_j}{v}n_1 \tag{A.1}$$

$$e'_{2,j} = (1-\delta)\, e_{2,j} + \lambda(\theta)\frac{v_j}{v}n_2 \tag{A.2}$$

$$e'_{3,j} = 0 \tag{A.3}$$

$$e'_{4,j} = (1-\delta)\, e_{4,j} + x\lambda(\theta)\frac{v_j}{v}(1-\delta)\, e_{4,r} + \lambda(\theta)\frac{v_j}{v}n_4, \tag{A.4}$$

while the corresponding law of motion for employment in type-r firms is:

$$e'_{1,r} = (1-\delta)\, e_{1,r} + x\lambda(\theta)\frac{v_r}{v}(1-\delta)(e_{1,c} + e_{1,s}) + \lambda(\theta)\frac{v_r}{v}n_1 \tag{A.5}$$

$$e'_{2,r} = 0 \tag{A.6}$$

$$e'_{3,r} = (1-\delta)\, e_{3,r} + \lambda(\theta)\frac{v_r}{v}n_3 \tag{A.7}$$

$$e'_{4,r} = \left(1 - x\lambda(\theta)\frac{v_c + v_s}{v}\right)(1-\delta)\, e_{4,r} + \lambda(\theta)\frac{v_r}{v}n_4. \tag{A.8}$$

The measures of non-employed workers evolve according to:

$$n'_1 = \delta(e_{1,r} + e_{1,c} + e_{1,s}) + (1 - \lambda(\theta))\, n_1 \tag{A.9}$$

$$n'_2 = \delta(e_{2,c} + e_{2,s}) + \left(1 - \lambda(\theta)\frac{v_c + v_s}{v}\right)n_2 \tag{A.10}$$

$$n'_3 = \delta e_{3,r} + \left(1 - \lambda(\theta)\frac{v_r}{v}\right)n_3 \tag{A.11}$$

$$n'_4 = \delta(e_{4,r} + e_{4,c} + e_{4,s}) + (1 - \lambda(\theta))\, n_4 \tag{A.12}$$

Let $n = \sum_i n_i$ denote the aggregate measure of non-employed workers. The stocks of vacant positions of firm types $j = c, s$ evolve according to:

$$v'_j = x\lambda(\theta)\frac{v_r}{v}(1-\delta)\, e_{1,j} + \left(1 - \frac{\lambda(\theta)}{\theta}\frac{n_1 + n_2 + n_4 + xe_{4,r}}{n + x(e_{1,c} + e_{1,s} + e_{4,r})}\right)v_j \\ + \delta(1-n)\gamma_j, \tag{A.13}$$

while the corresponding law of motion for r firms is:

$$v_r' = x\lambda(\theta)\frac{v_c+v_s}{v}(1-\delta)e_{4,r} + \left(1-\frac{\lambda(\theta)}{\theta}\frac{n_1+n_3+n_4+x(e_{1,c}+e_{1,s})}{n+x(e_{1,c}+e_{1,s}+e_{4,r})}\right)v_r + \delta(1-n)\gamma_j. \tag{A.14}$$

In Equations (A.13) and (A.14), $\delta(1-n)$ corresponds to the number of firms with a filled position hit by the destruction shock, which exit the market and are replaced by firms that draw their type from the distribution γ_j, with $j=c,s,r$.

Finally, we have:

$$e_{i,c}+e_{i,s}+e_{i,r}+n_i=\omega_i, \tag{A.15}$$

where ω_i is the share of type-i workers, and $\sum_i \omega_i = 1$.

B Vacancy elasticity of the matching function

We use U.K. hiring and job vacancy data from Patterson et al. [2016] to estimate the elasticity of the matching function with respect to vacancies. These data are available at the levels of 2-digit occupations, which enables us to extract information for the low-pay segment of the labor market. We use newly-formed matches ($M_{o,t}$), unemployment claims ($U_{o,t}$) and job vacancies ($V_{o,t}$) for the following occupations o: 'Administrative', 'Secretarial and related', 'Caring personal service', 'Leisure and other personal service', 'Process, plant and machine', 'Elementary trades, plant and storage related', and 'Elementary administration and service', to run the following linear regression:

$$\log\left(\frac{M_{o,t}}{U_{o,t}}\right) = \alpha_o + \varpi' g(t) + \psi \log\left(\frac{V_{o,t}}{U_{o,t}}\right) + \varepsilon_{o,t}. \tag{B.1}$$

α_o is an occupation fixed effect, $g(t)$ is a polynomial of time that allows for a flexible time trend, $\varepsilon_{o,t}$ is the regression residual, and ψ is the coefficient of interest. The data is monthly, seasonally adjusted, and runs from April 2004 through June 2012. Estimation results are reported in Table B1.

Table B1: Vacancy elasticity of the matching function

	Log- job finding ($\log({}^{M_{o,t}}/_{U_{o,t}})$)			
	(1)	**(2)**	**(3)**	**(4)**
Log- market tightness ($\log({}^{V_{o,t}}/_{U_{o,t}})$)	0.643***	0.701***	0.586***	0.703***
	(0.027)	(0.041)	(0.025)	(0.034)
R-squared	0.859	0.896	0.802	0.871
Time trend ($g(t)$)		✔		✔
Occupation fixed effect (α_o)			✔	✔

Notes: Author's calculations based on data from Patterson et al. [2016]. Each column reports coefficients from a linear estimation of Equation (B.1). Column (1) presents the univariate regression, column (2) adds a polynomial time trend, column (3) adds occupation fixed effects, and column (4) controls simultaneously for the time trend and occupation fixed effects. The time trend $g(t)$ is a 5-th order polynomial function. Standard errors in parentheses are clustered at the occupation level.

The coefficient on the log of market tightness (as measured by the ratio between job vacancies and unemployment claimants) is precisely estimated and statistically higher than the conventional value of 0.50. Depending on the specification used, it ranges from 0.59 to 0.70. In the model's calibration, we use the mid-point value of 0.65.

C Additional tables and figures

Table C1 reports the duration of unemployment spells computed from the LFS for workers in the low-pay segment of the labor market. Our model does a good job at matching the average duration of unemployment. However, as shown in the bottom panel of Figure 3, it does so by understating the shares of short (under 6 months) and very long (5 years and over) unemployment spells.

Table C1: Distribution of unemployment duration

Unemployment duration:	
Less than 3 months	35.2
3 to 6 months	16.0
6 to 12 months	14.0
1 to 2 years	14.2
2 to 3 years	4.9
3 to 4 years	2.4
4 to 5 years	1.6
More than 5 years	11.7

Notes: Authors' calculations based on data from the Labour Force Survey. All table entries are expressed in percent.

Table C2 compares the model-generated quarterly transition rates across labor market states to their empirical counterparts (displayed in Table 4, which is repeated here for reference). The model overstates the transition rate from Z to R jobs, which is likely due to a relatively high on-the-job search intensity (x is found to be 0.35, which is high compared to values of this parameter commonly used in the literature). It also overestimates the transition rate from non-employment to ZHCs. On the other hand, it matches perfectly the outflow rates of regular jobs. Notice that the model features a non-zero quarterly transition rate from R to Z jobs. This result is a fabrication of time aggregation. At the model's bi-weekly frequency, there is no such transitions, since the equilibrium does not feature any type-4 workers. But some type-1 workers who are employed in a R job at some point are employed in a Z job six months later after transitioning through unemployment.

Table C2: Model fit: Transition rates between employment status and labor contracts

	(a) Model					**(b) Data**			
	To:	$\boldsymbol{N}$	$\boldsymbol{Z}$	$\boldsymbol{R}$		**To:**	$\boldsymbol{N}$	$\boldsymbol{Z}$	$\boldsymbol{R}$
From:	$\boldsymbol{N}$	56.0	8.8	35.2	**From:**	$\boldsymbol{N}$	62.2	4.2	33.5
	$\boldsymbol{Z}$	4.8	83.6	11.6		$\boldsymbol{Z}$	6.2	87.3	6.5
	$\boldsymbol{R}$	4.4	0.3	95.3		$\boldsymbol{R}$	4.4	0.5	95.2

Notes: Panel (a): Model's predicted moments. Panel (b): Author's calculations based on data from the Labour Force Survey. N: Not employed, Z: Employed in a zero-hours contract, R: Employed not in a zero-hours contract. All table entries are expressed in percent.

Figure C1 plots the distributions $H_j\left(.\right)$'s from which each firm type j draws its demand of working hours, $\widetilde{h}$. Types c and s demand lower and more volatile hours relative to type-r firms.

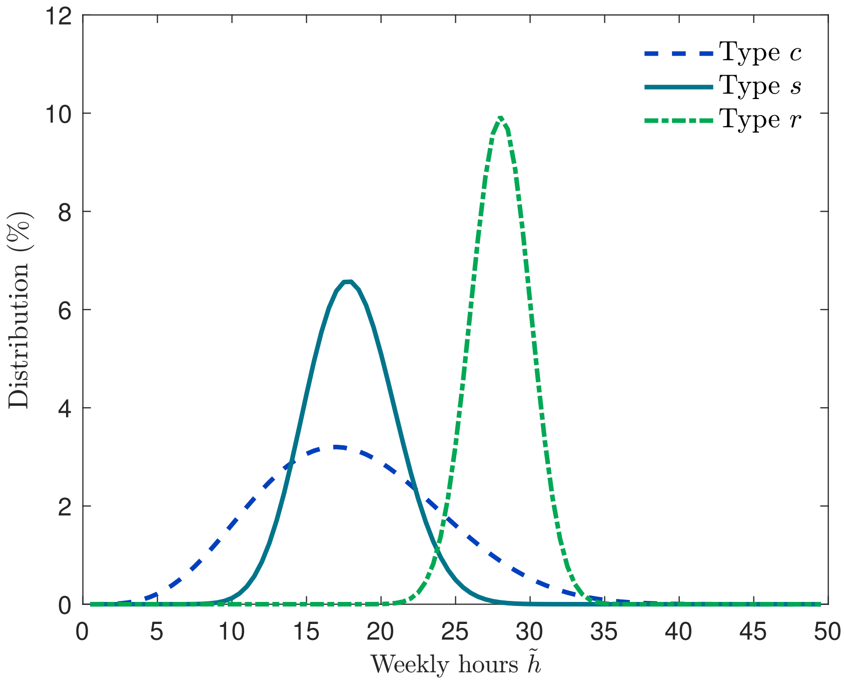

Figure C1: Firms' distribution of the demand of weekly working hours $\tilde{h}$

Notes: The figure shows the distribution of the demand of weekly working hours of firms of type $j = c, s, r$.

Figure C2, which is the counterpart of Figure 4 in the text, presents regions of the parameter space under the assumption that workers employed in ZHCs can decline hours worked above a. As can be seen, firms with a relatively high mean of hours and high standard deviation can no longer profitably substitute ZHCs for regular contracts. As a result, the region of type-r firms expands at the detriment of the region of type-s firms. Our baseline choices of μ_c, μ_s, μ_r and σ_c, σ_s, σ_r remain consistent with the taxonomy of firms' types in (2).

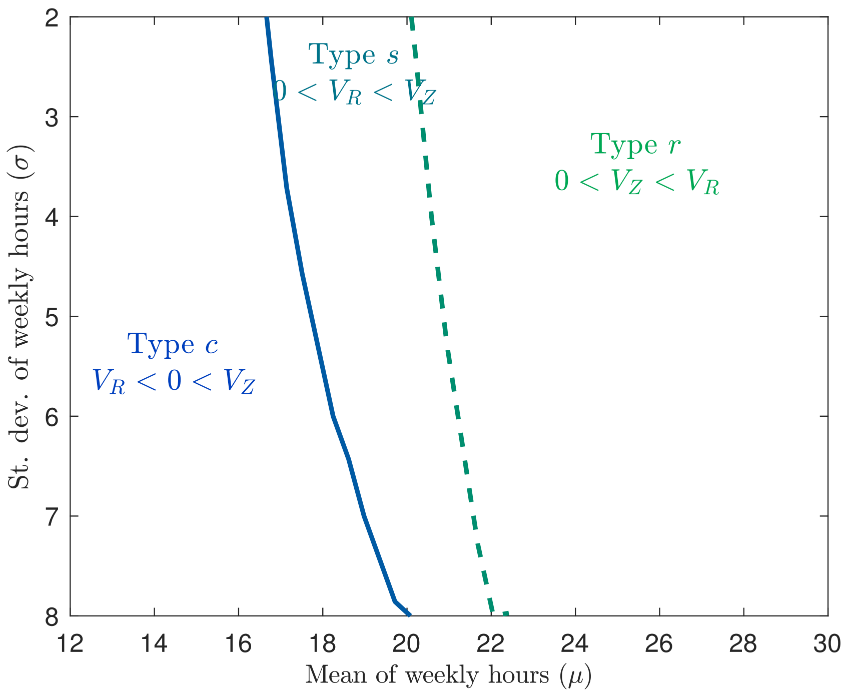

Figure C2: Firms' types across (some) regions of the parameter space

Notes: The figure shows firms' rankings of the asset values of advertising Z jobs (V_Z) and R jobs (V_R), as a function of firms' mean μ_j and standard deviation σ_j of the demand of weekly hours worked.

D Robustness checks

D.1 The job-filling rate. In the first step of the calibration, we target the job-finding rate observed in the LFS data and the mid-value of the U.K. job-filling rate from Kuhn et al. [2021] to pin down the values of matching efficiency M and market tightness θ. In this section, we investigate the sensitivity of the results to the empirical values bracketing the job-filling rate. While the baseline calibration targets a bi-weekly value of 0.20, we recalibrate the model under the assumption that $\lambda(\theta)/\theta$ is lower by 25 percent. This yields matching efficiency M equal to 0.1071 and a value of market tightness θ of 0.3172 (vs. $M = 0.1278$ and $\theta = 0.2412$ in the baseline calibration). The equilibrium and welfare effects of a ban on ZHCs under this calibration are reported in the upper panels of Table D1. The main difference with the baseline results is that job creation is less responsive to the policy reform, which in turn leads to smaller welfare effects. Conversely, a 25 percent higher calibration target for $\lambda(\theta)/\theta$, which yields recalibrated values of $M = 0.1478$ and $\theta = 0.1953$, amplifies the response of job creation and triggers larger welfare losses following a ban on ZHCs. These results are shown in the lower panels of Table D1. The baseline results seem to compare favorably to these two scenarios. On the one hand, the targeted job-filling rate of 0.15 yields expected durations of vacancies of almost one semester, which seems rather long, and a decrease of sectoral output in response to a ban on ZHCs by 'only' 1 percent. On the other hand, the targeted job-filling rate of 0.25 generates a reduction of sectoral output by 5-6 percent following a ban on ZHCs, which seems large given that ZHCs account for 6.5 percent of baseline employment and would be partially replaced by regular employment contracts after the policy reform.

D.2 Changes in the taper rate of welfare benefits. The taper rate of the U.K.'s Universal Credit (UC) will soon be changed from 63 percent to 55 percent, which will give eligible workers an extra 8 pence for every extra £1 that they earn. We investigate the consequences of a ban on ZHCs in a steady-state equilibrium that is characterized by $\tau = 0.55$. We assume that in such a steady state there would be a larger share of type-2 workers who participate in the labor market, given the stronger work incentives provided by the UC taper rate. Specifically, we recalibrate the model under the constraint that ω_2 be 50 percent higher than in the baseline equilibrium ($\omega_2 = 0.0467$ vs. $\omega_2 = 0.0311$ in our main analysis). This leads to a higher employment rate in steady state equilibrium, which in turn is important to appreciate the magnitude of the impact of a ban on ZHCs. As shown in Table D2, the ban causes a smaller change of the unemployment rate: it increases by between 1 and 1.5 p.p., which is twice as less as the baseline experiments. This partly explains why the welfare losses in Table D2 are lower compared to those in the main analysis. The other reason is that the more generous taper rate in this scenario mitigates the welfare losses for those workers who remain in the labor market after the policy reform.

D.3 Other determinants of welfare effects. The welfare effects of a ban on ZHCs depend workers' preferences, the endowment in terms of working time a, the taper rate and generosity of UI benefits b. In Table D3, we assess the sensitivity of the welfare effects to the assumption on working time availability, a. In the upper panel, individuals are available to work three days per week, where each day consists of 5 hours of daily work plus half an hour of travel-to-work. In the lower panel, they are endowed with five working days per week. Recall from Figure 5 that a matters for the bounds on the marginal utility cost of working excess hours, α_i, that separate different worker types. Thus, the *values* of α_i corresponding to the different percentiles are much different between the upper and lower panels of Table D3. As the table shows, however, the welfare effects computed at those values are very similar to the welfare effects reported in the baseline experiments.

Table D1: Robustness check: Equilibrium and welfare effects of a ban on Z contracts
Changing the targeted job-filling rate $\lambda(\theta)/\theta$

I. $\lambda(\theta)/\theta = 0.15$

	Baseline	**Share of type-c among Z jobs**				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(Ia) Equilibrium (re)allocation						
Employment rate (in %)	90.8	87.5	87.5	87.5	87.4	87.4
		-3.27	*-3.30*	*-3.34*	*-3.37*	*-3.40*
Unemployment rate (in %)	9.2	9.7	9.7	9.7	9.8	9.8
		0.47	*0.50*	*0.54*	*0.57*	*0.60*
Duration of R vacancies (in weeks)	8.5	11.5	11.5	11.5	11.5	11.5
		-2.25	*-2.27*	*-2.29*	*-2.32*	*-2.34*
Net output (1 = baseline)	1.00	0.99	0.99	0.99	0.99	0.99
(Ib) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-0.19	-0.19	-0.20	-0.21	-0.23
At 25th percentile of α	0.00	-0.23	-0.24	-0.26	-0.27	-0.28
At 50th percentile of α	0.00	-0.29	-0.30	-0.31	-0.32	-0.33
At 75th percentile of α	0.00	-0.35	-0.36	-0.37	-0.37	-0.38
At 99th percentile of α	0.00	-0.41	-0.42	-0.42	-0.43	-0.43

II. $\lambda(\theta)/\theta = 0.25$

	Baseline	**Share of type-c among Z jobs**				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(IIa) Equilibrium (re)allocation						
Employment rate (in %)	90.8	84.2	83.8	83.5	83.2	82.8
		-6.66	*-6.97*	*-7.30*	*-7.64*	*-8.01*
Unemployment rate (in %)	9.2	13.1	13.5	13.8	14.2	14.5
		3.96	*4.28*	*4.61*	*4.97*	*5.35*
Duration of R vacancies (in weeks)	8.5	5.9	5.8	5.7	5.6	5.5
		-2.60	*-2.68*	*-2.77*	*-2.86*	*-2.96*
Net output (1 = baseline)	1.00	0.95	0.95	0.94	0.94	0.94
(IIb) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-1.87	-2.01	-2.15	-2.30	-2.45
At 25th percentile of α	0.00	-1.87	-1.99	-2.12	-2.26	-2.39
At 50th percentile of α	0.00	-1.86	-1.97	-2.09	-2.22	-2.34
At 75th percentile of α	0.00	-1.85	-1.96	-2.06	-2.18	-2.29
At 99th percentile of α	0.00	-1.84	-1.94	-2.04	-2.14	-2.23

Notes: Sections I, II of the table correspond to different calibration targets for the job-filling rate $\lambda(\theta)/\theta$. Panels (Ia), (IIa) of the table report the equilibrium allocation effects of a ban on ZHCs. Panels (Ib), (IIb) of the table report the welfare consequences for workers who remain in the labor market after a ban on ZHCs. The first column called 'Baseline' describes the equilibrium with ZHCs. The other columns describe the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

Table D2: Equilibrium and welfare effects of a ban on Z contracts
Changing the taper rate τ and participation of less attached workers

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(a) Equilibrium (re)allocation						
Employment rate (in %)	90.6	85.4	85.3	85.2	85.0	84.9
		-5.91	*-6.02*	*-6.13*	*-6.25*	*-6.38*
Unemployment rate (in %)	9.4	10.4	10.5	10.7	10.8	10.9
		1.01	*1.13*	*1.25*	*1.38*	*1.51*
Duration of R vacancies (in weeks)	10.8	8.4	8.4	8.3	8.2	8.2
		-2.41	*-2.47*	*-2.53*	*-2.59*	*-2.65*
Net output (1 = baseline)	1.00	0.97	0.97	0.97	0.97	0.97
(b) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-0.57	-0.63	-0.69	-0.75	-0.81
At 25th percentile of α	0.00	-0.63	-0.68	-0.74	-0.80	-0.85
At 50th percentile of α	0.00	-0.69	-0.74	-0.79	-0.84	-0.89
At 75th percentile of α	0.00	-0.75	-0.80	-0.84	-0.89	-0.93
At 99th percentile of α	0.00	-0.81	-0.85	-0.89	-0.93	-0.97

Notes: Panel (a) of the table reports the equilibrium allocation effects of a ban on ZHCs. Panel (b) of the table reports the welfare consequences for workers who remain in the labor market after a ban on ZHCs. The first column called 'Baseline' describes the equilibrium with ZHCs. The other columns describe the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

Table D3: Robustness checks: Welfare effects of a ban on Z contracts
Changing available hours a

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
I. Welfare (in % of CEV) under $a = 16.5$ **hours**						
At 1st percentile of α	0.00	-0.91	-0.97	-1.03	-1.09	-1.15
At 25th percentile of α	0.00	-0.93	-0.98	-1.03	-1.09	-1.14
At 50th percentile of α	0.00	-0.94	-0.99	-1.04	-1.09	-1.13
At 75th percentile of α	0.00	-0.96	-1.01	-1.05	-1.09	-1.13
At 99th percentile of α	0.00	-0.98	-1.02	-1.05	-1.09	-1.12
II. Welfare (in % of CEV) under $a = 27.5$ **hours**						
At 1st percentile of α	0.00	-0.93	-0.99	-1.04	-1.08	-1.11
At 25th percentile of α	0.00	-0.96	-1.01	-1.05	-1.08	-1.09
At 50th percentile of α	0.00	-1.00	-1.03	-1.06	-1.08	-1.07
At 75th percentile of α	0.00	-1.03	-1.05	-1.07	-1.08	-1.06
At 99th percentile of α	0.00	-1.06	-1.08	-1.08	-1.07	-1.04

Notes: Sections I, II of the table correspond to different choices of values for available hours per week a. Each section of the table reports the welfare consequences of a ban on ZHCs for workers who remain in the labor market. The first column called 'Baseline' corresponds to the equilibrium with ZHCs. The other columns correspond to the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

Table D4: Robustness checks: Welfare effects of a ban on Z contracts
Changing unemployment benefits b

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
At 1st percentile of α	0.00	-1.07	-1.14	-1.22	-1.29	-1.37
At 25th percentile of α	0.00	-1.11	-1.17	-1.24	-1.31	-1.38
At 50th percentile of α	0.00	-1.14	-1.20	-1.26	-1.32	-1.38
At 75th percentile of α	0.00	-1.17	-1.23	-1.28	-1.34	-1.39
At 99th percentile of α	0.00	-1.21	-1.25	-1.30	-1.35	-1.40

Notes: The table reports the welfare consequences for workers of a ban on ZHCs for workers who remain in the labor market. The first column called 'Baseline' corresponds to the equilibrium with ZHCs. The other columns correspond to the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

Table D5: Robustness checks: Equilibrium and welfare effects of a ban on Z contracts
Changing (μ_c, μ_s, μ_r) and $(\sigma_c, \sigma_s, \sigma_r)$

I. $(\mu_c, \mu_s, \mu_r) = (18, 18, 28)$ **and** $(\sigma_c, \sigma_s, \sigma_r) = (8, 1, 2)$

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(Ia) Equilibrium (re)allocation						
Employment rate (in %)	90.8	86.9	86.6	86.2	85.8	85.4
		-3.93	*-4.26*	*-4.61*	*-4.99*	*-5.40*
Unemployment rate (in %)	9.2	10.3	10.7	11.0	11.4	11.8
		1.14	*1.48*	*1.84*	*2.23*	*2.65*
Duration of R vacancies (in weeks)	10.5	8.5	8.3	8.1	8.0	7.8
		-2.06	*-2.22*	*-2.39*	*-2.56*	*-2.73*
Net output (1 = baseline)	1.00	0.98	0.98	0.97	0.97	0.97
(Ib) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-0.51	-0.66	-0.81	-0.97	-1.12
At 25th percentile of α	0.00	-0.55	-0.69	-0.84	-0.98	-1.11
At 50th percentile of α	0.00	-0.60	-0.73	-0.86	-0.99	-1.10
At 75th percentile of α	0.00	-0.65	-0.77	-0.89	-1.00	-1.10
At 99th percentile of α	0.00	-0.70	-0.81	-0.91	-1.01	-1.09

II. $(\mu_c, \mu_s, \mu_r) = (16, 20, 28)$ **and** $(\sigma_c, \sigma_s, \sigma_r) = (4.5, 4.5, 2)$

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(IIa) Equilibrium (re)allocation						
Employment rate (in %)	90.8	87.7	87.4	87.1	86.8	86.4
		-3.14	*-3.42*	*-3.72*	*-4.03*	*-4.38*
Unemployment rate (in %)	9.2	9.5	9.8	10.1	10.4	10.8
		0.33	*0.61*	*0.92*	*1.25*	*1.60*
Duration of R vacancies (in weeks)	10.5	8.9	8.7	8.6	8.4	8.2
		-1.63	*-1.79*	*-1.95*	*-2.11*	*-2.28*
Net output (1 = baseline)	1.00	0.99	0.99	0.99	0.98	0.98
(IIb) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-0.16	-0.25	-0.38	-0.52	-0.66
At 25th percentile of α	0.00	-0.19	-0.30	-0.43	-0.56	-0.69
At 50th percentile of α	0.00	-0.23	-0.36	-0.48	-0.60	-0.72
At 75th percentile of α	0.00	-0.30	-0.41	-0.52	-0.64	-0.75
At 99th percentile of α	0.00	-0.36	-0.46	-0.57	-0.68	-0.79

III. $(\mu_c, \mu_s, \mu_r) = (18, 18, 28)$ **and** $(\sigma_c, \sigma_s, \sigma_r) = (2, 6, 4.5)$

	Baseline	Share of type-c among Z jobs				
		0.00	**0.25**	**0.50**	**0.75**	**1.00**
(IIIa) Equilibrium (re)allocation						
Employment rate (in %)	90.8	86.0	85.8	85.7	85.5	85.4
		-4.83	*-4.98*	*-5.13*	*-5.29*	*-5.46*
Unemployment rate (in %)	9.2	11.3	11.4	11.6	11.7	11.9
		2.07	*2.22*	*2.38*	*2.54*	*2.71*
Duration of R vacancies (in weeks)	10.5	8.0	8.0	7.9	7.8	7.8
		-2.49	*-2.55*	*-2.62*	*-2.68*	*-2.75*
Net output (1 = baseline)	1.00	0.97	0.97	0.97	0.97	0.97
(IIIb) Welfare (in % of CEV)						
At 1st percentile of α	0.00	-0.95	-1.02	-1.08	-1.15	-1.21
At 25th percentile of α	0.00	-0.98	-1.04	-1.10	-1.16	-1.22
At 50th percentile of α	0.00	-1.01	-1.06	-1.12	-1.17	-1.22
At 75th percentile of α	0.00	-1.04	-1.09	-1.13	-1.18	-1.23
At 99th percentile of α	0.00	-1.07	-1.11	-1.15	-1.19	-1.23

Notes: Sections I, II, III of the table correspond to different choices of values for (μ_c, μ_s, μ_r) and $(\sigma_c, \sigma_s, \sigma_r)$. Panels (Ia), (IIa), (IIIa) of the table report the equilibrium allocation effects of a ban on ZHCs. Panels (Ib), (IIb), (IIIb) of the table report the welfare consequences for workers who remain in the labor market after a ban on ZHCs. The first column called 'Baseline' describes the equilibrium with ZHCs. The other columns describe the equilibrium obtained under the ban on ZHCs, under the assumptions that type-c firms accounted for 0, 25, 50, 75, or 100 percent of all firms that offered ZHCs in the baseline equilibrium (the other providers of ZHCs are type-s firms). Welfare effects are computed in consumption equivalent variation (CEV) and expressed in percent.

Next, in Table D4, we study the sensitivity of the welfare effects to the replacement ratio of UI benefits. We calibrate b such that the replacement ratio is 70 percent, which is the replacement ratio for a minimum wage worker who is single and lives without children according to the OECD's net replacement rate calculator (`https://stats.oecd.org/Index.aspx?DataSetCode=NRR`). The welfare losses are higher than in the baseline experiments, which is driven by larger reductions in consumption during unemployment due to lower b.

D.4 Firms' demand of weekly working hours. We conducted extensive robustness analyses to assess the role of the means μ_c, μ_s, μ_r and standard deviations σ_c, σ_s, σ_r that govern firms' demand of working hours, $\tilde{h}$. Table D5 summarizes the results. The first set of changes consist in varying the difference between the standard deviations σ_c and σ_s, while holding the means μ_c and μ_s unchanged (μ_r and σ_r are also unchanged from the baseline calibration). Recall from Figure 4 that by choosing $\mu_c = \mu_s$, we can only select σ_c and σ_s along the vertical line $\mu = 18$ below and above the curve that separates type-c from type-s firms. In the upper panels of Table D5, we increase the value of σ_c and decrease that of σ_s by 2 units relative to the baseline calibration. The results remain broadly similar. Next, we investigate the consequences of $\sigma_c = \sigma_s$, which we set to mid-point value of 4.5 (in the baseline calibration, $\sigma_c = 6$ and $\sigma_s = 3$). $\sigma_c = \sigma_s$ implies that μ_c must be different from, and lower than, μ_s (see Figure 4). We use $\mu_c = 16$ and $\mu_s = 20$, so that the mean value of hours in Z contracts remains equal to 18 hours as in the baseline experiment. Again, μ_r and σ_r are also left unchanged. The impact of a ban on ZHCs on equilibrium allocations is dampened, and as a result the welfare losses are slightly lower compared to the main analysis. Last, we investigate the consequences of changing σ_r, the standard deviation of working hours demanded by type-r firms, while keeping μ_c, μ_s, μ_r as well as σ_c and σ_s unchanged from the baseline calibration. We set $\sigma_r = 4.5$, which is the mid-point value of σ_c and σ_s. The lower panels of Table D5 shows that the effect on the results is negligible. Overall, the main results are robust to changing μ_c, μ_s, μ_r and σ_c, σ_s, σ_r, partly because these changes imply that we must recalibrate ϕ, κ, K, and our calibration targets for those parameters put some discipline on the responsiveness of job creation. Note that changes to μ_c, μ_s, σ_c, σ_s also affect average labor earnings (through actual working hours, since $h = \widetilde{h}$ under a Z contract) and hence unemployment benefits, which matters for the assessment of welfare effects.