Discounts, Rationing, and Unemployment

Alex Clymo*

June 6, 2020

Abstract

How are changes in discount rates transmitted to unemployment, and are they a quantitatively relevant driver of the Great Recession? In this paper I answer these questions in a search and matching model featuring endogenous capital accumulation. Rising discount rates have a direct effect on unemployment, since hiring is an investment activity for firms in the presence of search frictions, and an indirect effect, since a rise in discounts reduces investment in capital, reducing the marginal product of labour and hence incentives to hire. I estimate changes in discount rates during the Great Recession using data on both stock markets and investment. Discounts measured from the stock market increase by more but recover faster, while discounts for capital investment increase by less but more persistently. Combined, the two discount measures can account for 52% of the peak rise in unemployment during the Great Recession. Finally, while capital discounts affect unemployment by changing labour productivity, their importance cannot be inferred from observed movements in labour productivity alone, and capital discounts raised unemployment even while observed labour productivity was rising.

Keywords: discounts, unemployment, investment, great recession

JEL codes: E24, E22, J64, G01

1 Introduction

It is well known that stock prices are more volatile than can be explained by changes in dividends in simple asset pricing models (Fama and French, 1988, Campbell and Shiller, 1988). The majority of movements in stock prices must therefore be explained by changes in how people effectively discount the future: changes in discount rates, or "discounts". In this paper, I investigate to what extent changes in discount rates can move not only the stock market, but the real economy.

In particular, stock prices tend to fall in recessions, when unemployment also rises, as shown in Figure 1. Hall (2017) argues that the same forces which cause people to discount stock returns more heavily in recessions may also cause them to discount the benefits of *hiring* more, leading to

*Department of Economics, University of Essex. Email: a.clymo@essex.ac.uk. I have benefitted from comments on various iterations of this project from Wouter den Haan, Sergio de Ferra, Axel Gottfries, Marcus Hagedorn, Andrea Lanteri, Kevin Sheedy, and seminar participants at the LSE, and EEA-ESEM conference 2017. I owe a particular debt of gratitude to my supervisor, Wouter den Haan, for invaluable support during my PhD, where I began work on an earlier version of this project. I thank the ESRC and Paul Woolley Centre for financial support during my PhD. All errors are my own.

reduced hiring and increased unemployment. Of particular note is the striking coincident increase in unemployment (Figure 1, top right panel) and decline in stock prices (Figure 1, bottom right panel) during the Great Recession of 2008.

Focusing on the Great Recession, however, the behaviour of capital (Figure 1, bottom left panel) is perhaps equally striking. Capital began falling immediately in the Great Recession (relative to its pre-crisis trend) as investment collapsed, and has now fallen by more than 15%. If increased discount rates can cause a decrease in hiring, it is equally plausible that they should cause a fall in investment in physical capital. Indeed, high discount rates may proxy for tightened financial frictions which make investors effectively more impatient, and the financial crisis has motivated a large literature (e.g. Jermann and Quadrini, 2012, Brinca et al., 2016) arguing that financial frictions were an important driver of investment and output during the Great Recession. Moreover, capital remained depressed permanently following the crisis, even after stock markets had recovered, suggesting that the discount rates estimated from these two markets may in fact be different.

In this paper I synthesise research on how the financial crisis affected unemployment and capital investment by considering both through the lens of discounts. I show that an increase in discount rates increases unemployment both *directly*, by increasing the amount by which the benefits of hiring are discounted, and *indirectly* by reducing capital investment and hence the marginal productivity of labour. I estimate discounts from both stock markets and capital investment and show that, combined, these discounts can account for 52% of the peak rise in unemployment during the Great Recession, and that both discount rates are important but their effects peak at different times.

To understand the key ideas of this paper, consider first a simple search and matching model of unemployment, à la Pissarides (2000). Anticipating Michaillat's (2012) terminology, I will informally refer to this as a model of *frictional unemployment*. To hire workers, firms must first post vacancies at a flow cost, which are only filled when firms match with a worker. After matching, the worker produces their marginal product of labour (MPL) for the firm each period, and is paid a wage, until the match eventually ends. Thus, the cost of hiring is paid upfront, while the benefit to the firm, marginal product less wage payments, is spread over the future. If firms discount the future more heavily, they will therefore discount the future benefits of hiring more relative to the costs of hiring, which are paid now. This reduces the incentives to hire, and hence raises equilibrium unemployment, and is the core mechanism in the papers of Hall (2017) and Kehoe et al. (forthcoming).

Consider now a model without search frictions but where output is produced using both labour and capital. The production function is $y = zk^{\alpha}l^{1-\alpha}$, where z denotes total factor productivity (TFP). Importantly, this model features diminishing marginal product of labour and, for a given wage, firms hire until the marginal product of labour falls to equal the real wage: $(1-\alpha)zk^{\alpha}l^{-\alpha} = w$. If TFP or capital are too low, and the real wage is sticky and unable to fall, firms will not hire the whole labour force and unemployment will ensue. Following Michaillat (2012), I will refer to this form of unemployment as *rationing unemployment*.

In this simple example hiring decisions are static, but discounts indirectly affect hiring incentives by changing investment in capital. If discount rates rise, firms will invest in less capital. This will reduce the marginal product of labour, reducing hiring incentives and thus increasing unemployment. Given the scale of the decline in investment and capital during the Great Recession, this channel is potentially very important. As shown in Figure 1, detrended capital had fallen by 4.5% from 2008 to 2010, and by more than 15% by 2018. Starting from a 5% unemployment rate, this simple rationing model predicts that a 4.5% fall in capital will raise unemployment to $1 - 0.95 \times (1 - 4.5\%) = 9.3\%$ if the real wage does not adjust.[1] This is not far off the 10% maximum unemployment rate observed in 2010, and hence movements in capital in the Great Recession are potentially large enough to contribute meaningfully to unemployment.

Figure 1: Motivating data

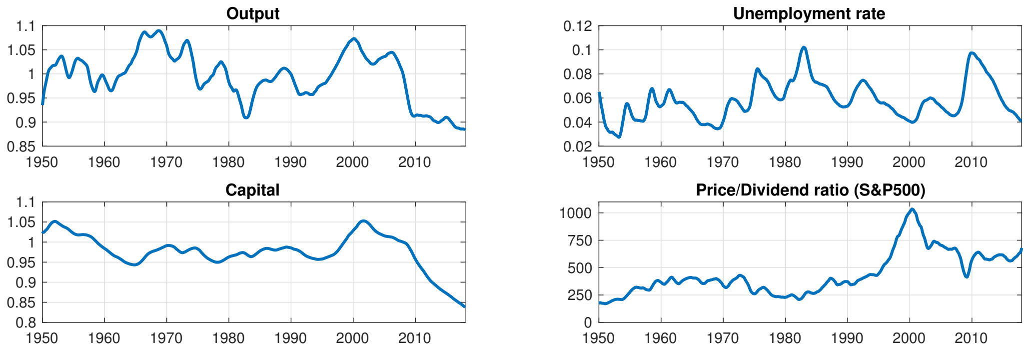

Data for output and capital are real, seasonally adjusted, and detrended with 1950Q1-2008Q1 growth rates, and expressed as deviations from 2008Q1 values. See Section 4.1 for additional details.

These two examples give dramatically different views on how discounts could affect unemployment, and in the remainder of this paper I build a model which nests both mechanisms. Specifically, I extend Hall's (2017) analysis and build a search and matching model featuring endogenous capital accumulation, wage rigidity, and time-varying discount rates for both capital and labour.

My first contribution is to provide analytical results clarifying how discounts are transmitted into unemployment in the presence of capital. I first show that endogenous capital amplifies the response of unemployment to discounts relative to existing models because an additional effect, the rationing unemployment described above, operates. The transmission through capital therefore works fundamentally differently from the effects in the linear search model, which I illustrate by applying Michaillat's (2012) decomposition of unemployment into "rationing" and "frictional" components in my model, formalising the two examples given above.

My second contribution is to estimate discounts from both stock markets and capital, and calculate their impact on unemployment during the Great Recession through the lens of a quantitative model. Series for both discounts are estimated through my structural model, alongside four other shocks which are used to fully explain the time series data on key economic aggregates. Intuitively, stock market discounts are estimated from the stock market Euler equation as the sequence of discount rates required to match the time series for stock prices which cannot be explained by movements in dividends. Similarly, capital discounts are estimated from the capital Euler equation as the sequence of discount rates required to match the time series for investment which cannot be explained by other shocks.

I show that both discounts behave similarly at the beginning of the crisis, spiking rapidly from 2008 until mid-2009, but that the initial rise in stock market discounts is nearly twice as large as the rise in capital market discounts. However, discounts measured from the stock market recover relatively quickly, which reflects the fast recovery in stock prices in the data. On the other hand,

[1] Suppose that the total labour force is equal to 1, and the real wage is fully sticky at w. We can solve for employment as $l = \min\{((1-\alpha)z)^{\frac{1}{\alpha}} w^{-\frac{1}{\alpha}} k, 1\}$, and unemployment as $u = \max\{1-((1-\alpha)z)^{\frac{1}{\alpha}} w^{-\frac{1}{\alpha}} k, 0\}$. If employment is initially equal to 0.95, and TFP and the real wage are constant, a 4.5% reduction in capital will reduce it to $0.95 \times (1-4.5\%)$, raising unemployment to $1 - 0.95 \times (1-4.5\%) = 9.3\%$

discounts measured from capital rise much more persistently, and remain elevated throughout the remainder of the sample. This reflects the persistent decline in investment and capital since the crisis, as shown in Figure 1. Thus, the two sources of data imply different paths for discount rates, and focusing only on stock markets might understate the role of discounts by ignoring the less dramatic, but more persistent, movements in discounts estimated from investment data.

With my estimated discount series in hand, my main numerical results assess the quantitative relevance of stock market and capital discounts in driving unemployment during the Great Recession. To do this, I feed my two estimated discount series through the model, while turning off all other shocks. I have two forward-looking investment choices for the non-financial side of my economy represented by the Euler equation for capital and, due to matching frictions, the Euler equation for employment. A natural question is therefore how to apply the two estimated discount series to each of these decisions. We might expect firms to discount these capital and labour investments differently because of their different cost and cashflow timing structures,[2] and I therefore allow firms to discount employment and capital decisions differently by assigning them possibly different discount rates.

For my baseline results, I apply capital discounts to the capital Euler equation. This is natural, since these discounts were measured from this equation directly to explain movements in capital not attributable to other shocks. Since I do not directly measure a discount rate from the employment Euler equation, I follow Hall (2017) and apply stock market discounts to the employment Euler. This is motivated by the striking correlation between the stock market and unemployment (Figure 1) and job finding rates (Hall, 2017, Figure 8), suggesting that the same forces which drive stock market discounts drive the discounting of hiring decisions, but I show that my results are entirely robust to allocating discounts in different ways.[3]

My main result is that the combined effect of discounts on unemployment is powerful, and they raise the unemployment rate sharply by just under 2.5 percentage points between 2008M1 and 2009M9, and persistently throughout the rest of the sample. Given that unemployment in the data peaked at a maximum of just under 10% in late 2009, discounts alone can therefore explain 52% of the peak rise in unemployment during the Great Recession. Thus, I find that increased discounting of the future is a quantitatively important driver of unemployment during the Great Recession. Furthermore I show that taking into account discounts measured from capital greatly increases the effect of discounts on unemployment, by accounting for around half of their peak impact, and essentially all of their persistence.

Investigating the mechanism, I show that elevated discounts cause a rise in *frictional* unemployment early in the crisis, and *rationing* unemployment later in the crisis, where the effect via rationing is absent in search models which abstract from capital. Additionally, search frictions mean that firms stop hiring early in the crisis in anticipation of future falls in capital, a feature absent from financial frictions models of investment with static labour markets. Finally, I show that while capi-

[2]Capital investments tend to be large and generate returns over long horizons (typical depreciation rates are 6.5-10% *annually*), while estimated hiring costs tend to be smaller and employment relationships shorter (typical job separation rates are around 3% *monthly*).

[3]Hall (2017, p306) argues that "the flow of benefits from a newly hired worker has financial risk comparable to corporate earnings". An alternative approach to the one taken in this paper would have been to *directly* estimate discount rates for employment from the employment Euler equation, $\frac{\kappa}{q_t} = \beta_t^l \mathrm{E}_t \left[mpl_{t+1} - w_{t+1} + \frac{(1-\rho)\kappa}{q_{t+1}} \right]$, where κ/q_t denotes vacancy posting cost over vacancy filling rate, mpl_{t+1} the MPL, w_{t+1} the wage, and β_t^l the estimated discount rate. However, this is more complicated than doing so for capital, since to measure the discount we require a measure of the MPL *net of wages*. This would render discount estimates sensitive to the measure, or model, of wages used, about which there is some disagreement in the literature. This remains an interesting avenue for future research.

tal discounts raised unemployment by reducing labour productivity, this is entirely consistent with the observation that labour productivity was actually rising early in the crisis, since other shocks increased labour productivity during this period both directly and through endogenous changes in the marginal product of labour.

Related Literature. This paper mainly contributes to the literature on the role of discounts in driving unemployment over the business cycle. An early contribution by Mukoyama (2009) questioned the ability of discounts to meaningfully drive unemployment. More recently, Hall (2017) demonstrated that making wages stickier increased the power of variations in discounts. Kehoe et al. (forthcoming) show that adding on-the-job human capital accumulation amplifies the effects of a rise in discounts on unemployment, because the benefits of hiring are felt further in the future, making them more sensitive to the discount rate. Martellini et al. (2019) investigate the relationship between discounts and unemployment in a lifecycle model, and document discrepancies with the data across age groups.

Many search models do include capital accumulation, including contributions by Merz (1995), Andolfatto (1996), Pissarides (2000), and den Haan et al. (2012). Closest to my work are papers which consider discounts and frictions in both investment and hiring. Building on earlier work (Yashiv, 2000, Merz and Yashiv, 2007), Yashiv (2016) estimates a model of joint adjustment costs in capital and labour. He finds that variations in both hiring and investment are driven mostly by their respective expected returns, thus finding an important role for variations in discounts. Relative to this paper, he provides a more detailed specification of the joint capital-labour adjustment costs, but he purposefully does not provide a full structural model of the economy. Hall (2016) extends the linear search and matching model of Hall (2017) to include endogenous capital and discounts for both capital and labour. However, unemployment is modelled in a reduced form manner and does not depend on capital and hence capital discounts, which is the focus of this paper.

Discounts can be understood as a reduced-form way to represent the role of financial frictions in making hiring or capital investment more expensive for firms. For hiring, there is a large recent literature on financial frictions, with notable papers including Wasmer and Weil (2004), Petrosky-Nadeau (2014), Petrosky-Nadeau and Wasmer (2015), Quadrini and Sun (2015), Schoefer (2016), and Carrillo-Tudela et al. (2018). For capital, early contributions include Bernanke and Gertler (1989), Kiyotaki and Moore (1997), and Bernanke, Gertler, and Gilchrist (1999).

In this paper I consider how discounts affect investment in both capital and labour. Existing papers, such as Christiano et al. (2011) and Mumtaz and Zanetti (2016), have built frameworks including financial frictions, capital, and labour market matching which could be reinterpreted through the lens of discounts. My time-varying capital discount rate is similar to the investment wedge in Chari et al.'s (2007) business cycle accounting approach. They model distortions in the labour market as a tax on labour income whereas I explicitly model search frictions. Brinca et al. (2016) apply this approach to the Great Recession, and find an important role for the investment wedge during this episode.

The rest of the paper is structured as follows. In Section 2 I set up the model, and in Section 3 I perform analytical comparative statics. In Section 4 I analyse the Great Recession using my quantitative model, and in Section 5 I conclude.

2 Model

In this section I set up the baseline model to be used throughout the rest of the paper. The model is an extension of the standard search and matching model to include capital, sticky wages, and time-varying discounts. Time is discrete and the horizon is infinite. There are two agents in the model: a representative (multi-worker) firm, and a representative household.

The assumption of time-varying discounts is a stand in for any financial or contracting frictions, or preferences, which induce risk-adjusted time preferences to change. Rather than spell these out, I follow Hall (2017) and simply assume that discount factors are time varying. Since variations in discount rates stand in for all risk-adjusted time preferences, I assume that agents are risk neutral, and markets in the model are complete so that I do not have to consider the firm's financial structure.

2.1 Firms

The representative firm produces output using a Cobb-Douglas production function $y_t = z_t k_{t-1}^{\alpha} l_{t-1}^{1-\alpha}$, featuring one period time-to-build in both capital and labour. There is no intensive margin for labour, so l_{t-1} refers to employment. z_t is a common productivity shock, which follows a Markov stochastic process.

The firm hires labour subject to a standard matching friction. At time t the firm posts v_t vacancies at flow-cost κ per vacancy. The firm takes the vacancy-filling probability q_t as given. I assume that the law of large numbers holds, so a firm posting v_t vacancies receives $q_t v_t$ new employees with certainty, who produce from the next period. Workers separate from the firm at the exogenous rate ρ, and the firm's labour stock thus evolves as

$$l_t = q_t v_t + (1-\rho) l_{t-1}. \tag{1}$$

The firm additionally accumulates a stock of capital according to

$$k_t = i_t + (1-\delta) k_{t-1}, \tag{2}$$

where i_t is investment and δ is proportional depreciation. I assume that adjustment costs in capital are external to the firm, so that capital trades at a price p_t^k which the firm takes as given. The price is given by the weakly increasing function $p_t^k = p^k(i_t/k_{t-1})$, which can be derived as the solution to the profit maximisation problem of a competitive capital goods producing sector with convex adjustment costs. Let variables without time subscripts denote steady state values. I normalise the price of capital to one in steady state, by noting that $i/k = \delta$ in steady state, and assuming that $p^k(\delta) = 1$. For my quantitative work, I choose the commonly used quadratic adjustment cost function for capital, which implies that the price of capital is linear in the investment rate:

$$p^k\left(\frac{i_t}{k_{t-1}}\right) = 1 + \psi_k \left(i_t/k_{t-1} - \delta\right). \tag{3}$$

$\psi_k \geq 0$ controls the degree of capital adjustment costs, and the derivation is relegated to Appendix A.1. Cashflow, e_t, is given by output less wages, vacancy posting costs, and investment: $e_t = z_t k_{t-1}^{\alpha} l_{t-1}^{1-\alpha} - w_t l_{t-1} - p_t^k i_t - \kappa v_t$. Plugging in the equations for the evolution of capital and labour, (1) and (2), gives cashflow in terms of the stock variables:

$$e_t = z_t k_{t-1}^{\alpha} l_{t-1}^{1-\alpha} - w_t l_{t-1} - p_t^k (k_t - (1-\delta)k_{t-1}) - \frac{\kappa}{q_t}(l_t - (1-\rho)l_{t-1}). \tag{4}$$

In a standard setup, the firm maximises the discounted sum of cashflows, subject to the sequence of constraints (1), (2), and (4), discounted by the households constant discount factor, β. The solution to this problem would lead to Euler equations for capital and labour: $p_t^k = \beta \mathrm{E}_t[\alpha z_{t+1} k_t^{\alpha-1} l_t^{1-\alpha} + \ldots$ $p_{t+1}^k(1-\delta)]$ and $\frac{\kappa}{q_t} = \beta \mathrm{E}_t\left[(1-\alpha)z_{t+1}k_t^{\alpha}l_t^{-\alpha} - w_{t+1} + \frac{(1-\rho)\kappa}{q_{t+1}}\right]$. I instead assume that the discount rates for capital and labour are time-varying, and potentially different for each asset. The discount rates between t and $t+1$ for capital and labour are given by β_t^k and β_t^l respectively, and they are assumed to follow exogenous stochastic Markov processes. Thus, the Euler equations are replaced with

$$p_t^k = \beta_t^k \mathrm{E}_t\left[\alpha z_{t+1} k_t^{\alpha-1} l_t^{1-\alpha} + p_{t+1}^k (1-\delta)\right] \tag{5}$$

$$\frac{\kappa}{q_t} = \beta_t^l \mathrm{E}_t\left[(1-\alpha) z_{t+1} k_t^{\alpha} l_t^{-\alpha} - w_{t+1} + \frac{(1-\rho)\kappa}{q_{t+1}}\right] \tag{6}$$

Time varying discount rates now affect the incentives to invest in each asset. All else equal, an increase in either discount rate (i.e. a reduction in the discount factor β_t^i) will lead to a reduction in investment in that asset, since the future benefit is now discounted more heavily against the current cost of acquiring the asset. For both assets, the firm trades off the adjustment cost of acquiring one more unit of the asset today with the discounted benefit of the marginal cashflow generated tomorrow, and the adjustment cost saved tomorrow. Search frictions thus naturally act as a labour hiring cost, with it costing the firm $\frac{\kappa}{q_t}$ to acquire one more worker, just as it costs p_t^k to acquire one more unit of capital. Due to the Cobb Douglas assumption, the MPL is simply equal to a constant fraction of measured labour productivity: $(1-\alpha)z_{t+1}k_t^{\alpha}l_t^{-\alpha} = (1-\alpha)y_{t+1}/l_t$.[4] The worker separation rate, ρ, acts as a depreciation rate, symmetrically to the capital depreciation rate, δ, in the firm's hiring decision. I assume that the firm takes the wage as given when choosing labour and capital.[5]

2.2 Stock market and dividends

The firm's stock price is determined by a standard Euler equation for pricing stocks. The firm maintains a unit mass of shares, such that p_t^s is both the price per stock and the total market value of the firm. As discussed in Section 1, I follow Hall (2017) and assume that the discount rate in the stock market is equal to the discount rate firms apply to hiring, β_t^l. This keeps my model as close as possible to Hall's model, so that discount rates for hiring are measured directly from the stock market. I will therefore refer to "stock market discounts" or "labour discounts" interchangeably.

The Euler equation pricing stocks is

$$p_t^s = \beta_t^l \mathrm{E}_t\left[p_{t+1}^s + d_{t+1}\right], \tag{7}$$

[4]Iterating forward on (6) yields the standard free entry condition modified to include capital in production and time-varying discounts. At time 0 we have $\frac{\kappa}{q_0} = \sum_{t=0}^{\infty}\left(\Pi_{j=0}^{t}\beta_{j-1}^l\right)((1-\alpha)y_{t+1}/l_t - w_{t+1})$ and similar for $t > 0$.

[5]In a multi-worker-firm setting, this is a simplification, as firms may understand that hiring more workers allows them to lower the wage (see, e.g., Brügemann et al., 2019). However, as I show in Appendix A.2, this model can also be exactly reformulated as a single-worker-firm search model with endogenous capital, as in Pissarides (2000), in which case the fixed wage is appropriate. It is not the multi-worker-firm aspect which is crucial for my results, but rather the inclusion of endogenous capital accumulation.

where d_t is both dividends per share and the firm's total dividend payout. As previously mentioned, I do not explicitly build a model of the firm's financial structure, and therefore also do not have a model of the firm's dividend payout policies. Given this, I specify dividends as an exogenous stochastic Markov process, and measure dividends directly from the data.[6] This has the advantage ensuring that I correctly measure discounts in the data by matching stock prices in (7), which would not be the case in a model which generated a counterfactual dividend process. The dividend shock only affects stock prices, and has no effect on any other variables in the equilibrium of the model.

2.3 Labour market and matching technology

I assume a standard constant returns to scale (CRS) matching function for creating new employment matches. This takes the stock of unemployed workers, defined as u_t, and posted vacancies, v_t, to create $m_t = m(u_t, v_t)$ new matches. Market tightness is defined as $\theta_t \equiv v_t/u_t$. Given the CRS assumption, this gives the vacancy filling rate as a function $q_t = q(\theta_t) \equiv m(\theta_t^{-1}, 1)$. The job finding rate, λ_t, is given by $\lambda_t = \lambda(\theta_t) \equiv m(1, \theta_t)$. I maintain as an assumption throughout this paper that the matching function is such that the job finding rate is strictly increasing in tightness.

The Great Recession saw large changes in labour market participation, which make a two-state model problematic for analysing this period. While a fully-fledged three-state model is beyond the scope of this paper, I allow for exogenous changes in participation to improve the accounting of the model. As I show in Appendix E, in general equilibrium different participation assumptions have limited effects on my main results.

I normalise the population to one. Recalling that l_{t-1} is employment at t (chosen at $t-1$) and u_t is unemployment at t, I define total non-participation at t as n_t. These three states sum up to the whole population: $l_{t-1} + u_t + n_t = 1$. Agents are participants in the labour market if they are employed or unemployed, and I define total participation at t as $p_t \equiv l_{t-1} + u_t$. Accordingly, $p_t + n_t = 1$, and unemployment is given by $u_t = p_t - l_{t-1}$. The unemployment rate is defined as $\tilde{u}_t \equiv u_t/p_t$.

Since the population is normalised to one, p_t is both total participation and the participation rate. I assume that the participation rate evolves exogenously according to a Markov process. In terms of flows, I assume that only unemployed workers search for a job, and workers thus cannot transition directly from non-participation to employment.[7] The job separation rate, ρ, gives the total probability that a worker separates from a job into either the unemployment or non-participation pools. Since the participation stock is exogenous, for a given ρ, precisely how I allocate the remaining flows does not affect the equilibrium of the model.

2.4 Wage setting

In this paper I explore two wage-setting protocols. My baseline quantitative model follows Hall (2017), where I assume that firms and workers use the alternating offer bargaining protocol of Hall and Milgrom (2008). For my analytical results, and as a robustness exercise, I also follow Michaillat

[6] Specifically, behind the scenes the firm chooses its financial structure. For example, assuming no issuance of new stocks and letting x_t denote the firm's saving in risk-free bonds, we would have $e_t + r_t x_{t-1} = d_t + x_t$. The stochastic process for dividends therefore controls the firm's mix of debt versus equity, but I do not need to track debt to solve the model since firms real decisions are simply given by (5) and (6).

[7] This assumption is restrictive, since many workers flow from non-participation to employment in the data. In Appendix E I make an alternative participation rate assumption and show that the results are very similar.

(2012), and specify wages in my model in terms of a calibrated exogenous wage rule, rather than as the outcome of an endogenous bargaining process. Importantly, both protocols allow me to directly control the degree of wage rigidity, which is an important object in determining how discount rate shocks are transmitted to the economy. I discuss each protocol in turn below.

2.4.1 Wage setting assumption 1: alternating offer bargaining

In my baseline model I follow Hall (2017) and assume that wages are set by alternating offer bargaining. This bargaining structure serves to create endogenous wage rigidity, since the worker's value is now insulated from the value of unemployment. In my quantitative work I exploit this feature, by calibrating parameters of the bargaining rule in order to match the observed degree of wage rigidity of new hires found in the data. I briefly sketch details of the bargaining game here; see Hall and Milgrom (2008) for a more detailed discussion.[8]

Denote the discounted value to the worker of unemployment by U_t. The value to the worker of having a job is split into two components: W_t denotes the discounted value of wage payments during the job, and C_t the "career value", which gives the expected value following the ending of the job. These are given by

$$U_t = b_t + \beta_t^l \mathrm{E}_t \left[\lambda_t(W_{t+1} + C_{t+1}) + (1-\lambda_t)U_{t+1}\right] \tag{8}$$

$$C_t = \beta_t^l \mathrm{E}_t \left[\rho U_{t+1} + (1-\rho)C_{t+1}\right] \tag{9}$$

$$W_t = w_t + \beta_t^l (1-\rho)\mathrm{E}_t \left[W_{t+1}\right] \tag{10}$$

In my baseline specification I follow Hall (2017) and assume that workers discount the future at β_t^l, which is the same rate that firms apply to hiring decisions. b_t gives the flow benefit to workers of unemployment. I assume that this value is time-varying and follows a stochastic Markov process, which I discuss further below. The marginal value to the firm of hiring a worker is given by the discounted sum of marginal products less wage payments, $J_t = Z_t - W_t$, where

$$Z_t = (1-\alpha)\frac{y_t}{l_{t-1}} + \beta_t^l (1-\rho)\mathrm{E}_t \left[Z_{t+1}\right] \tag{11}$$

Firms and workers bargain over the discounted sum of wage payments, W_t. They take turns making a proposal each period, which the other side may accept or reject. If an offer is rejected the match is delayed by one period and bargaining resumes in the next period unless bargaining breaks down at exogenous probability ψ. The cost of delay to the worker is that they do not receive their wage and receive b_t for one period. The cost to delay to the firm is assumed to be equal to a fraction χ of marginal product, which can be motivated by assuming that the firm can still temporarily operate production, at reduced efficiency, without the worker. In equilibrium firms and workers only

[8] Hall and Milgrom's (2008) bargaining game applies to single-worker firms, where a single worker bargains with a single firm. In my multi-worker firm model, I assume that wages are set as if the firm bargained with a single worker over the marginal product of labour. This can be rationalised in two ways. Firstly, I can assume that the firm negotiates with a trade union of all workers over the share of profit not allocated to capital. Alternatively, this model is isomorphic to a model with single-worker firms, as shown in Appendix A.2. Both yield identical bargaining outcomes to those in the text.

make offers which the other side is indifferent about accepting or rejecting, giving

$$W_t^w + C_t = \psi U_t + (1-\psi)\left[b_t + \beta_t^l \mathrm{E}_t\left(W_{t+1}^f + C_{t+1}\right)\right] \tag{12}$$

$$Z_t - W_t^f = (1-\psi)\left[-\chi(1-\alpha)\frac{y_t}{l_{t-1}} + \beta_t^l \mathrm{E}_t\left(Z_{t+1} - W_{t+1}^w\right)\right] \tag{13}$$

where W_t^w is the worker's indifference value, offered by the firm, and W_t^f is the firm's indifference value, offered by the worker. Following Hall (2017) I do not take a stand on which side is allowed to make the first offer, and instead assume that the equilibrium wage is equal to the average of the two indifference values: $W_t = \frac{1}{2}\left(W_t^w + W_t^f\right)$. Combining this with (12) and (13) gives the equilibrium wage value as

$$W_t = \frac{1}{2}\left[(1-\psi)\left(b_t + \chi(1-\alpha)\frac{y_t}{l_{t-1}}\right) + \psi U_t + Z_t - C_t\right] + (1-\psi)\beta_t^l \mathrm{E}_t\left[W_{t+1} + \frac{C_{t+1} - Z_{t+1}}{2}\right]. \tag{14}$$

Given a solution for W_t, the period-by-period equilibrium wage can backed out from (10). The delay parameter ψ controls the degree of wage rigidity. When $\psi = 1$ bargaining is guaranteed to break down within one period, and the solution reduces to the Nash bargaining solution with equal bargaining weights. As ψ is made smaller the wage becomes more insulated from the value of unemployment, which enters as ψU_t, and the wage becomes more rigid.

The bargained wage is increasing in the worker's flow value of unemployment, b_t, since this improves the worker's bargaining position. Since the wage naturally affects hiring incentives, changes in b_t will affect equilibrium hiring and unemployment. In my quantitative work I perform an exact decomposition of post-war economic data, backing out a series of shocks to exactly replicate data on key economic series. I treat b_t as a shock, which I use in the decomposition to exactly match the observed employment data. Intuitively, I treat this shock as a residual, capturing any movements in unemployment which the other shocks in the model are not able to generate, and also refer to it informally as a "wage shock".[9]

2.4.2 Wage setting assumption 2: Exogenous wage rule

My second wage setting assumption is simply to assume an exogenous wage rule linking the wage to the marginal product of labour. This follows Michaillat (2012), and is useful for two reasons. Firstly, it is a simple and flexible way to specify wages while directly controlling the degree of wage rigidity. Secondly, unlike bargaining-based notions of wage setting this wage rule will always be valid by assumption, even in the limit of no search frictions, making it useful for defining Michaillat's (2012) notion of rationing unemployment.

To construct the exogenous wage rule, consider the market clearing wage in this model in the absence of matching frictions. The market clearing wage is simply the marginal product of labour when all workers are employed ($l_{t-1} = p_t$) and for a given level of capital and technology: $w_t^{mc} = (1-\alpha) z_t k_{t-1}^{\alpha} p_t^{-\alpha}$. I assume that the wage rule is given by

$$w_t = \omega\left(z_t k_{t-1}^{\alpha} p_t^{-\alpha}\right)^{\gamma}, \tag{15}$$

[9] There are many potential reasons that hiring incentives could change, aside from the shocks already in the model, for example downwards nominal wage rigidity, or changes in the value of unemployment. Since I am primarily interested in the role of discounts, I use this single reduced form shock to capture these other factors.

with ω controlling the average level of the wage. Recalling that steady-state values of a variable are denoted by an absence of subscripts, the steady state wage is given by $w = \omega\left(zk^{\alpha}p^{-\alpha}\right)^{\gamma}$. Note that we must set ω such that the steady state wage is below the MPL in steady state ($w < w^{mc} = (1-\alpha)zk^{\alpha}p^{-\alpha}$) to compensate the firm for vacancy posting costs.

In response to changes in either z_t or k_{t-1}, the wage stickiness parameter γ controls how far the wage is able to adjust towards the new market clearing wage. When $\gamma = 0$ the wage is fully rigid and fixed at the steady state wage: $w_t = w$. When $\gamma = 1$ the wage is flexible, and adjusts one-for-one with changes in the market clearing wage: $w_t = w_t^{mc} \times (w/w^{mc})$. Intermediate values of γ give partial wage flexibility.

2.5 Stochastic processes

There are six exogenous stochastic variables driving the economy: the two discounts, β_t^k and β_t^l, TFP, z_t, dividends, d_t, participation, p_t, and the flow value of unemployment, b_t. I assume that all of these variables follow independent AR(1) processes in logs. That is, each variable $x \in \{\beta^k, \beta^l, z, d, p, b\}$ evolves according to $\log x_t = (1-\rho_x)\log x + \rho_x \log x_{t-1} + e_t^x$, where x is the steady state value, ρ_x controls the autocorrelation, and $e_t^x \sim N(0, \sigma_x^2)$. Together these define the stochastic process for the state $s_t \equiv (\beta_t^k, \beta_t^l, z_t, d_t, p_t, b_t)$.

2.6 Definition of equilibrium

I define competitive equilibrium in the case of alternating offer wage setting, and a similar definition applies when wages are set by the exogenous wage rule. Let s^t denote the history of shocks up to time t. Equilibrium is defined as a sequence of functions for prices and allocations which satisfy the restrictions of the model:

Definition 1. *Fix initial conditions (l_{-1}, k_{-1}, s_0). A sticky-wage equilibrium is defined as a sequence of functions*

$$\{y_t(s^t), l_t(s^t), k_t(s^t),\ q_t(s^t), v_t(s^t), u_t(s^t), \theta_t(s^t), w_t(s^t), i_t(s^t), ... \\ p_t^k(s^t), p_t^s(s^t), U_t(s^t), W_t(s^t), C_t(s^t), Z_t(s^t), w_t(s^t)\}_{t=0}^{\infty} \quad (16)$$

satisfying: Production function, $y_t = z_t k_{t-1}^{\alpha} l_{t-1}^{1-\alpha}$; Labour evolution, (1); Capital evolution, (2); Euler equations (5) and (6); Asset pricing, (7); Unemployment definition, $u_t = p_t - l_{t-1}$ and $\tilde{u}_t = u_t/p_t$; Tightness definition, $\theta_t = v_t/u_t$; Wage rule, equations (8), (9), (10), (11), (14); and the stochastic process for s_t.

Since the discount rates are given exogenously, and are not a function of consumption, we can state the equilibrium in the labour and capital markets without reference to market clearing in the goods market. If the capital-production technology behind the function $p^k(i_t/k_{t-1})$ is specified then consumption can additionally be found as the residual from output after vacancy posting and investment costs.

2.7 Rationing and frictional unemployment

A key idea in this paper is that discount shocks may be transmitted into unemployment in different ways. To characterise this more concretely, I utilise Michaillat's (2012) distinction between *rationing* and *frictional* unemployment.

Rationing unemployment is defined as the amount of unemployment a model generates even in the absence of matching frictions. Not all models will generate rationing unemployment, since it requires a mechanism to stop the labour market from clearing even when hiring costs vanish. Michaillat (2012) shows that the combination of wage rigidity and diminishing marginal product of labour delivers rationing unemployment. Intuitively, when hiring frictions vanish firms hire until the wage equals the MPL, and rationing unemployment occurs if the wage is stuck above the market clearing level.

Defining rationing unemployment therefore requires a notion of wage stickiness. Additionally, this notion of wage determination must be well defined even in the limit as hiring frictions vanish, and unemployment (potentially) disappears. Michaillat (2012) therefore assumes that wages are determined by an exogenous (calibrated) rule which controls how much the real wage is able to respond to the marginal product of labour. In order to define rationing unemployment consistently, I follow Michaillat (2012) and define it based on the exogenous wage setting rule (15). Rationing unemployment is then defined as:

Definition 2. *For a given level of capital, k_{t-1}, productivity, z_t, and participation, p_t, rationing unemployment is the equilibrium level of unemployment in the model when matching frictions are removed ($\kappa = 0$). It is given by*

$$u_t^R = \max\left\{0, p_t - \left(\frac{(1-\alpha)z_t^{1-\gamma}k_{t-1}^{\alpha(1-\gamma)}p_t^{\alpha\gamma}}{\omega}\right)^{\frac{1}{\alpha}}\right\}. \tag{17}$$

To derive this expression, note that when hiring frictions disappear, the labour optimality condition (6) reduces to the static condition that the wage equals the MPL: $w_t = (1-\alpha)z_t k_{t-1}^{\alpha} l_{t-1}^{-\alpha}$. This is combined with the wage rule, $w_t = \omega\left(z_t k_{t-1}^{\alpha} p_t^{-\alpha}\right)^{\gamma}$, to find a fully static solution to the level of employment in the absence of rationing frictions, denoted l_{t-1}^R. The formula for rationing unemployment in (17) is simply the level of unemployment implied by this hiring, or zero if this is greater: $u_t^R = \max\left\{0, p_t - l_{t-1}^R\right\}$. The rationing unemployment *rate* is defined as $\tilde{u}_t^R \equiv u_t^R/p_t$.

Whether rationing unemployment exists (i.e. is positive) at a given point in time depends on the level of productivity and capital. Inspecting (17), we see that rationing unemployment will be non-zero for low enough productivity or capital. Intuitively, when either is low, the MPL is also low, which reduces the incentives to hire (given that wages are sticky and do not fully adjust) and hence increases rationing unemployment. Similarly, whenever rationing unemployment is non-zero it is decreasing in productivity and capital.

This formula is a generalisation of Michaillat's (2012; equation 15) definition of rationing unemployment to a model with capital. His definition is simply the special case of this formula with capital and participation constant at one ($k_{t-1} = p_t = 1$). It is interesting to note that Michaillat's (2012) model had no capital in production, so diminishing MPL translated into decreasing returns to scale in the single production factor, labour. My model instead features constant returns in the long run, once capital has been able to adjust.[10]

Finally, frictional unemployment is defined to be the difference between actual unemployment

[10] An alternative definition of rationing unemployment in a model with capital would also allow capital to adjust to its new optimal level in the absence of hiring frictions. This notion of rationing unemployment continues to be well defined, but is more complicated to operationalise in a dynamic model. I discuss this further in Appendix A.3.

and rationing unemployment. Formally:

Definition 3. *For a given level of unemployment, u_t, and rationing unemployment, u_t^R, frictional unemployment is the excess of unemployment over rationing unemployment:*

$$u_t^F = u_t - u_t^R. \tag{18}$$

This is the excess unemployment which cannot be explained by rationing aline, with the frictional unemployment rate defined as $\tilde{u}_t^F \equiv u_t^F / p_t$.

3 Steady state analysis: analytical results

In this section I provide analytical results to illustrate the main workings of the model. My focus is on how an increase in the discount rate is transmitted into unemployment directly via hiring incentives and indirectly through capital, and to decompose this into frictional and rationing components. I focus on a steady state analysis, and perform comparative statics with respect to discount rates. In this section I assume that wages are determined according to the exogenous wage rule (15). This wage rule has the advantage of being analytically tractable, and also means that actual equilibrium unemployment and rationing unemployment are computed using the same wage rule. I additionally suppose that participation is constant at $p = 1$, and maintain two basic assumptions on the parameters: Firstly, parameters are such that employment is interior in steady state across any comparative statics considered. That is, $u, l \in (0, 1)$. Secondly, the discount rate is such that the overall discount rate applied to matches is always positive: $r + \rho > 0$.

Rather than considering each discount rate separately, in this section I assume that capital and labour discounts are equal: $\beta^k = \beta^l = \beta$ and perform comparative statics across the common discount rate, $r \equiv \beta^{-1} - 1$.[11] In order to build intuition about how the introduction of capital affects the model relative to existing work, I instead consider two cases which differentiate between the behaviour of capital. I first consider a version of the model with a fixed capital stock, where capital does not respond to a change in the discount rate. This reduces the model to a search model with labour as the only input and decreasing returns to scale, exactly as in Michaillat (2012). I then consider my full model with endogenous capital, where capital also adjusts to the discount rate through the Euler equation.

3.1 Fixed capital

In this section I consider the model with a fixed amount of capital, but where the labour market has time to adjust to its new steady state.[12] Since I will perform comparative statics across different different steady states, I enhance the notation slightly to improve clarity, and use starred variables to refer to values in the initial steady state: k^*, r^* and so on. I thus hold capital at its initial steady state value, k^*, and vary the discount rate r, studying the effect on the equilibrium. I use a superscript s to denote the new steady state values in the "short run" with fixed capital: $l^s(r), u^s(r), u^{R,s}(r), u^{F,s}(r)$, and so on. I suppress the dependence of these variables on the value of r where it will not cause confusion.

[11] I consider separate discounts in my numerical work, and analytically in Appendix B.2.

[12] Given the high job finding and separation rates in the US, it is often argued that the assumption that the labour market is in steady state provides a reasonable approximation to labour market dynamics (Shimer, 2012).

In steady state, we can write the job filling rate as a strictly decreasing function of equilibrium employment: $q = \hat{q}(l)$ with $\hat{q}'(l) < 0$. Intuitively, if steady-state employment is higher then more vacancies must be being posted, reducing the job filling probability (see Appendix B for derivation).

The steady state level of employment with a fixed capital stock can be found by plugging $\hat{q}(l)$ and the wage rule, (15), into the labour optimality condition, (6), taken in steady state. This gives employment as the implicit solution to

$$\frac{\kappa}{\hat{q}(l^s)} = \frac{(1-\alpha)z(k^*/l^s)^\alpha - \omega z^\gamma (k^*)^{\alpha\gamma}}{r+\rho}, \tag{19}$$

which is the standard condition that the expected vacancy posting cost should equal the discounted surplus accruing to the firm. Unemployment is then defined as $u^s = 1 - l^s$. How does unemployment respond to a rise in discounts, i.e. an increase in r, when capital is fixed? Discounts enter the equation in only one way: in the denominator discounting the firm's surplus. Implicitly differentiating (19) reveals that when discounts rise unemployment must rise. The intuition is the same as in existing search models. The right hand side of (19) gives the discounted benefit of posting a vacancy in steady state. Raising the discount rate lowers the present value of this stream of future benefits to the firm. To encourage vacancy posting this requires the cost of posting a vacancy to fall on the left hand side, which is achieved by lowering tightness, which increases the job filling rate and also increases unemployment.[13]

How is this increase in unemployment divided between frictional and rationing unemployment? It turns out that when capital is fixed the increase must be entirely frictional, as rationing unemployment does not change. To see this, recall the definition of rationing unemployment in (17). For a fixed level of capital (and productivity) this is simply a constant,

$$u^{R,s} = \max\left\{0, 1 - \left(\frac{(1-\alpha)z^{1-\gamma}(k^*)^{\alpha(1-\gamma)}}{\omega}\right)^{\frac{1}{\alpha}}\right\}, \tag{20}$$

and hence a rise in discounts can have no effect on rationing unemployment. This is very intuitive: rationing unemployment is simply the level of unemployment which equates the wage to the MPL in the absence of hiring costs. For a given level of capital, the MPL is a fully static object which does not depend on discounts.

Frictional unemployment is defined as $u^{F,s}(r) = u^s(r) - u^{R,s}(r)$, and since the rise in discounts causes unemployment to rise without a rise in rationing unemployment, the entire increase must be frictional. This is also intuitive. As explained above, the rise in unemployment was due to the discounted benefit of hiring falling relative to vacancy posting costs. This is a purely frictional phenomenon, which is naturally categorised as a rise in frictional unemployment. I summarise the findings from this section in the following proposition.

Proposition 1. *When capital is fixed, following a rise in discounts ($\uparrow r$): 1) unemployment rises ($\uparrow u^s$), 2) rationing unemployment is constant ($\overline{u^{R,s}}$), 3) frictional unemployment rises ($\uparrow u^{F,s}$).*

Proof. See text above and Appendix B.1. □

[13] Relative to the linear (constant returns to scale) model of Hall (2017), there is also a dampening effect in my model due to decreasing returns to scale, which is represented by the positive term $\alpha(1-\alpha)z(k^*)^\alpha l^s(r)^{-\alpha-1}$ in the denominator.

3.2 Flexible capital

In this section I move on to considering my full model, when the level of capital is also free to adjust to a rise in discounts. The addition of endogenous capital is the novelty of this paper, and I show that the transmission of the discount rise to unemployment is very different. I use a superscript l to denote "long-run" steady state values with endogenous capital: $l^l, u^l, u^{R,l}, u^{F,l}$ and so on. I consider an increase in r, but now allow capital, k, to optimally adjust according to the Euler equation, (5).

Recall that with my choice of capital adjustment cost function the price of capital is constant at $p^k = 1$ across all steady states. The capital Euler equation, (5), in steady state gives a solution for the capital-labour ratio in steady state:

$$\frac{k}{l} = \left(\frac{\alpha z}{r+\delta}\right)^{\frac{1}{1-\alpha}} \equiv \mathcal{K}(r). \tag{21}$$

I use $\mathcal{K}(r)$ to summarise the capital-labour ratio as a function of the discount rate (suppressing the dependence on productivity, which I take as constant. (21) shows us that $\mathcal{K}'(r) < 0$, since when agents are more impatient they are less willing to invest in capital, and the capital-labour ratio must fall to raise the MPK. Combining $\mathcal{K}(r)$ with the labour market Euler, (6), and wage rule, (15), gives employment with endogenous capital as the implicit solution to

$$\frac{\kappa}{\hat{q}(l^l)} = \frac{(1-\alpha)z\mathcal{K}(r)^\alpha - \omega z^\gamma (\mathcal{K}(r)l^l)^{\alpha\gamma}}{r+\rho}, \tag{22}$$

where I also used that $k^l = \mathcal{K}(r)l^l$. Just as with fixed capital, this condition states that employment must be such that expected vacancy posting costs equal the discounted sum of surpluses to the firm. However, capital is now endogenous, and is replaced with the optimal value coming from the capital Euler equation. Unemployment with flexible capital is symmetrically defined as $u^l = 1 - l^l$.

As before, raising discounts must raise unemployment, which can be seen by implicitly differentiating (22). Intuitively, (22) shows that when capital is endogenous there are now two ways that discounts affect unemployment. Firstly, and commonly with the model with fixed capital, the r in the denominator on the right hand side means that raising discounts directly lowers the firm's discounted surplus from hiring. Secondly, and differently from the model with fixed capital, raising discounts lowers capital intensity ($\mathcal{K}'(r) < 0$) which reduces the MPL through capital shallowing, and hence reduces the firm's surplus when wages are sticky. Notice instead that when capital is fixed, capital intensity actually goes *up* following a rise in discounts, as employment falls while capital is fixed.[14]

How does the increase in unemployment with endogenous capital map into changes in frictional and rationing unemployment? Rationing unemployment with flexible capital is given by the definition, (17), when capital is equal to its long run steady state value:

$$u^{R,l}(r) = \max\left\{0, 1 - \left(\frac{(1-\alpha)z^{1-\gamma}\left(\mathcal{K}(r)l^l(r)\right)^{\alpha(1-\gamma)}}{\omega}\right)^{\frac{1}{\alpha}}\right\}, \tag{23}$$

[14]The difference between the response of the capital-labour ratio in the "short" versus "long run" of the model is consistent with the data during the Great Recession, where the capital-labour ratio initially rises before eventually permanently falling.

where I again used that $k^l = \mathcal{K}(r)l^l$ to replace the level of capital. Differently from the model with fixed capital, rationing unemployment now responds to discounts through the response of capital. In particular, if rationing unemployment is non-zero, then raising discounts will lower capital intensity and total capital, and hence increase rationing unemployment by reducing the MPL and the incentive to hire. Formally, this can be seen by differentiating (23) when $u^{R,l}(r) > 0$. This is a key novelty of my model relative to Michaillat (2012). In his model there is no capital, and hence no way that discounts can affect rationing unemployment. In my model with capital, discounts can affect rationing unemployment by changing capital intensity.

The degree to which rationing unemployment responds to discounts depends crucially on wage rigidity. In particular, the less rigid are wages (i.e. the larger is γ) the less rationing unemployment responds to discounts. This can be seen by the dependence of rationing unemployment on the exponent $\alpha(1-\gamma)$ in (23). In the limit of fully flexible wages ($\gamma = 1$) rationing unemployment does not respond at all to the level of capital, and hence the level of discounts. On the other hand, when wages are fully rigid (i.e. $\gamma = 0$), rationing unemployment responds strongly to discounts, because in this case rationing unemployment moves linearly with capital.

Frictional unemployment is defined as before, as the residual unemployment not explained by rationing unemployment: $u^{F,l}(r) = u^l(r) - u^{R,l}(r)$. What happens to frictional unemployment following a rise in discounts thus depends on whether total unemployment or frictional unemployment rises more. Consider the more interesting case of what happens to frictional unemployment when rationing unemployment is already positive. In this case, it turns out that frictional unemployment could either rise or fall following a rise in discounts, and this depends crucially on the degree of wage flexibility. In particular, in the limit of fully rigid wages, frictional unemployment must fall, and in the limit of fully flexible wages, it must rise. These results are summarised in the following proposition.

Proposition 2. *Suppose that parameters are such that rationing unemployment is currently positive ($u^{R,l}(r) > 0$). When capital is flexible, and chosen according to the Euler equation, (5), following a rise in discounts ($\uparrow r$): 1) unemployment rises ($\uparrow u^l$), 2) rationing unemployment rises ($\uparrow u^{R,l}$), 3) frictional unemployment may either rise or fall ($\uparrow\downarrow u^{F,s}$).*

Proof. See text above and Appendix B.1. □

Finally, given the ambiguous response of frictional unemployment when capital is flexible, one might wonder whether the response of unemployment to a rise in discounts is greater or smaller than when capital is fixed. In fact, as I summarise in the final proposition, it is possible to show that the increase in unemployment is always greater when capital is flexible.

Proposition 3. *Following a rise in discounts ($\uparrow r$) from the initial steady state level r^*, unemployment increases more when capital is flexible than in the model where capital is fixed. That is, $\frac{d}{dr}u^l(r^*) > \frac{d}{dr}u^s(r^*)$.*

Proof. See Appendix B.1. □

In summary, in this section I demonstrated two results. Firstly, adding endogenous capital to the search model amplifies the effect of rising discounts on unemployment. Secondly, endogenous capital also fundamentally changes the transmission mechanism from discounts to unemployment. When capital is fixed, the transmission is through frictional unemployment, while it is through rationing unemployment (with ambiguous response of frictional) when capital is endogenous.

4 Discounts in the Great Recession

In this section I estimate discounts during the Great Recession and assess their effects on unemployment. I move on from the comparative statics exercises of the previous section, and now use the full dynamic version of my model. I perform counterfactual exercises by passing estimated shock series through the model, which is solved both linearly and nonlinearly.

4.1 Data

I first describe the data used for the exercises. Since I am focusing on the Great Recession, which falls at the end of the sample for current data, using filtering to extract the business cycle component of the data is problematic. Instead, I use the simpler approach of taking the cyclical component of the data to be the deviation from the pre-crisis trend.[15]

National accounts data is quarterly, while labour market and stock market data is available at a monthly frequency. I solve my model at a monthly frequency, and linearly interpolate national accounts data to create monthly data. After combining my monthly and quarterly data I have a monthly dataset running from 1950M1 to 2018M1. Data sources are standard, and so I relegate details to Appendix C, and note here only that 1) I use Fernald's (2014) national accounts and utilisation adjusted TFP data, and 2) I use data on both aggregate employment and unemployment, and reconcile the two by backing out an appropriate measure of labour market participation.

The data for the Great Recession is plotted in Figure 2 (data for the whole sample is given in Figure 11 in the appendix). Data for all series except for the unemployment rate are deviations from their 2008M1 values. The bottom right panel shows a peak decline in employment of nearly 10% by 2010, which never fully recovers even by 2018. The top right panel gives the unemployment rate. Unemployment rises up to a peak of just below 10% by 2010, before gradually declining and recovering by 2016. In the bottom centre panel we see that capital falls more gradually, but has fallen nearly 5% by 2010, and continues to fall to a deviation of over 15% by the end of the sample. The top left panel shows adjusted and unadjusted TFP. These two series show a discrepancy at the beginning of the crisis, where raw TFP falls by 3% in 2009 before completely recovering by 2010. Mechanically, this behaviour is because output falls around one year before employment, which eventually catches up. Thus, it appears to be driven by temporary under-utilisation of resources in the first year of the crisis. Accordingly, Fernald's (2014) method to correct for utilisation changes does not find a fall in TFP, but instead finds a temporary rise at the beginning of the crisis, which is also then undone. I choose to use utilisation adjusted TFP as my measure of productivity when estimating the model.

The bottom-left panel shows labour productivity – output over employment – computed both using the raw output data and using the same utilisation adjustment series as for TFP. This behaves qualitatively similarly to TFP, but the movements are more pronounced, since labour productivity also responds to changes in the capital-labour ratio. Labour productivity rises nearly 4% at the beginning of the crisis, before falling nearly 8% by the end of the sample.

[15] In practice, pre-crisis the results of my simple detrending versus using the HP-filter with the literature's standard smoothing parameter of 10^5 are very similar. However, during the crisis, the well known end-point problem of the HP-filter means that it finds that the cyclical component of GDP has fully recovered by the end of the sample, despite GDP being more than 10% below the pre-crisis trend. See Appendix D.5 for details.

Figure 2: Great recession data

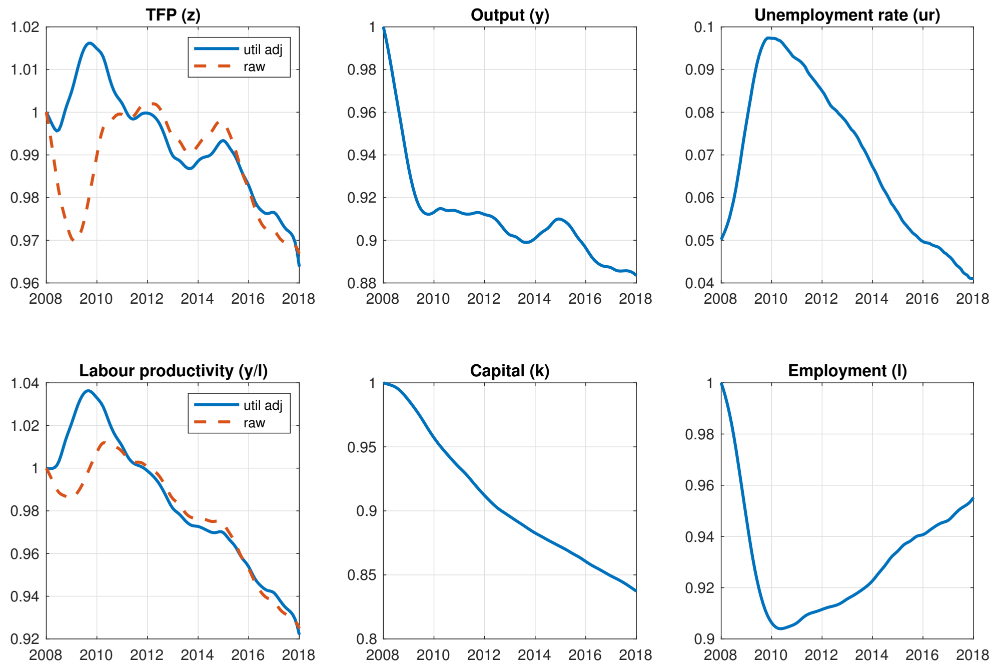

Data for aggregates during Great Recession period. Data for TFP, labour productivity, employment, output, and capital are real, seasonally adjusted, and detrended with 1950Q1-2008Q1 growth rates, and expressed as deviations from 2008M1 values. TFP and labour productivity are given as both raw (blue) and utilisation adjusted (dashed red) series. Unemployment is the seasonally adjusted civilian unemployment rate. See Appendix C for additional details.

4.2 Calibration and estimation

My core quantitative exercise is to investigate the importance of discounts, relative to other shocks, in driving unemployment during the Great Recession. I do this using a full decomposition exercise for the Great Recession, where I use the model to extract several shocks from the data in order to exactly match the paths of key economic aggregates. In order to do this, I return to the full model as specified in Section 2, with wage setting determined by alternating offer bargaining.

The model is calibrated and estimated in two steps. Firstly, most parameters are calibrated to target moments in the non-stochastic steady state of the model. Secondly, the remaining parameters are estimated using an iterative procedure, to match dynamic features of the data. Since my focus will be on the model's ability to replicate the Great Recession, I estimate the model using pre-crisis data. For the simulation used in my estimation and baseline results, the model is solved using first order perturbation techniques.[16] I perform robustness using a nonlinear perfect-foresight solution in Appendix E.

The model is calibrated at a monthly frequency. All parameters are calibrated to the non-stochastic steady state of the model, apart from i) the wage bargaining negotiation breakdown

[16]Log-linearisation is performed using the Dynare toolbox (Adjemian et al., 2011) and model simulation and estimation are carried out manually using Matlab. I simulate capital, employment, and the real wage using their log-linearised policy functions, and compute all other variables using their exact nonlinear equations.

parameter, χ, which controls the degree of wage rigidity and ii) the stochastic processes for the shocks. A detailed discussion of the parameter values and methodology is given in Appendix D.1, and I provide a summary below. Calibrated parameter values are summarised in Table 1.

I choose standard values on the production side to match the labour share and rate of capital depreciation. Following Bernanke et al. (1999) and Brinca et al. (2016) I choose the adjustment cost parameter ψ_k to match an elasticity of the price of capital to the investment rate of 0.25 in the steady state. I also perform robustness allowing for higher and lower adjustment costs. The steady state capital discount is chosen to match an investment-output ratio of 18% as in the data, which gives a 5.37% yearly return on capital. For the stock market, I target a steady state price to (monthly) dividend ratio of 407.5, equal to the pre-crisis average of my stock market data. The steady-state labour discount is measured from the stock market, given this price-dividend ratio, which gives a 3.0% yearly return.[17]

The labour market is parameterised to stay close to Hall's (2017) strategy. I take the job separation rate to be $\rho = 3.45\%$, following Shimer (2005), and choose the job finding rate to match an unemployment rate of $\tilde{u} = 0.05$. I assume a Cobb-Douglas matching function, $m_t = \psi_0 u_t^{\psi_1} v_t^{1-\psi_1}$, with standard elasticity $\psi_1 = 0.5$ (Petrongolo and Pissarides, 2001). The steady state participation rate is calibrated to match the average employment rate in Fernald's (2014) data between 1950M1 and 2008M1. I set the vacancy cost equal to 32% of the steady state wage, following Michaillat (2012). Combined with the Euler equation for labour, this implies that the wage is 98.2% of the marginal product of labour in steady state. I set the worker's value of unemployment as 40% of the steady state wage, motivated by empirical replacement rates and following Hall (2017) and Shimer (2005). I back out the firm's bargaining disruption cost, χ, so that the targeted real wage is supported as the equilibrium of the wage bargaining game.

The remaining parameters are estimated on simulated data using an iterative procedure. The estimation exploits the invertibility of the model, and backs out the required paths for the six underlying shocks in order to exactly match the historical data for six key macroeconomic aggregates: the unemployment rate, $\tilde{u}_t$, output, y_t, capital, k_t, employment, l_t, stock prices, p_t^s, and dividends, d_t. The autocorrelations and variances of shocks are estimated using OLS regressions on the recovered shocks.[18] For the estimation of labour discounts, stock price data displays both high and low frequency movements. Since I am interested primarily in high frequency movements (for example, during the Great Recession the S&P500 crashed and had fully recovered within 2.5 years) I estimate the AR(1) process for β_t^l on first differences using a lagged-IV strategy, which can consistently estimate ρ_{β^l} in the presence of an unobserved mean which follows a random walk. I discuss this estimation strategy, and robustness to applying it to other variables, in Appendix D.2. This yields a value of $\rho_{\beta^l} = 0.92$, which implies a half life of innovations of around 9 months. This is less persistent than capital discounts, consistent with the path for estimated discounts in Section 4.3.1.

Finally, the degree of wage rigidity is crucial for the model's ability to generate volatile unemployment in response to shocks. This is controlled by the bargaining breakdown probability, ψ, which I estimate by targeting the degree of wage flexibility observed in the data. Haefke et al. (2013)

[17] The return on the stock market is lower than that measured in the data due to detrending, since the model does not account for the average growth rate of dividends. The returns on stocks and bonds should be interpreted as returns net of the average growth rate of dividends and TFP respectively.

[18] For the participation rate and dividends the Augmented Dickey-Fuller test fails to reject the hypothesis of a unit root, and OLS estimates an autoregressive coefficient above one. I choose to set the autocorrelation of participation to a large but stationary number, and set the autocorrelation of dividends to that of TFP, motivated by the long-run dependence of dividends on economic conditions. Results are robust to reasonable variations in these parameters.

estimate that, after controlling for composition effects, the elasticity of real wages of new hires to aggregate labour productivity is 0.7 for production and supervisory workers. I generate an equivalent dataset to theirs by simulating my model subject to all shocks apart from the wage shock, and calibrate ψ to match their estimated elasticity. It is worth noting that these wages are relatively flexible.

Table 1: Calibrated parameters

Parameter	Interpretation	Value	Source / target
z	s.s. tfp	0.5140	Normalise $y = 1$
α	capital elasticity	0.3213	Labour share 2/3
δ	depreciation	0.0056	Annual depreciation 6.5%
β^l	s.s. labour discount	0.9976	$p/d = 407.5$
β^k	s.s. capital discount	0.9956	$i/y = 0.18$
ψ_k	capital adj. costs	44.762	$\epsilon_{p^k,i/k} = 0.25$ (Brinca et al., 2016)
ρ	job separation rate	0.0345	Shimer (2005)
ψ_0	match efficiency	0.6555	Normalise $\theta^* = 1$
ψ_1	match elasticity	0.5	Petrongolo and Pissarides (2001)
κ	vacancy posting cost	0.4143	$\kappa = 0.32w^*$ (Michaillat, 2012)
b	worker flow u. value	0.5178	$b = 0.4w$
d	s.s. dividend	0.0025	Normalise $p = 1$
χ	firm delay cost	0.5674	$\tilde{u} = 0.05$
ψ	prob of bargaining collapse	0.0225	wage flex. (Haefke et al., 2013)
ρ_{β^l}	autocorr. labour discount	0.9251	directly estimated
ρ_{β^k}	autocorr. capital discount	0.9907	directly estimated
ρ_z	autocorr. TFP	0.997	directly estimated
ρ_p	autocorr. participation	0.999	directly estimated
ρ_d	autocorr. dividends	0.997	autocorrelation of TFP
ρ_b	autocorr. unemp. value	0.9961	directly estimated
σ_{β^l}	sd. labour discount	0.0021	directly estimated
σ_{β^k}	sd. capital discount	1.27e-04	directly estimated
σ_z	sd. TFP	0.0011	directly estimated
σ_p	sd. participation	8.24e-04	directly estimated
σ_d	sd. dividends	0.0041	directly estimated
σ_b	sd. unemp. value	0.0022	directly estimated

Baseline calibration and estimation. Model is calibrated at monthly frequency. See Section 4.2 for a detailed discussion of calibration choices. First block gives parameters which are calibrated in steady state, and second block those which are estimated from model simulations.

4.3 Results

Having calibrated the model and estimated paths for the shocks driving the economy, I now move on to my main results. I first discuss my estimated discount series, and then investigate their role in driving unemployment during the Great Recession.

4.3.1 Estimated discounts during the Great Recession

The historical paths for the discount rate shocks for capital and labour are backed out from the data in the estimation procedure described in the last section. The path for labour discounts is the path required to match the path for stock prices in the data, as per the stock pricing equation, (7) . The path for capital discounts is backed out from the full model simultaneously with the other shocks in order to jointly match the data on capital and employment. However, intuitively it is identified as the path for capital discounts required to explain the path for capital, as per the capital Euler equation, (5).

Figure 3: Estimated discounts from the stock market and capital

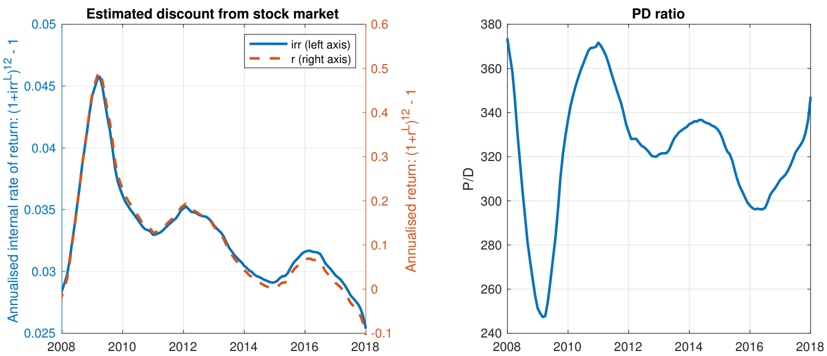

(a) Stock discounts

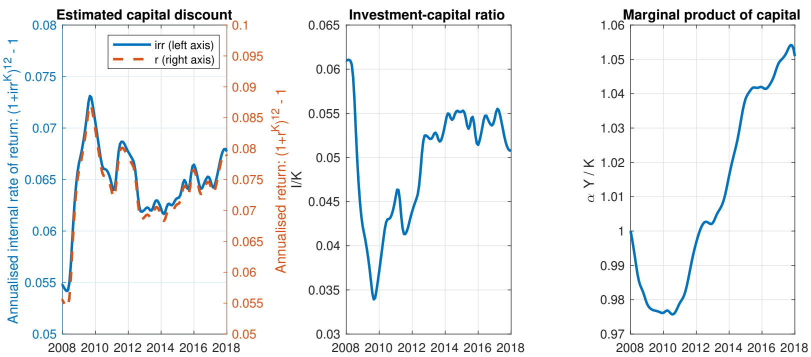

(b) Capital discounts

Estimated discount shocks and key identifying data for the Great Recession. Panel (a) gives discounts estimated from the stock market, and panel (b) discounts estimated from capital. Discounts are expressed both as annualised returns, and as annualised internal rates of return. Data are those inputted into the model, and are thus detrended and seasonally adjusted. See Section 4.1 for additional details.

The paths for the discount shocks during the Great Recession are plotted in Figure 3. For each discount $i = \{l, k\}$, the results are presented in two ways. Firstly, I plot the actual value of the monthly discount, β_t^i. However, when discounts are time varying, the month-to-month discount at t is not fully informative about, for example, the incentives to invest in capital at time t. Since installed capital will persist for many periods, until it is fully depreciated, all of the discount rates over the near future will be important for calculating the total discounted returns from the investment. Therefore, following Hall (2017), I also plot the internal rate of return, which summarises the total discount rate over the life of an investment. This is calculated as the constant discount rate which sets the net present value of the investment to zero, and is effectively an average of the rates of return over the life of the investment weighted by the size of the associated payoffs at each date. For details, see

Appendix D.3.

The leftmost panels of Figure 3 plot the estimated discounts. The behaviour of the monthly discounts and internal rates of return are almost identical, except that the movements in the internal rates of return are of a smaller scale since they represent averages of expected future monthly discounts. The top left panel plots labour discounts, estimated from the stock market. Starting in 2008, these show a sharp rise of just over 60% by 2009M4, from an initial internal rate of return of 2.8% to 4.6%. Discounts measured from the stock market recover relatively quickly: they are only elevated by 15% by 2011M1, and have fully recovered by the end of the sample. This is due to the behaviour of stock prices over this period. The top right panel plots the price-dividend ratio in my stock market data. The sharp rise is driven by the collapse of stock prices in during 2008, since higher discount rates cause stock prices to fall by reducing the demand for stocks as future dividends are more heavily discounted. Similarly, the relatively quick recovery is driven by the recovery of stock prices, which began in 2009.

The bottom left panel plots capital discounts, estimated from the model. Starting in 2008, these show a more modest rise of 34% by 2009M9, from an initial internal rate of return of 5.4% to 7.3%. While the rise is smaller than the rise in stock discounts, the rise in capital discounts is much more persistent: they recover less than half of their value by 2011M1 and remain severely elevated for the remainder of the sample. Inspecting the capital Euler equation, (5), the behaviour of capital discounts can be understood through the behaviour of capital prices, p_t^k, and the marginal product of capital. The bottom centre panel plots the investment-capital ratio, which drives the price of capital. This collapses at the beginning of the recession as investment slowed, but begins a strong recovery from 2009M9, and is the key driver of the initial spike in capital discounts. The collapse in investment causes the price of capital to fall which requires estimated discounts to rise, since higher discounts reduce the demand for, and hence price of, capital. The persistent rise in capital discounts is driven by the fall in capital and hence the increase in the marginal product of capital. The bottom right panel shows that the marginal product of capital has been persistently elevated since 2012. This explains the permanent rise in discounts, since, all else equal, a higher marginal product of capital must be matched with a rise in discounts in the steady state of the Euler equation. Intuitively, since investment and capital have been persistently lower since the crisis than can be explained by other shocks, this must be because the discount rates applied to capital have been higher.

4.3.2 Role of discounts in driving unemployment

Given the linearity of the approximated estimated model, I am able to exactly decompose the paths of all endogenous variables into contributions from each observed shock, with the total summing to reproduce the data. My main results concern the role of discounts in driving unemployment during the Great Recession. To generate these results, I create counterfactual simulations by feeding only the estimated discount shocks through the model.

I construct the counterfactual for any shock $x = \{\beta^l, \beta^k, z, d, p, b\}$ as follows. I first initialise the model in 2008M1 using the observed values of the states and all shocks. I then feed the observed paths for the shock in question, $\{x_t\}_{t=2008M1}^{2018M1}$, through the model, while setting the innovations for all other shocks, $\{e_t^{x'}\}_{t=2008M1}^{2018M1}$ for $x' \neq x$, to zero so that the variables converge back to their steady steady values. I then simulate the model, calculating the counterfactual values for all endogenous variables of interest. Finally, I subtract the counterfactual values from the counterfactual where all shocks (including x) are switched off to remove the effect of natural drift of shocks and endogenous states back to steady state.

The main results are given in Figure 4, which plots the response of key aggregates to the estimated discount shocks during the Great Recession. The data are given in the thin black line, the response of the economy to both discounts combined is given in the solid blue line, and the responses to capital and labour discounts individually are given in dashed red and dash-dotted yellow respectively. The top left panel plots the counterfactual unemployment rate. The combined effect of discounts is powerful: the unemployment rate rises sharply by just under 2.5 percentage points from 5% to 7.45% between 2008M1 and 2009M9. Given that unemployment in the data peaked at a maximum of just under 10% in late 2009, discounts alone can therefore explain 52% of the peak rise in unemployment during the Great Recession. Moreover, this effect is very persistent, and higher discounts cause elevated unemployment throughout the rest of the sample. Thus, I find that increased discounting of the future is a quantitatively important driver of unemployment during the Great Recession.

Figure 4: Response to estimated discount shocks

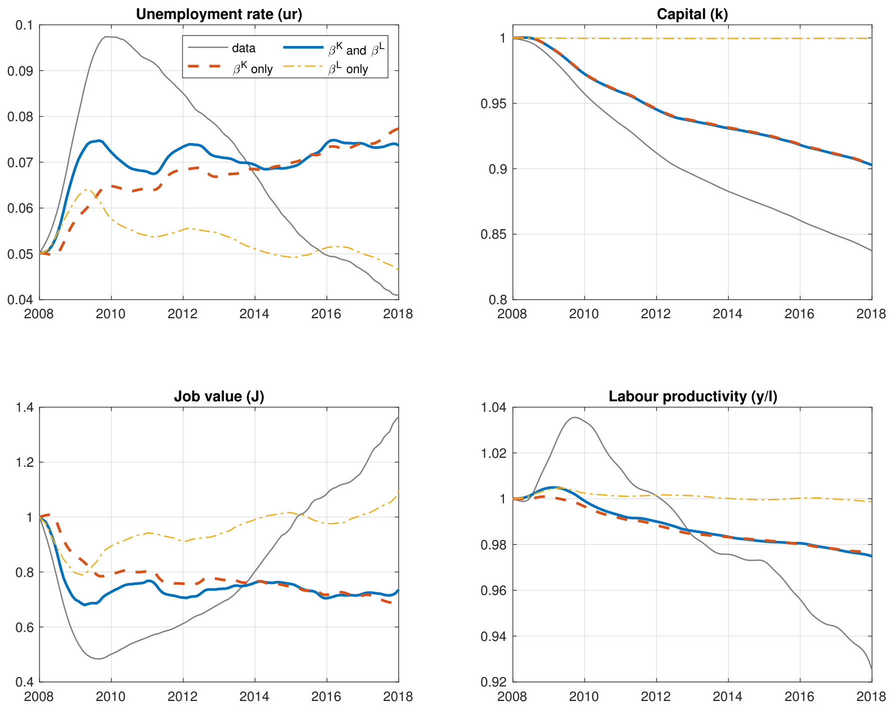

Response of economy to estimated discount shocks. Variables are given as deviations from 2008M1 values, apart from the unemployment rate. Thin black line gives the data. Solid blue line gives response to both discount shocks combined. Dashed red and dash-dotted yellow decompose this into capital discount and labour discount shocks respectively. See Section 4.3.2 for details of counterfactual construction.

Decomposing this into the contributions of the two individual discounts, there are interesting differences in the timing of when each discount contributes to unemployment. The rise in labour discounts causes unemployment to rise immediately and powerfully, leading to 1.4pp rise by 2009M5. This rise follows the sharp increase in labour discounts early in the crisis, estimated from the stock market collapse. Following the relatively quick recovery of the stock market, the increase in unemployment caused by labour discounts also reverses relatively quickly, and has completely recovered by the end of the sample.

The rise in capital discounts causes a more gradual, but much more persistent, rise in unemployment. The increase in unemployment caused by capital discounts lags the effect of labour discounts, but is of a similar magnitude early in the crisis. However, later in the recession capital discounts are still placing upwards pressure on unemployment even as labour discounts have faded, causing increases in unemployment of at least 2pp in the later years of the sample. Again, this reflects the estimated path for capital discounts, which rose less strongly but much more persistently than labour discounts, reflected in the persistent fall in investment and capital seen since the crisis. In terms of relative importance, the rise in labour discounts causes essentially all of the early increase in unemployment in 2008, and explains more than 50% of the rise in unemployment caused by both discounts up to 2009M6. From then on the effect of labour discounts fades, and capital discounts explain essentially all of the discount-led unemployment in the latter half of the sample. This shows the importance of including capital when modelling the role of discounts on unemployment, since capital discounts explain half of their combined peak impact, and all of their persistence.

The remaining panels of Figure 4 give the counterfactual series for additional aggregates which explain the role of these discounts in driving unemployment. The bottom left panel plots the marginal job value, J_t. This is defined as the right hand side of the labour Euler equation, (6), such that tightness is determined as $\frac{\kappa}{q(\theta_t)} = J_t$. Discounts increase unemployment by lowering this job value, leading to lower vacancy posting and hiring. This series thus closely tracks the unemployment rate itself, in the top left panel. What is different across the two discount rates is precisely *how* they affect the job value.

The rise in capital discounts lowers the job value by reducing investment and hence lowering the marginal product of labour. Capital is plotted in the top right panel. The rise in capital discounts causes capital to fall by 9.7% by then end of the sample. Comparing this to the total fall in capital in the data, capital discounts cause 59.5% of the total 16.2% fall. Through capital shallowing, this reduction in capital makes workers less productive: recall that the MPL is given by $y_t/l_{t-1} = z_t(k_{t-1}/l_{t-1})^\alpha$. Labour productivity is plotted in the bottom right panel, and capital discounts reduce labour productivity from 2009 onwards, and by 2.4% by the end of the sample. The reduced productivity of workers reduces the job value and hence the incentive to hire, which causes the rise in unemployment. Interestingly, the forward looking nature of hiring in the search and matching model means that capital discounts raise unemployment early in the crisis – 1.5pp by 2009M10, or 54% of the total rise by the end of the sample – even though capital has only fallen by 2.4pp, or 25% of its eventual total fall. This is due to search frictions, as firms stop hiring early in the crisis in anticipation of further declines in capital and labour productivity. This feature is absent in models where discounts (or financial frictions) affect capital but without search frictions in labour markets.

Rising labour discounts lead to negligible changes in capital, due to the low estimated persistence of labour discounts and the presence of capital adjustment costs. Accordingly, since they raise unemployment and lower employment, labour discounts actually lead labour productivity to *rise* during the crisis. Accordingly, the decline in job value and rise in unemployment due to labour discounts is predominantly due to the direct effect of the increased labour discount rate reducing the discounted value of future firm surpluses.

4.3.3 Role of rationing and frictional unemployment

Having established that rising discounts are quantitatively important drivers of unemployment during the Great Recession, I investigate how this decomposes into frictional and rationing unemployment. Within any counterfactual, the definition of rationing unemployment remains as given in equation (17), and can be calculated from the counterfactual levels of TFP and capital. The results are plotted in Figure 5. The left panel looks at the combined effect of both discount shocks, and decomposes the rise in unemployment into frictional and rationing components. I find that the rise in frictional unemployment explains essentially all of initial run-up of (counterfactual) unemployment, and 49% of the peak in 2009M9. From then on, the level of frictional unemployment is actually falling, and becomes lower than the pre-crisis level from 2010M9 onwards.

Pre-crisis, rationing unemployment was almost nonexistent, explaining just 0.3pp of the 5% unemployment rate. However, from 2008M8 the rationing unemployment rate starts monotonically rising, and explains at least 50% of the counterfactual increase in unemployment from 2009M11 onwards. To understand why, recall the definition of rationing unemployment in (17). Rationing unemployment is the level of unemployment which would prevail only due to wage rigidity, in the absence of search frictions. It rises whenever TFP or capital falls, since these reduce the marginal product of labour. In Figure 4 I showed that discounts cause capital to start falling from 2008M8 onwards. The rise in rationing unemployment simply mirrors this fall in capital, as it reduces the fundamental benefit to firms of hiring.

Figure 5: Rationing vs. frictional unemployment rate decomposition

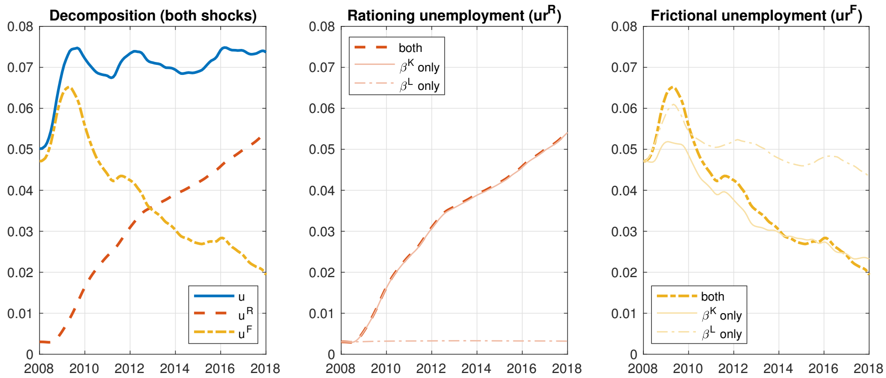

Response of economy to estimated discount shocks: decomposition into frictional and rationing unemployment. Left panel plots the counterfactual unemployment rate path when the economy is subjected to both discount shocks (solid blue line) and then decomposes this into the rationing (dashed red) and frictional (dash-dotted yellow) unemployment rate. The remaining two panels further decompose rationing and frictional unemployment into their contributions from the two separate discount shocks. See Section 4.3.2 for details of counterfactual construction.

Thus, high rationing unemployment explains the persistence of the effect of discounts on unemployment, while frictional unemployment is more important for the initial rise in unemployment. The remaining panels of Figure 5 explore the underlying drivers of frictional and rationing unemployment. The centre panel decomposes rationing unemployment into the role of labour and capital discounts separately. Unsurprisingly, capital discounts drive the entire rise in rationing unemployment. This simply reflects that capital discounts alone drive the decline in capital in Figure 4. The right panel similarly decomposes frictional unemployment. I find that labour discounts drive 74%

of the initial rise in frictional unemployment. Recall that frictional unemployment is the residual unemployment which cannot be explained by rationing, and is therefore explained by hiring frictions. Rising labour discounts directly make hiring frictions worse, because they mechanically reduce the discounted *future* benefit of hiring ($q(\theta_t)J_t$) relative to the *current* cost of vacancy posting (κ). Rising capital discounts explain the downwards trend in frictional unemployment, as reduced investment reduces the marginal product of labour, reducing the fundamental desire to hire and making hiring frictions less relevant. However, capital discounts do also initially raise frictional unemployment. This is because in the presence of hiring frictions firms stop hiring at the beginning of the crisis in anticipation of future declines in investment and the MPL, as I discussed above.

Overall, my results show that search frictions (i.e. *frictional unemployment*) are important for understanding the transmission of discounts to unemployment in the early phases of the crisis. They explain half of the initial increase in discount-led unemployment, and transmit direct increases in labour market discounting and expected future decreases in capital investment into unemployment. This result contrasts with Michaillat's (2012) results in a model driven only by TFP shocks, where he finds that frictional unemployment must fall during a recession. Thus, discount rate shocks provide an important case where frictional unemployment remains quantitatively important even during recessions. On the other hand, the longer-term impact of discounts on unemployment is well explained by wage rigidity and a fundamental lack of investment (i.e. *rationing unemployment*). This transmits persistently elevated capital discounts into persistently higher unemployment.

4.3.4 Decomposition of employment, capital, and labour productivity

My final exercise is to investigate the quantitative relevance of discounts relative to other shocks during the Great Recession, paying particular attention to employment, capital, and labour productivity. As part of my estimation procedure I recover series for my shocks which can exactly replicate the data for key economic aggregates. In Figure 6 I use these shocks to decompose the changes in the these three variables during the Great Recession into contributions from my two discount measures, and the other shocks. In the interest of space, decompositions for other key variables are relegated to Figure 14 in the appendix.

The left panel plots the employment decomposition. The data is plotted in thin black, and each line gives counterfactual employment when simulating the economy subject only to the path of the given shock. In line with unemployment, employment initially sharply falls by over 9.5% before starting a recovery in 2010. Unlike unemployment, employment remains 4.5% below trend even at the end of the sample, largely associated with permanent declines in participation. As is to be expected given their importance in driving unemployment, discounts remain important in driving employment. After removing the decline in employment driven by exogenous participation (thin green line) the peak decline in employment is 5.6% in 2009M11. Combined, the initial peak reduction in employment due to both discounts is 2.6% in 2009M6, which is 46.5% of the participation-adjusted peak in the data.

Of the other shocks, the permanent decline in participation places the most downwards pressure on employment, and can alone fully explain the decline in employment by the end of the sample. Since utilisation adjusted TFP actually rises early in the crisis, it places upwards pressure on employment (dash-dotted yellow line) until later in the crisis when TFP falls permanently below trend. Finally, the unemployment value shock captures movements in employment not attributable to the other shocks in the model. It therefore causes a large decline in employment during the initial fall in employment, but causes a large *rise* in employment above trend during the recovery, in order to offset the drags on employment caused by capital discounts, TFP, and participation.

Figure 6: Full decomposition of employment, capital, and labour productivity

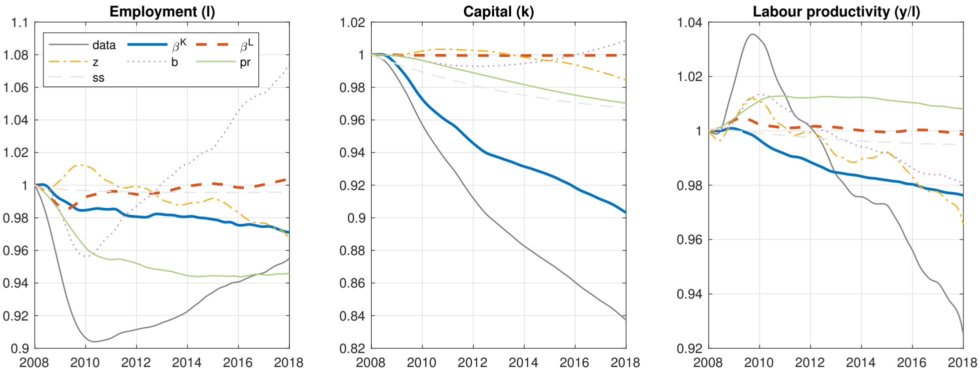

Each panel decomposes one variable during the Great Recession (data: thin black line) into counterfactual contributions from each shock. All variables are expressed as deviations from their 2008M1 values. Counterfactuals are given for capital discounts (thick blue), labour discounts (dashed red), TFP (dash-dotted yellow), unemployment value (dotted purple), participation rate (solid thin green), and drift of shocks and state variables back to steady state in the absence of shocks (dashed grey).

The right panel plots the labour productivity decomposition. Labour productivity initially sharply rises during the crisis, by 3.6% by 2009M10. As shown in Figure 2, a lot of this is driven by the utilisation adjustment of TFP, but even non-utilisation adjusted labour productivity rises by up to 1.2% by 2010. Following this, labour productivity starts a secular decline. It has fallen back to its pre-crisis value by 2012M3, and has fallen by 7.5% by the end of the sample. In terms of the decomposition, rising capital discounts consistently place downwards pressure on labour productivity, and labour discounts have limited effects on labour productivity, as previously discussed in Section 4.3.2. Accordingly, the remaining other shocks to the economy drive the bulk of the early increase in labour productivity, and contribute to the longer-term decline not explained by the capital discount shock.

In particular, increasing utilisation-adjusted TFP, the increased unemployment value shock, and declining labour market participation all put roughly equal upwards pressure on labour productivity during the early phase of the crisis. For TFP, this effect is mostly mechanical, as rising TFP increases output and hence labour productivity. For the other two shocks, the effects instead work through reducing employment, and hence raising labour productivity due to the diminishing marginal product of labour inherent in the aggregate production function. Declining participation reduces employment directly by shrinking the pool of workers, and the wage shock reduces employment by raising wages, lowering hiring, and hence lowering employment. Importantly, in the presence of capital adjustment costs these shocks do not quickly change capital, as shown in the centre panel, which means that they lead to a rise in the capital-labour ratio and hence labour productivity.

The decomposition of labour productivity illustrates an important conceptual point. Even though rising capital discounts raise unemployment by lowering labour productivity, their importance cannot simply by measured by inspecting observed labour productivity. In fact, capital discounts had increased unemployment by 1.8pp by 2012M3, even though labour productivity had been increasing during this whole period. The reasons for this are threefold. Firstly, as discussed above, capital discounts were putting downwards pressure on labour productivity during this period, but this effect was masked by rises in labour productivity coming from utilisation-adjusted TFP and the increased capital-labour ratio caused by the massive decline in employment itself. Secondly, the decline in employment caused by capital discounts themselves also offsets some of their negative effect on labour productivity: if employment had not adjusted in response to capital discounts the 9.7% total decline in capital they caused would lead labour productivity to decline by $0.097^{\alpha} = 3.2\%$. Since employment fell in response to the shock, counterfactual labour productivity only falls by 2.4%. Thirdly, in the presence of search frictions hiring is not a static decision. While observed labour productivity did not immediately fall during the crisis, it has since fallen by 7.5% relative to its pre-crisis trend, and the expectation of the future fall reduced hiring at the beginning of the crisis.

4.4 Further exercises and robustness

In the appendix I consider a wide array of additional and robustness exercises. In Appendix D.6 I provide an analysis of the role of discounts over the whole postwar sample and show that they are equally important for explaining the volatility of HP-filtered unemployment and its low correlation with labour productivity. In Appendix E, I show that my results are robust to using a non-linear (perfect foresight) solution method, changing the allocation of estimated discounts to Euler equations in the decompositions, reasonable variations in key parameters such as wage rigidity, replacement rates, and capital adjustment costs, adding job-separation shocks, using unadjusted TFP data, and accounting for changes in labour market participation.

5 Conclusion

In this paper I investigated the role of rising discount rates in driving unemployment during the Great Recession. I built a search and matching model with capital accumulation and time-varying discount rates for both capital and labour, and argued that the unprecedented decline in investment is an important contributor to unemployment during this period.

After incorporating capital, which fell by more than 15% during this period, I found that increased discount rates estimated from stock markets and investment can explain 52% of the peak rise in unemployment during the Great Recession. Importantly, I find that capital discounts raise unemployment by reducing labour productivity, even while observed labour productivity was actually rising early in the crisis. Thus, the importance of capital discounts for unemployment cannot be inferred from observed labour productivity alone, and must be estimated using capital data itself. These results speak to the importance of explicitly modelling capital in search and matching models, especially when trying to explain unemployment during financial recessions.

In this paper I have shown that one of the reasons that unemployment rises when discounts rise is that hiring requires the creation of costly, long-lived, capital for workers to produce with. A natural limitation of my approach is that I worked with a representative firm, while the root causes of high discounts would naturally have heterogeneous effects across firms. For example, tightened financial frictions are more likely to affect highly leveraged firms, who may be younger or smaller on average. Extending this analysis to a heterogeneous firm model is a natural next step. In addition, much job creation happens at young firms (Haltiwanger et al., 2013), suggesting an important role for the creation of not just capital, but also *entirely new firms*, in driving job creation and unemployment. Future work could investigate this idea further, by considering the impact of discount rates on unemployment through firm entry and growth in a full model of firm dynamics.

References

[1] Acemoglu, Daron, and Robert Shimer, 1999. *Efficient Unemployment Insurance.* Journal of Political Economy, 1999, vol. 107, issue 5, 893-928.

[2] Stphane Adjemian, Houtan Bastani, Michel Juillard, Frédéric Karamé, Junior Maih, Ferhat Mihoubi, George Perendia, Johannes Pfeifer, Marco Ratto and Sébastien Villemot, 2011. *Dynare: Reference Manual, Version 4.* Dynare Working Papers, 1, CEPREMAP.

[3] Andolfatto, David, 1996. *Business Cycles and Labor Market Search.* American Economic Review 86(1): 112-32

[4] Basu, Susanto, John G. Fernald and Miles S. Kimball, 2006. *Are Technology Improvements Contractionary?* American Economic Review, Vol. 96, NO. 5, December 2006 (pp. 1418-1448)

[5] Bernanke, Ben, and Mark Gertler, 1989. *Agency Costs, Net Worth, and Business Fluctuations.* The American Economic Review, Vol. 79, No. 1 (Mar., 1989), pp. 14-31

[6] Bernanke, Ben S, Gertler, Mark & Gilchrist, Simon, 1999. *The financial accelerator in a quantitative business cycle framework.* Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 21, pages 1341-1393 Elsevier.

[7] Brinca, P., V. V. Chari, Patrick J. Kehoe, Ellen R. McGrattan, 2016. *Accounting for Business Cycles.* Handbook of Macroeconomics, Volume 2, 2016, Pages 1013-1063

[8] Brügemann, Björn, Pieter Gautier, and Guido Menzio, 2019. *Intra Firm Bargaining and Shapley Values.* The Review of Economic Studies, Volume 86, Issue 2, March 2019, Pages 564-592.

[9] Campbell, John Y., and Robert J. Shiller, 1988. *The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors.* Review of Financial Studies 1 (3): 195-228.

[10] Carrillo-Tudela, Carlos, Michael Graberb, and Klaus Waelde, 2018. *Unemployment and vacancy dynamics with imperfect financial markets.* Labour Economics, Volume 50, March 2018, Pages 128-143.

[11] Chari, V. V., Patrick J. Kehoe, Ellen R. McGrattan, 2007. *Business Cycle Accounting.* Volume 75, Issue 3, May 2007, Pages 781-836.

[12] Christiano, Lawrence J, Trabandt, Mathias & Walentin, Karl, 2011. *Introducing financial frictions and unemployment into a small open economy model.* Journal of Economic Dynamics and Control, Elsevier, vol. 35(12), pages 1999-2041.

[13] den Haan, Wouter, J., Garey Ramey, and Joel Watson. 2000 *Job Destruction and Propagation of Shocks.* American Economic Review, 90 (3): 482-498.

[14] Elsby, Michael, Ryan Michaels and David Ratner, 2017. *Vacancy Chains.* 2017 Meeting Papers 888, Society for Economic Dynamics.

[15] Fama, Eugene F., and Kenneth R. French, 1988. *Dividend Yields and Expected Stock Returns.* Journal of Financial Economics 22 (1) 3-25.

[16] Farmer, Roger, 2012. *Confidence, Crashes and Animal Spirits.* Economic Journal, Royal Economic Society, vol. 122(559), pages 155-172, 03.

[17] Fernald, John, 2014. *A Quarterly, Utilization-Adjusted Series on Total Factor Productivity.* FRBSF Working Paper 2012-19.

[18] Gomme, Paul and Damba Lkhagvasuren, 2013. *Calibration and Simulation of DSGE Models.* Handbook of Research Methods and Applications in Empirical Methods in Macroeconomics, Nigar Nasimzade and Michael Thornton, eds. (Edward Elgar), pp. 575-592.

[19] Haefke, Christian, Marcus Sonntag, and Thijs van Rens, 2013. *Wage Rigidity and Job Creation.* Journal of Monetary Economics, 60(8).

[20] Hall, Robert E, 2016. *Macroeconomics of Persistent Slumps.* National Bureau of Economic Research Working Paper 22230.

[21] Hall, Robert, 2017. *High Discounts & High Unemployment.* American Economic Review 2017, 107(2): 305330.

[22] Hall, Robert E., and Paul R. Milgrom, 2008. *The Limited Influence of Unemployment on the Wage Bargain.* American Economic Review 98 (4): 1653-74.

[23] Haltiwanger, John, Ron Jarmin, and Javier Miranda, 2013. *Who Creates Jobs? Small versus Large versus Young.* The Review of Economics and Statistics, Vol. XCV(2), May 2013.

[24] Jermann, Urban, and Vincenzo Quadrini, 2012. *Macroeconomic Effects of Financial Shocks.* American Economic Review, Vol. 102, No. 1, Feb 2012 (pp. 238-71)

[25] Kehoe, Patrick, Viriliu Midrigan, and Elena Pastorino, forthcoming. *Debt Constraints and Employment.* Journal of Political Economy, forthcoming.

[26] Kiyotaki, Nobuhiro & Moore, John, 1997. *Credit Cycles.* Journal of Political Economy, University of Chicago Press, vol. 105(2), pages 211-48, April.

[27] Martellini, Paolo, Guido Menzio, and Ludo Visschers, 2019. *Revisiting the Hypothesis of High Discounts and High Unemployment.* IZA Discussion Paper No. 12441.

[28] Merz, Monika, 1995. *Search in the Labor Market and the Real Business Cycle.* Journal of Monetary Economics 36(2): 269-300

[29] Merz, Monika, and Eran Yashiv, 2007. *Labor and the Market Value of the Firm.* American Economic Review 97 (4): 1419-31.

[30] Michaillat, Pascal, 2012. *Do Matching Frictions Explain Unemployment? Not in Bad Times.* American Economic Review, American Economic Association, vol. 102(4), pages 1721-50, June.

[31] Mukoyama, Toshihiko, 2009. *A Note on Cyclical Discount Factors and Labor Market Volatility.* `https://sites.google.com/site/toshimukoyama/` (accessed December 21, 2018)

[32] Mumtaz, Haroon & Zanetti, Francesco , 2013. *The Effect of Labor and Financial Frictions on Aggregate Fluctuations.* Economics Series Working Papers 690, University of Oxford, Department of Economics.

[33] Petrongolo, Barbara & Christopher A. Pissarides, 2001. *Looking into the Black Box: A Survey of the Matching Function.* Journal of Economic Literature, American Economic Association, vol. 39(2), pages 390-431, June.

[34] Petrosky-Nadeau, Nicolas, 2014. *Credit, Vacancies and Unemployment Fluctuations.* Review of Economic Dynamics, Elsevier for the Society for Economic Dynamics, vol. 17(2), pages 191-205, April.

[35] Petrosky-Nadeau, Nicolas and Wasmer, Etienne, 2015. *Macroeconomic dynamics in a model of goods, labor, and credit market frictions.* Journal of Monetary Economics, Elsevier, vol. 72(C), pages 97-113.

[36] Pissarides, Christopher A., 2000. *Equilibrium Unemployment Theory, 2nd Edition.* MIT Press Books, The MIT Press, edition 1, volume 1, number 0262161877, June.

[37] Pissarides, Christopher A. 2009. *The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?* Econometrica 77 (5): 1339-69.

[38] Quadrini, Vincenso, and Sun, Qi, 2015. *Credit and Hiring.* Mimeo, available at `http://www-bcf.usc.edu/~quadrini/papers/UPpap.pdf`. Last accessed 30/03/2015.

[39] Rogerson, Richard, and Robert Shimer. 2010. *Search in Macroeconomic Models of the Labor Market.* Handbook of Labor Economics, Elsevier.

[40] Schoefer, Benjamin, 2015. *The Financial Channel of Wage Rigidity.* Mimeo, available at `http://scholar.harvard.edu/files/schoefer/files/jmp_bschoefer.pdf`. Last accessed 30/03/2015.

[41] Shimer, Robert, 2005. *The Cyclical Behavior of Equilibrium Unemployment and Vacancies.* American Economic Review, American Economic Association, vol. 95(1), pages 25-49, March.

[42] Shimer, Robert, 2012. *Reassessing the ins and outs of unemployment.* Review of Economic Dynamics. Volume 15, Issue 2, April 2012, Pages 127-148

[43] Wasmer, Etienne and Philippe Weil, 2004. *The Macroeconomics of Labor and Credit Market Imperfections.* American Economic Review, Vol. 94, No. 4, September 2004 (pp. 944-963)

[44] Winkler, Fabian, 2015. *The Role of Learning for Asset Prices, Business Cycles and Monetary Policy.* Mimeo, available at `http://personal.lse.ac.uk/winklerh/jmp.pdf`. Last accessed 27/03/2015.

[45] Yashiv, Eran, 2000. *Hiring as Investment Behavior.* Review of Economic Dynamics 3 (3): 486-522.

[46] Yashiv, Eran, 2016. *Capital values and job values.* Review of Economic Dynamics 19 (2016) 190209.

Appendices

A Appendix to Section 2

A.1 Capital price function derivation

A representative capital producing firm purchases the output good and converts it into capital at a cost, which it then sells to final goods firms. In particular, to create i_t units of new capital the capital producer must purchase i_t units of output and then pay an additional cost $\frac{\psi_k}{2}\left(\frac{i_t}{k_{t-1}}-\delta\right)^2 k_{t-1}$ denominated in the output good. This is a quadratic cost of deviating the investment rate away from the steady state investment rate, $i/k=\delta$, which is scaled by the size of the capital stock. ψ_k controls the degree of adjustment costs. The capital producing firm's maximisation problem can be stated as:

$$\pi = \max_{i_t} p_t^k i_t - i_t - \frac{\psi_k}{2}\left(\frac{i_t}{k_{t-1}}-\delta\right)^2 k_{t-1} \tag{24}$$

The first order condition for investment implies that the equilibrium capital price is linear in the investment rate:

$$p_t^k = 1 + \psi_k\left(\frac{i_t}{k_{t-1}}-\delta\right) \tag{25}$$

The general price function $p^k(i_t/k_{t-1})$ holds for generic cost functions of the form $f\left(\frac{i_t}{k_{t-1}}-\delta\right)k_{t-1}$.

A.2 Equivalence to single worker firm model

Pissarides (2000) shows how the basic single-worker-firm search model can be augmented to include capital. In this section, I show that the multi-worker-firm model is isomorphic to a single-worker-firm model with capital, conditional on firms taken wages as given.

In the single-worker-firm model, modified to match my notation, we can think of each match producing output $y_t = z_t\tilde{k}_{t-1}^{\alpha}$, where $\tilde{k}_{t-1}$ is capital per match, and total capital is thus $k_t = \tilde{k}_t l_t$. Matches rent capital flexibly each period at rental rate r_t^k, and each worker-firm match chooses the level of capital each period to maximise static output net of capital rental payments: $\pi(z_t, r_t^k) = \max_{\tilde{k}_{t-1}} z_t\tilde{k}_{t-1}^{\alpha} - r_t^k\tilde{k}_{t-1}$. Optimality gives $r_t^k = \alpha z_t\tilde{k}_{t-1}^{\alpha-1}$. Plugging this in to profirs yields $\pi(z_t, r_t^k) = (1-\alpha)z_t\tilde{k}_{t-1}^{\alpha} = (1-\alpha)y_t/l_{t-1}$. The optimal choice of capital thus implies that the static output net of capital rental per match is $(1-\alpha)y_t/l_{t-1}$, which is identical to the MPL in the multi worker firm model. The standard free entry condition then gives a condition identical to (6). The household's Euler equation for capital is $p_t^k = \beta_t^k\left(r_{t+1}^k + (1-\delta)p_{t+1}^k\right)$. Replacing r_{t+1}^k using $r_t^k = \alpha z_t\tilde{k}_{t-1}^{\alpha-1}$ and $k_t = \tilde{k}_t l_t$ yields (5).

A.3 Alternative definition of rationing unemployment

As discussed in the text, the extension to a model with capital reveals a subtlety in how rationing unemployment should be defined, since capital is itself an endogenous object. Should rationing unemployment be simply the level of unemployment in the absence of hiring frictions given the existing level of TFP and capital stock today? Or should we recognise that removing hiring costs would lead to different investment choices and hence a potentially different level of capital today? I opt to take the first definition in the main text. In this section, I briefly outline the alternative definition of rationing unemployment, show that it is well defined, but explain why it is harder to use outside of steady state.

Intuitively, the alternative definition of rationing unemployment is the equilibrium level of unemployment in the model when matching frictions are removed ($\kappa = 0$) and capital is allowed to adjust to the level firms would like to set given that matching frictions have been removed. In practice, this is more complicated than it sounds, since it is not not clear how we are letting capital adjust. Capital adjustment is a dynamic process subject to adjustment costs, so the "optimal level of capital at time t in the absence of hiring costs" will naturally depend on the level of capital last period, and the optimal level of investment from the capital Euler equation (5).

So, one sensible alternative definition is to suppose that there had never been hiring costs for all of time (since $t = 0$), and that capital has evolved optimally according to these shocks. Rationing unemployment is the level of unemployment generated by current TFP and this level of capital. This makes clear that one needs to define a whole new dynamic model in order to define this alternative notion of rationing unemployment.

Specifically, consider the model defined by the following equations. Capital evolution:

$$k_t = i_t + (1 - \delta)k_{t-1} \tag{26}$$

Euler equation for capital:

$$p_t^k = \beta_t^k \mathrm{E}_t \left[\alpha z_{t+1} k_t^{\alpha-1} l_t^{1-\alpha} + p_{t+1}^k (1 - \delta) \right] \tag{27}$$

where $p_t^k = p^k(i_t/k_{t-1})$. Exogenous wage rule:

$$w_t = \omega \left(z_t k_{t-1}^{\alpha} p_t^{-\alpha} \right)^{\gamma} \tag{28}$$

Unemployment definition:

$$u_t = p_t - l_{t-1} \tag{29}$$

Finally, static labour demand in the absence of hiring costs is expressed as an inequality constraint to allow for firms to be rationed in the case where labour demand exceeds the population. That is, we have the labour demand condition,

$$w_t \geq (1 - \alpha) z_t k_{t-1}^{\alpha} l_{t-1}^{-\alpha} \tag{30}$$

and non-negative unemployment, $u_t \geq 0$, with complementary slackness: $(w_t - (1-\alpha) z_t k_{t-1}^{\alpha} l_{t-1}^{-\alpha}) u_t = 0$.

This is the full model from the main text, but with hiring costs removed ($\kappa = 0$) and wage setting replaced with the exogenous wage rule. The model is closed with stochastic processes for TFP, participation, and capital discounts, and a transversality condition for capital. The equilibrium of the model is defined in the usual way, and gives sequences for the endogenous variables as functions of initial capital and a realised sequence of shocks.

Definition 4. *Fix an initial capital stock k_{-1}, and sequence of shocks $s^t \equiv \{z_s, \beta_s^k\}_{s=0}^t$. Rationing unemployment at time t given this shock sequence is defined as the solution to unemployment in the model defined by equations (26) to (30): $u_t^R \equiv u_t(s^t)$.*

Relative to the definition of rationing unemployment used in the text, (17), there are a few differences to note. Firstly, rationing unemployment here depends not only on the current TFP shock, z_t, participation shock, p_t, and observed level of capital, k_{t-1}, but instead the whole history of TFP shocks up to time t. This is because the counterfactual level of capital in the absence of hiring costs, which affects the MPL and hence level of unemployment, is chosen in response to the whole history of shocks. Secondly, note that this definition of rationing unemployment also therefore depends on the history of capital discount shocks, since they also affect the counterfactual capital level.

In practical terms, this means that calculating this alternative measure of rationing unemployment requires solving the equilibrium policy functions of this dynamic model. This is feasible, since the model admits a recursive representation in the states $(k_{t-1}, z_t, p_t, \beta_t^k)$. However, it is somewhat more involved than the simpler definition (17) used in the text, which is static and has a simple analytical solution, which is why I use it as my main definition.

In order to demonstrate that this alternative definition of rationing unemployment still implies that positive rationing unemployment can exist despite allowing capital to adjust, I now show how to calculate it in steady state. The definition simplifies considerably in steady state, since now we have an analytical solution for capital using the capital Euler equation in steady state.

In particular, taking (27) in steady state gives the solution for the capital-labour ratio $\frac{k}{l} = \mathcal{K}(r)$ given in (21), where r is the level of the steady state capital discount.

Assume participation is $p = 1$ for simplicity. First suppose positive unemployment, and combine $\mathcal{K}(r)$ with the labour demand condition, (30), and wage rule, (28) to solve for employment:

$$(1-\alpha) z \mathcal{K}(r)^{\alpha} = \omega z^{\gamma} (\mathcal{K}(r) l)^{\alpha\gamma} \implies l = \left(\frac{1-\alpha}{\omega}\right)^{\frac{1}{\alpha\gamma}} \left(z^{\frac{1}{\alpha}} \mathcal{K}(r)\right)^{\frac{1-\gamma}{\gamma}} \tag{31}$$

Finally, using $u = 1 - l$ and incorporating the possibility of zero unemployment gives the alternative definition of rationing unemployment in steady state:

$$u^R = \max\left\{0, 1 - \left(\frac{1-\alpha}{\omega}\right)^{\frac{1}{\alpha\gamma}} \left(z^{\frac{1}{\alpha}} \mathcal{K}(r)\right)^{\frac{1-\gamma}{\gamma}}\right\} \tag{32}$$

This equation makes it clear that rationing unemployment can still be positive in this alternative definition. To see this, note that the expression must be positive for low enough z or $\mathcal{K}(r)$, as long as wages are sticky ($\gamma < 1$). Rationing unemployment by this definition will typically be lower than the definition in the main text, since counterfactual capital will rise when hiring costs are removed, increasing the MPL and incentives to hire.

B Appendix to Section 3

Derivation of $\hat{q}(l)$ function: In steady state, we can write the job filling rate as a strictly decreasing function of equilibrium employment: $q = \hat{q}(l)$ with $\hat{q}'(l) < 0$. Taking (1) in steady state gives steady state employment as a function of tightness: $l = \frac{\lambda(\theta)}{\lambda(\theta)+\rho}$. Since I assumed that $\lambda(\theta)$ was strictly increasing, l is a strictly increasing function of θ, which can be inverted to give steady-state θ as a strictly increasing function of steady-state l: $\theta = \hat{\theta}(l)$ with $\hat{\theta}'(l) > 0$. Since $q'(\theta) < 0$ we have that $\hat{q}(l) \equiv q(\hat{\theta}(l))$ is decreasing in l.

B.1 Proofs

Proof of Proposition 1: To prove that total unemployment is increasing in discounts, implicitly differentiating (19) gives:

$$l^{s\prime}(r) = \frac{-\kappa/\hat{q}(l^s)}{\alpha(1-\alpha)z(k^*)^\alpha l^s(r)^{-\alpha-1} - (r+\rho)\kappa\frac{\hat{q}'(l^s)}{\hat{q}(l^s)^2}} < 0 \tag{33}$$

which is negative because $\hat{q}'(l) < 0$. Since $u^s = 1 - l^s$ we thus have that

$$u^{s\prime}(r) = -l^{s\prime}(r) > 0 \tag{34}$$

□

Proof of Proposition 2: To prove that total unemployment is increasing in discounts, differentiating (22), gives

$$l^{l\prime}(r) = \frac{-\frac{\kappa}{\hat{q}(l^l)} + \frac{\alpha\mathcal{K}'(r)}{\mathcal{K}(r)}\left((1-\alpha)z\mathcal{K}(r)^\alpha - \gamma\omega z^\gamma(\mathcal{K}(r)l^l)^{\alpha\gamma}\right)}{\alpha\gamma\omega z^\gamma\mathcal{K}(r)(\mathcal{K}(r)l^l)^{\alpha\gamma-1} - (r+\rho)\kappa\frac{\hat{q}'(l^l)}{\hat{q}(l^l)^2}} < 0. \tag{35}$$

To show that this derivative is negative, note that the denominator is positive because $\hat{q}'(l^l) < 0$. The numerator is negative because 1) $\mathcal{K}'(r) < 0$ and 2) the term in brackets is positive if employment is positive in equilibrium.[19] This shows that employment falls as discounts rise. The definition of unemployment then trivially shows that unemployment rises as discounts rise:

$$u^{l\prime}(r) = -l^{l\prime}(r) > 0. \tag{36}$$

To show that rationing unemployment is increasing in discounts, differentiate (23) when $u^{R,l}(r) > 0$ to yield

$$u^{R,l\prime}(r) = -(1-\gamma)\left(\frac{(1-\alpha)z^{1-\gamma}}{\omega}\right)^{\frac{1}{\alpha}}\left(\mathcal{K}'(r)l^l(r) + \mathcal{K}(r)l^{l\prime}(r)\right)\left(\mathcal{K}(r)l^l(r)\right)^{-\gamma} > 0. \tag{37}$$

This is positive because $\mathcal{K}'(r) < 0$ and $l'(r) < 0$, which simply means that total capital (not just capital intensity) falls when discounts rise.

To show that the response of frictional unemployment is ambiguous we can compare two limit cases. As the first limit case, consider the case of fully flexible wages: $\gamma = 1$. Inspecting (23) shows that in this case rationing unemployment is a function only of parameters and is hence unrelated to discounts: $u^{R,l} = \max\left\{0, 1 - \left(\frac{(1-\alpha)}{\omega}\right)^{\frac{1}{\alpha}}\right\}$. Consider a calibration where $\omega > 1-\alpha$, so that rationing unemployment is then positive.[20] Since unemployment increases when discounts rise even when wages are fully flexible (from (35) we still have $l^{l\prime}(r) < 0$ when $\gamma = 1$) the entire increase in

[19] Under the maintained assumption of interior employment, we have that the firm's surplus from posting vacancies is positive: $\frac{\kappa}{\hat{q}(l^l)} = \frac{(1-\alpha)z\mathcal{K}(r)^\alpha - \omega z^\gamma(\mathcal{K}(r)l^l)^{\alpha\gamma}}{r+\rho} > 0$. Since $r+\rho > 0$, this implies that $(1-\alpha)z\mathcal{K}(r)^\alpha - \omega z^\gamma(\mathcal{K}(r)l^l)^{\alpha\gamma} > 0$. Since $\gamma \leq 1$, this also implies that $(1-\alpha)z\mathcal{K}(r)^\alpha - \gamma\omega z^\gamma(\mathcal{K}(r)l^l)^{\alpha\gamma} > 0$, as required.

[20] Recalling the wage rule (15), a high value of ω corresponds to a high wage intercept and hence high average wage. For $\omega > 1-\alpha$ the wage is high enough that unemployment would be positive even in the absence of hiring frictions: i.e. rationing unemployment is positive.

unemployment must be frictional. Intuitively, when wages are very flexible there is limited scope for rationing unemployment, which derives from wage stickiness, and so any increase in unemployment must be frictional.

As the second limit case, consider fully sticky sticky wages: $\gamma = 0$. Continue to consider parameters such that rationing unemployment is positive. In this case, plugging (23) and $l^l \equiv 1 - u^l$ into the definition of frictional unemployment, $u^{F,l\prime}(r) = u^{l\prime}(r) - u^{R,l\prime}(r)$, gives

$$u^{F,l}(r)\Big|_{\gamma=0} = l^l(r)\left[\left(\frac{(1-\alpha)z}{\omega}\right)^{\frac{1}{\alpha}}\mathcal{K}(r) - 1\right]. \tag{38}$$

Under the maintained assumption that employment is positive, frictional unemployment must also be positive in this case.[21] This expression reveals that frictional unemployment declines as discounts rise ($u^{F\prime}(r) < 0$) because both $l^{l\prime}(r) < 0$ and $\mathcal{K}'(r) < 0$. Intuitively, when wages are very sticky, rationing unemployment is very powerful. For a fully fixed wage, any fall in capital, and associated fall in MPL, will translate into a large rise in rationing unemployment. Frictional unemployment falls because, as discussed in Michaillat (2012), search frictions become consequently less important in driving unemployment when unemployment is high. This is because when unemployment is high the vacancy filling rate is high, and it is less costly for firms to fill vacancies, rendering search frictions less important. □

Proof of Proposition 3: Recall that in the initial steady state capital was optimally chosen from the Euler equation, meaning that we have $k^* = \mathcal{K}(r^*)l^*$. Additionally note that when $r = r^*$ we have $l^s(r^*) = l^l(r^*) = l^*$. Plugging these into (33) and (35) gives:

$$l^{s\prime}(r^*) = \frac{-\kappa/\hat{q}(l^*)}{\alpha(1-\alpha)z\mathcal{K}(r^*)^\alpha(l^*)^{-1} - (r^*+\rho)\kappa\frac{\hat{q}'(l^*)}{\hat{q}(l^*)^2}} \tag{39}$$

$$l^{l\prime}(r^*) = \frac{-\frac{\kappa}{\hat{q}(l^*)} + \frac{\alpha\mathcal{K}'(r^*)}{\mathcal{K}(r^*)}\left((1-\alpha)z\mathcal{K}(r^*)^\alpha - \gamma\omega z^\gamma(\mathcal{K}(r^*)l^*)^{\alpha\gamma}\right)}{\alpha\gamma\omega z^\gamma\mathcal{K}(r^*)^{\alpha\gamma}(l^*)^{\alpha\gamma-1} - (r^*+\rho)\kappa\frac{\hat{q}'(l^*)}{\hat{q}(l^*)^2}} \tag{40}$$

Comparing (39) and (40) shows that the fall in employment following an increase in discounts is larger in the long run. This is because 1) in the numerator, (40) contains an extra negative term, and 2) the denominator of (40) is smaller.

Both of these follow from the maintained assumption that unemployment is positive. Under the maintained assumption of interior employment, we have that the firm's surplus from posting vacancies is positive: $\frac{\kappa}{\hat{q}(l^*)} = \frac{(1-\alpha)z\mathcal{K}(r^*)^\alpha - \omega z^\gamma(\mathcal{K}(r^*)l^*)^{\alpha\gamma}}{r^*+\rho} > 0$. Since $r^* + \rho > 0$, this implies that $(1-\alpha)z\mathcal{K}(r^*)^\alpha - \omega z^\gamma(\mathcal{K}(r^*)l^*)^{\alpha\gamma} > 0$. Since $\gamma \leq 1$, this also implies that $(1-\alpha)z\mathcal{K}(r^*)^\alpha - \gamma\omega z^\gamma(\mathcal{K}(r^*)l^*)^{\alpha\gamma} > 0$, as required for point (1). Multiplying both sides by $\alpha(l^*)^{-1}$ gives $\alpha(1-\alpha)z\mathcal{K}(r^*)^\alpha(l^*)^{-1} > \alpha\gamma\omega z^\gamma(\mathcal{K}(r^*))^{\alpha\gamma}(l^*)^{\alpha\gamma-1}$, as required for point (2). □

[21] Under the maintained assumption of interior employment, now with $\gamma = 0$, we have that the firm's surplus from posting vacancies is positive: $\frac{\kappa}{\hat{q}(l^l)} = \frac{(1-\alpha)z\mathcal{K}(r)^\alpha - \omega}{r+\rho} > 0$. Since $r + \rho > 0$, this implies that $(1-\alpha)z\mathcal{K}(r)^\alpha > \omega$. This implies that the bracket in (38) is positive. Since employment is also positive, $l^l(r) > 0$, (38) show that frictional unemployment must be positive.

B.2 Separate discounts

In this section I analytically analyse the case of separate discounts. I focus on the long run with endogenous capital, where the distinction is meaningful, and allow the discount rates on capital, $r_K \equiv \beta_K^{-1} - 1$, and labour, $r_L \equiv \beta_L^{-1} - 1$, to differ.

In this case, the equilibrium is described by the following modified versions of (21) and (22):

$$\frac{k}{l} = \left(\frac{\alpha z}{r_K + \delta}\right)^{\frac{1}{1-\alpha}} \equiv \mathcal{K}(r_K) \tag{41}$$

$$\frac{\kappa}{q(l^l)} = \frac{(1-\alpha)z\mathcal{K}(r_K)^\alpha - \omega z^\gamma (\mathcal{K}(r_K)l^l)^{\alpha\gamma}}{r_L + \rho}. \tag{42}$$

Now, the discount on labour only appears in the denominator of the job value in (42). The discount on capital only directly affects the capital-labour ratio via (41), and then indirectly affects the job value. Following the arguments of Section 3, an increase in either discount rate will reduce employment and hence increase unemployment. But how does the transmission vary for the two discounts? Differentiating (42) with respect to each discount shows how employment responds to each:

$$l^l_{r_L} = \frac{-\frac{\kappa}{q(l^l)}}{\alpha\gamma\omega z^\gamma \mathcal{K}(r_K)(\mathcal{K}(r_K)l^l)^{\alpha\gamma-1} - \frac{q'(l^l)\kappa(r_L+\rho)}{q(l^l)^2}} \tag{43}$$

$$l^l_{r_K} = \frac{\frac{\alpha\mathcal{K}'(r_K)}{\mathcal{K}(r_K)}\left((1-\alpha)z\mathcal{K}(r_K)^\alpha - \gamma\omega z^\gamma(\mathcal{K}(r_K)l^l)^{\alpha\gamma}\right)}{\alpha\gamma\omega z^\gamma \mathcal{K}(r_K)(\mathcal{K}(r_K)l^l)^{\alpha\gamma-1} - \frac{q'(l^l)\kappa(r_L+\rho)}{q(l^l)^2}}. \tag{44}$$

As should be expected, these two responses sum up to the total response of employment in the model with a common discount given in (35). This reveals how each discount operates on total employment and hence on unemployment. The denominators are common, and each derivative differs only in the numerator. The magnitude of the effect of the labour discount is controlled by search frictions, via the term $-\frac{\kappa}{q(l^l)}$. The magnitude of the effect of the capital discount is controlled by the responsiveness of the capital-labour ratio to the capital discount, $\mathcal{K}'(r_K)$, moderated by the degree of wage rigidity, γ.

How do the two frictions affect rationing and frictional unemployment? It turns out that rationing unemployment depends on both the capital and labour discounts:

$$u^{R,l}(r_K, r_L) = \max\left\{0, 1 - \left(\frac{(1-\alpha)z^{1-\gamma}\left(\mathcal{K}(r_K)l^l(r_K,r_L)\right)^{\alpha(1-\gamma)}}{\omega}\right)^{\frac{1}{\alpha}}\right\}. \tag{45}$$

Thus, there is not an exact separation between the capital and labour discounts in affecting frictional and rationing unemployment, as one might expect. Intuitively, this is because rationing unemployment is defined as the level of unemployment which would prevail in the absence of search frictions, but for the current level of capital. In the long run, the level of capital is an endogenous object which depends on both discounts: $k(r_K, r_L) = \mathcal{K}(r_K)l^l(r_K, r_L)$. The capital discount affects total capital by changing the capital-labour ratio, while the labour discount affects the level of capital by changing the level of labour, and hence indirectly the level of capital for a given capital-labour ratio. Additionally, the capital discount will affect frictional unemployment since this is defined as the residual level of unemployment net of rationing unemployment.

C Data appendix

C.1 Data sources

The national accounts data is quarterly, while labour market and stock market data is available at a monthly frequency. I solve my model at a monthly frequency, and linearly interpolate national accounts data to create monthly data. For the national accounts I take non-seasonally adjusted data, and all data is then seasonally adjusted using a five-quarter moving average filter. Monthly data is seasonally adjusted with a 13-month moving average filter. Data is converted to per-capita terms using population data from the Bureau of Labor Statistics (BLS).

National accounts data are taken from the database maintained by Fernald (2014). I use data from 1950Q1 to 2018Q1. This data has the advantage of having a utilisation adjustment measure, based on the Basu et al. (2006) methodology. The data covers the "Business sector" of the US economy, and accordingly excludes the government and non-profit sectors. I construct indices for output, y_t, capital, k_t, and factor utilisation from this dataset by cumulating forward the quarterly growth rates. Investment is calculated from the capital stock using my assumed depreciation rate.

For employment, l_t, in order to match the national accounts data, I must use data on business-sector employment. These are constructed from the "Employment by major industry sector" files from the Bureau of Labor Statistics (BLS) by adding Farm and Non-farm Business sector employment. I next construct the model-relevant notion of TFP, which is defined over capital and employment (not hours, since my model assumes constant hours per worker). To do this, I extract TFP as the residual of a Cobb Douglas production function in capital and employment with the capital share α taken from my calibration. Utilisation adjusted TFP, z_t^u, is calculated by dividing TFP by the utilisation index. Since I use utilisation adjusted TFP as my data when estimating the model, in order for the data inputted into the model to be mutually coherent when using adjusted TFP, I need to adjust one of output, employment, or capital in order for the production function to hold. I choose to adjust output, and recalculate it using actual employment and capital and the adjusted TFP series. All series are per capita, and then detrended using their average growth rates from 1950Q1 to 2008Q1.

I take the monthly Civilian Unemployment Rate from the BLS to be my measure of the unemployment rate. One complication arises when bringing the model to the data. During the Great Recession there were large changes in labour market participation, which make analysing the labour market through just the lens of unemployment problematic. To complicate matters further, the change in participation consists of both components related to demographics (population ageing) and also components plausibly endogenous to the business cycle (declines in participation within age groups). In my model both total employment and the unemployment rate are important variables. The former affects, for example, the capital labour ratio, and the latter is my key object of interest. Given that participation changes in the data, it is not possible to exactly match data on both employment and unemployment simultaneously in a model with constant participation. Accordingly, I back out a measure of the labour market participation rate as the rate implied by my employment and unemployment data, as explained in Appendix C.2.1.[22]

[22]Constructing measures of unemployment for any model requires taking a stand on the nature of participation. I show in the appendix that all of the results are robust to alternative definitions of unemployment. I perform the

While I do not use any time series of real wage data in the calibration or decomposition of the model, I do construct two measures of real wages to compare to my model-implied wage. My first measure comes directly from the national accounts and covers my whole sample. I take total nominal labour compensation for the business sector from the BLS' MSPC tables. I then construct a measure of wages as real earnings per employee by dividing by my measure of employment and deflating with the GDP deflator. My second measure uses micro-data from the Atlanta Fed's wage tracker based on data from the Current Population Survey, but is only available from 1997M10. The data gives the median 12 month growth rate of nominal hourly wages across sampled workers, and I take the data for all workers. These growth rates are cumulated forward to form a wage index, which I convert to real earnings per employee by multiplying by total hours and dividing by employment, and deflating with the GDP deflator. I discuss the model-implied wage data further in Appendix D.4.

Finally, for my stock market data I take monthly data for the S&P500 portfolio from Robert Shiller's website (`http://www.econ.yale.edu/~shiller/`). I use his real price and dividend series, detrend them both, and remove seasonality with a 13-month moving average filter. In order to preserve a key feature of the data – the price-dividend ratio – post-detrending, I normalise the average pre-crisis stock price to one, and normalise dividends so that the average price-dividend ratio is equal to its data average of 407.5 (or $407.5/12 \simeq 34$ in annual terms). Data for the whole sample are plotted in Figure 1, and for the Great Recession in Figure 3. After combining my monthly and quarterly data I have a monthly dataset running from 1950M1 to 2018M1.

C.2 Construction of participation rate and unemployment

In this section I give the details of how the participation rate and unemployment are constructed from the data.

C.2.1 Baseline approach

In my baseline calibration, I back out the level of labour market participation required to exactly match my data on employment and the unemployment rate. That is, I use data on l_{t-1} and $\tilde{u}_t$, and back out $u_t = \tilde{u}_t l_{t-1}/(1 - \tilde{u}_t)$ and $p_t = u_t + l_{t-1}$. If I was using data on employment for the whole economy, this would trivially recover the exact participation rate from the data. Since I use employment only in the business sector, this back out a measure of the participation rate within the business sector.

C.2.2 Robustness: alternative participation adjustment

As a robustness exercise, I also consider the following alternative measure of participation. The approach I follow here is to construct a counterfactual measure of labour market participation related only to changing demographics. I construct this measure using data on participation rates by age from the BLS. I hold participation within each age group at its 2008 level and simulate the change in the demographic makeup of the workforce. This series measures what aggregate participation would be if workers of each age participated as much as they did in 2008, loosely as if the Great Recession "had not happened".

same exercise using demographic data to strip out changes in participation related to demographics. I then back out the model-implied unemployment rate, and use this instead of unemployment data. This changes the measurement of unemployment over the Great Recession, but the role of discounts is quantitatively almost identical.

The BLS provides data on labour force participation both for the whole population, and for different age groups. Denote the whole economy participation rate by p_t^{raw}. Let lf_t denote the total labour force, pop_t the population, and index age groups by i. Then by construction we can decompose total participation to

$$p_t^{raw} = \frac{lf_t}{pop_t} = \frac{\sum_i lf_{i,t}}{pop_t} = \sum_i p_{i,t} s_{i,t} \tag{46}$$

with $s_{i,t} = pop_{i,t}/pop_t$ being the population shares of each age group, and $p_{i,t} = \frac{lf_{i,t}}{pop_{i,t}}$ the participation rates of each group. I construct a counterfactual overall participation rate, which gives the level of participation predicted by demographic change only, assuming that the participation rate within every age group is permanently fixed at its 2008 level, $p_{i,2008}$. Denote this counterfactual series by p_t^{cf}. It is given by

$$p_t^{cf} = \sum_i p_{i,2008} s_{i,t}. \tag{47}$$

Note that $p_{2008}^{cf} = p_{2008}^{raw}$ by construction. For this exercise I use the most detailed population breakdown available in the data. I take the whole population to be all those 16 years and older. The population groups are: 16 to 19, 20 to 24, 25 to 29, 30 to 34, 35 to 39, 40 to 44, 45 to 49, 50 to 54, 55 to 59, 60 to 64, 65 to 69, 70 to 74, and 75 plus. The raw and counterfactual participation rates are plotted for the whole sample in Figure 7.

Figure 7: Counterfactual participation rate

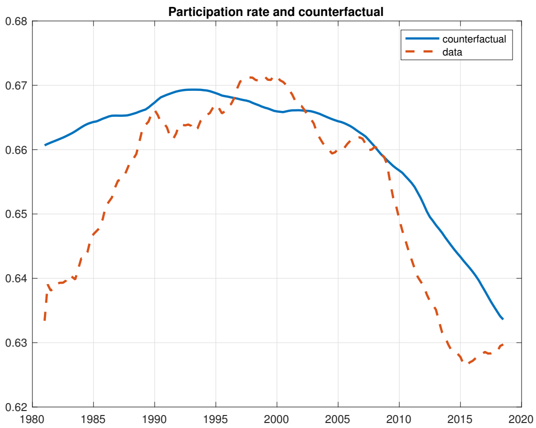

Figure compares the actual participation rate in the data (red line) with the counterfactual rate coming just from demographic changes (blue line). See text for details.

The two rates coincide by construction in 2008, when participation is 66%. By 2018 the participation rate had fallen to 63%, but amazingly the counterfactual rate due just to demographics also falls to 63.5%. Hence, by the end of the sample the majority of the aggregate decline in participation is caused simply by demographics.

However, during the middle of the crisis the actual participation rate, p_t^{raw}, falls by more than the counterfactual rate, p_t^{cf}. I interpret this additional decline as involuntary non-participation by people who want to work but drop out of the labour force during the recession. Since I only have a two-state model of the labour market, I cannot analyse participation decisions directly. Hence, in order to recognise the endogenous nature of participation, I choose to add this extra decline in participation to the definition of unemployment. In terms of the model, this means that some of the workers who are non-participants in the data will be added to the unemployment pool, and will thus be allowed to actively search for jobs.

In the exercise above I constructed a measure of counterfactual participation for the whole economy from the BLS data. Since the rest of my data is for the Business Sector only, I finally need to construct a measure of participation for just the Business Sector. Unfortunately, the data is only given for the whole economy. In the absence of direct data on this, I back out this series under two assumptions. Firstly, I assume that unemployment in the Business Sector was at the economy-wide level, 5%, just before the crisis in 2007M11. Secondly, I assume that from then onwards participation in the Business Sector evolves proportionately to participation in the whole economy.

For a participation series p_t, we have $u_t + l_{t-1} = p_t \Rightarrow u_t = p_t - l_{t-1}$. The first assumption allows me to back out participation in 2007M11 as

$$p_{2007M11} = u_{2007M11} + l_{2007M10}, \tag{48}$$

where $u_{2007M11} = 0.05 l_{2007M10}/(1 - 0.05)$. The second assumption then allows me to construct participation for the rest of the sample as $p_t = p_{2007M11} \times (p_t^{cf}/p_{2007M11}^{cf})$. Finally, unemployment for the rest of the sample is given by $u_t = p_t - l_{t-1}$. The results when using this alternative participation and unemployment definition are discussed in Section E.2 and plotted in Figure 23.

D Appendix to Section 4

D.1 Calibration and estimation details

Steady state calibration: The model is calibrated at a monthly frequency. I choose steady-state productivity, z, to normalise output to $y = 1$. I choose α to match a labour share of income of $wl/y = 2/3$. Since firms extract extra surplus due to the matching frictions, this corresponds to a value of $\alpha = 0.3213$.

On the investment side, I choose values standard to the Real Business Cycle literature, chosen to match postwar US data (see Gomme and Lkhagvasuren, 2013, for a recent discussion). I match an investment to output ratio of $i/y = 0.18$, which gives $i = 0.18y = 0.18$. I take depreciation to be 6.5% at an annual frequency, giving $\delta = 1 - (1 - 0.065)^{1/12} = 0.0056$ at a monthly frequency. In steady state, investment must equal depreciated capital, giving $k = i/\delta = 0.18/\delta$. I choose the steady-state capital discount rate in order to match this steady-state level of capital. Since the price of capital is one in steady state ($p^k(\delta) = 1$), the steady state return on capital is $r^k = \alpha y/k + 1 - \delta = 0.0044$, which equates to a 5.37% yearly return on capital. I thus set $\beta^k = 1/1.0044 = 0.9956$. Following Bernanke et al. (1999) and Brinca et al. (2016) I choose the adjustment cost parameter ψ_k to match an elasticity of the price of capital to the investment rate of 0.25 in the steady state. This requires setting $\psi_k = 0.25/\delta = 44.762$.[23] I also perform robustness allowing for higher and lower adjustment

[23]The scale of this parameter reflects the definition of the length of a period. Since my periods are only a month long, and investment is measured as a flow while capital is a stock, the investment-capital ratio is lower than a model with a quarterly calibration. If the model was calibrated quarterly the parameter would instead be $\psi_k = 0.25/\delta_q = 15.0$ where $\delta_q = 1 - (1 - 0.065)^{1/4}$, and the calibrated elasticity would remain identical.

costs.

For the stock market, I normalise steady-state stock prices to one, giving $p = 1$. I target a steady state price-dividend ratio of 407.5, equal to the pre-crisis average of my stock market data, implying steady state dividends of $p/d = 407.5 \Rightarrow d = 0.0025$. The steady-state labour discount is measured from the stock market, and satisfies $p = \beta^l(p + d)$, giving $\beta^l = 0.9976$. These discount rates imply a steady state annualised return on the stock market and capital of 3.0% and 5.4% respectively.[24]

The labour market is parameterised to stay close to Hall's (2017) strategy. I take the job separation rate to be $\rho = 0.0345$, following Shimer (2005). I target an unemployment rate of $\tilde{u} = 0.05$. Since the participation rate is constant in steady state, matching steady state flows in and out of unemployment implies a monthly job finding rate of $\lambda_w = (1 - \tilde{u})\rho/\tilde{u} = 0.6555$. I assume a Cobb-Douglas matching function, $m_t = \psi_0 u_t^{\psi_1} v_t^{1-\psi_1}$, with standard elasticity $\psi_1 = 0.5$ (Petrongolo and Pissarides, 2001). Following Shimer (2005), this allows me to normalise steady state tightness to $\theta^* = 1$ and I pick the matching efficiency parameter to match the required job finding rate: $\psi_0 = \lambda_w \theta^{\psi_1 - 1} = 0.6555$. The steady state participation rate is calibrated to match the average employment rate in Fernald's (2014) data between 1950M1 and 2008M1. Recalling that this is employment in the business sector only, this gives an employment rate of 51.5% and participation rate of $p = 0.5421$.

I set the steady state real wage, w, following Michaillat (2012). Based on empirical estimates, he requires that the steady state flow recruiting cost, κ, is equal to 0.32 of a workers steady state wage: $\kappa = 0.32w$. This allows me to solve for w and κ from the Euler equation for labour, (6), given a value of the steady-state discount rate and the targeted rate of unemployment. This gives values $w = 1.2946$ and $\kappa = 0.4143$. In steady state, this means that the wage is 98.2% of the marginal product of labour. This real wage must be supported as the equilibrium of the wage bargaining game. I assume a 40% replacement rate of unemployment benefits and set the steady state level of the worker's value of unemployment to $b = 0.4w = 0.5178$. I calibrate the level of the firm's disruption cost to $\chi = 0.5674$ (meaning that delay costs the firm 57% of marginal product) so that this wage satisfies (14) in steady state, for the value of ψ estimated in the next section.

Simulation and estimation of remaining parameters and shock paths: The remaining parameters are estimated using an iterative simulation procedure. I guess parameter values, simulate the model, re-estimate parameters from simulated data, and repeat until convergence. The remaining parameters to be estimated are the autocorrelation and variances of my six shocks, and the parameter ψ which controls the degree of wage rigidity, making 13 parameters in total. For the simulation used in estimation and my baseline results, the model is solved using first order perturbation techniques.[25] I perform robustness using a nonlinear perfect-foresight solution in Appendix E.

The estimation exploits the invertibility of the model, and backs out the required paths for the six underlying shocks shocks in order to exactly match the historical data for six key macroeconomic aggregates: the unemployment rate, $\tilde{u}_t$, output, y_t, capital, k_t, employment, l_t, stock prices, p_t^s, and dividends, d_t. Post-estimation, these shock series are then used to assess the importance of each shock at driving unemployment during the Great Recession.

Estimation is made easier since parameters can be estimated in three separate blocks. Firstly,

[24] The return on the stock market is lower than that measured in the data due to detrending, since the model does not account for the average growth rate of dividends. The returns should be interpreted as returns net of the average growth rate of dividends and TFP respectively.

[25] Linearisation is performed using the Dynare toolbox (Adjemian et al., 2011) and model simulation and estimation are carried out manually using Matlab.

since TFP, the participation rate, and dividends are observed in the data, the stochastic processes for these variables can be directly estimated on the data. I estimate the parameters of the log-AR(1) processes for each parameter separately using OLS regressions on data from 1950M1 to 2008M1. For TFP I estimate the annualised autocorrelation to be 0.9642, which is in line with usual estimates of TFP persistence. For both the participation rate and dividends an Augmented Dickey-Fuller test fails to reject the hypothesis of a unit root, and the log-AR(1) regressions give coefficients above unity. Since I solve the model under the assumption of stationarity, I choose to set the autoregressive coefficient of participation to $\rho_p = 0.999$, which is stationary but very persistent. For dividends, I choose to set the autocorrelation equal to the estimated autocorrelation of TFP, motivated by long-run dependence of dividends on economic conditions.[26]

In the next step, I estimate the process for labour market discounts using the stock market block of the model. Since dividends follow an independent process, equilibrium stock prices in the model are simply a function of discounts and dividends only. They are given by the function $p_t^s = p(\beta_t^l, d_t)$ which satisfies (7) given the processes for β_t^l and d_t. Since I observe both stock prices and dividends in the data, this allows me to back out the history of discount rate shocks required to match the data by inverting $p(\beta_t^l, d_t)$. Once I have an estimated series of discounts, $\{\beta_t^l\}_{1950M1}^{2008M1}$, I can run regressions on this data to estimate the parameters ρ_{β^l} and σ_{β^l}. The challenge is that constructing the policy function using first-order perturbation requires knowing the autocorrelation coefficient, in order to forecast future stock prices. I overcome this problem using an iterative estimation procedure. I first guess a coefficient ρ_{β^l}, and use it to construct the policy function $p(\beta_t^l, d_t)$. I then back out the series of discount shocks required to match the history of stock-price data, and use this data to estimate a new value of ρ_{β^l}. I repeat until convergence. Stock price data displays both high and low frequency movements. Since I am interested primarily in high frequency movements (for example, during the Great Recession the S&P500 crashed and had fully recovered within 2.5 years) I estimate the AR(1) process for β_t^l on first differences using a lagged-IV strategy, which can consistently estimate ρ_{β^l} in the presence of an unobserved mean which follows a random walk. I discuss this estimation strategy, and robustness to applying it to other variables, in Appendix D.2. This yields a value of $\rho_{\beta^l} = 0.92$, which implies a half life of innovations of around 9 months.

In the final step, I estimate the remaining parameters, $(\psi, \rho_{\beta^k}, \sigma_{\beta^k}, \rho_b, \sigma_b)$, using the full model. I again use an iterative procedure, guessing values of the parameters, simulating the model, backing out the shocks and updating parameters, and repeating until convergence. I use the model to solve for policy functions for capital and labour as functions of the model's state variables: $k_t = k(k_{t-1}, l_{t-1}, s_t)$ and $l_t = l(k_{t-1}, l_{t-1}, s_t)$. Each period I back out the required values of the final two remaining shocks, β_t^k and b_t, in order to match the data on these two variables exactly, given the endogenous state variables (k_{t-1}, l_{t-1}), and the estimated values of the other shocks. I estimate the AR(1) processes for β^k and b using OLS on the recovered shocks.

During the same iteration, I also update the bargaining breakdown probability, ψ, which controls the degree of wage rigidity. Intuitively, the less likely bargaining is to breakdown between periods, the more insulated are outside options, and hence wages, from current conditions, making them less responsive. I estimate ψ by targeting the degree of wage flexibility observed in the data. Haefke et al. (2013) estimate that, after controlling for composition effects, the elasticity of real wages of new hires

[26] The results are robust to reasonable variations in these parameters. Estimating the autoregressive coefficients using the IV strategy outlined below further lowers the parameters to $\rho_d = 0.9611$ and $\rho_p = 0.9476$ and strengthens the results, increasing the initial peak effect of discounts on unemployment from 2.5pp to over 3pp.

to aggregate labour productivity is 0.7 for production and supervisory workers.[27] They use quarterly data from 1979 to 2006, and regress log differenced labour productivity on log differenced real wages. I follow their estimation procedure exactly on simulated data from my model, and therefore construct quarterly data from my model by aggregating up monthly data. I create counterfactual wage and labour productivity data between 1979 and 2006 by simulating the model subject to all observed values of the shocks apart from the worker's unemployment value, b_t, so that the data is being driven by the shocks which drive productivity, and not wages.[28] This results in a value of $\psi = 0.0225$, which is of a similar magnitude to the value of 0.013 used by Hall (2017).

Finally, I estimate the parameters of the reduced-form wage rule used to calculate rationing unemployment. The wage rule is $w_t = \omega \left(z_t k_{t-1}^{\alpha} p_t^{-\alpha} \right)^{\gamma}$, and I estimate ω and γ via an OLS regression of $\log w_t$ on $\log \left(z_t k_{t-1}^{\alpha} p_t^{-\alpha} \right)$ and a constant term. This done on the same simulated data used to estimate the bargaining breakdown parameter, and yields coefficients of $\omega = 0.9436$ and $\gamma = 0.4903$. The reduced-form wage rule approximates the exact model wages very well, which is confirmed by a good visual fit with the data in Figure 9 and an R^2 of 96.4% in the estimation regression.

D.2 Estimation of labour discounts using first differences

I estimate the autoregressive parameter for the stock market discount shock, β_t^l, using an IV strategy in first differences. This allows me to focus on high-frequency variation in the data, and ignore the persistent, non-stationary movements in stock prices over the whole sample period. This is motivated by the fact that the Augmented Dickey-Fuller test fails to reject the hypothesis of a unit root for both stock prices and dividends over the sample period.

I use a lagged IV strategy in first differences, which is capable of estimating the autocorrelation coefficient of an AR(1) in the presence of a mean which follows a random walk. In particular, consider trying to estimate the AR(1) coefficient ρ of a generic variable x_t which follows the following process:

$$x_t = \mu_t + \rho x_{t-1} + \epsilon_t \tag{49}$$

where ϵ_t is a stationary error which satisfies the usual property of being contemporaneously uncorrelated from x_{t-1}. μ_t is an intercept which follows an independent random walk:

$$\mu_t = \mu_{t-1} + \nu_t \tag{50}$$

where ν_t is assumed to be independent of ϵ_t at all leads and lags. The non-stationary intercept means that an OLS regression of the form $x_t = \mu + \rho x_{t-1} + \epsilon_t$ will be misspecified, since μ_t is unobserved and correlated with x_{t-1}. However, the following specification can consistently estimate ρ. Take

[27] This corresponds to the estimates for earnings per person rather than wages, since I hold hours per worker fixed in my model. This is a conservative choice, since they find that wages are less flexible than earnings, with an elasticity of 0.57. This number is also more flexible than estimates of the flexibility of earnings in existing matches. Using the same data, they find an elasticity of 0.39 for all workers. Pissarides (2009) reports numbers in the range of 0.2-0.5 for existing workers in the US.

[28] The role of the shock b_t in the estimation is to allow the model to match the path for unemployment in the data. It does so by creating artificial movements in real wages to change hiring incentives, whenever the other shocks hitting the economy do not fully explain the level of unemployment. Thus, the wage data generated by shocks to b_t are not directly comparable to the data, as they are in fact a stand in any unmodelled shocks to hiring incentives, which is why I exclude the shock when calculating the correlation. Running the regressions on simulated data including the path for b_t leads to a higher correlation between wages and productivity, and hence my results are conservative since including b_t in the simulated data would require an even larger degree of fundamental wage rigidity (lower ψ) in order to match Haefke et al.'s (2013) data.

first differences of (49) to yield

$$x_t - x_{t-1} = \nu_t + \rho(x_{t-1} - x_{t-2}) + \epsilon_t - \epsilon_{t-1} \tag{51}$$

where I used $\nu_t = \mu_t - \mu_{t-1}$. Alternatively, relabel this

$$x_t - x_{t-1} = \rho(x_{t-1} - x_{t-2}) + e_t \tag{52}$$

where $e_t = \mu_t + \epsilon_t - \epsilon_{t-1}$ is the compound error term. Consider regressing

$$x_t - x_{t-1} = \hat{\mu} + \hat{\rho}(x_{t-1} - x_{t-2}) + \hat{e}_t \tag{53}$$

where $\hat{e}_t$ is the estimated analogue of the compound error term e_t. Sadly, $\hat{\rho}$ is not an unbiased estimator of ρ since the X variable, $(x_{t-1} - x_{t-2})$, is correlated with the compound error, $e_t = \nu_t + \epsilon_t - \epsilon_{t-1}$, since x_{t-1} is driven by ϵ_{t-1} in the true data generating process (49). Fortunately, this can be overcome with an IV strategy, since the instrument $(x_{t-2} - x_{t-3})$ is uncorrelated with the error term e_t, but correlated with the X variable, $(x_{t-1} - x_{t-2})$, allowing us to consistently estimate ρ. Note that under these assumptions $(x_t - x_{t-1})$ is stationary and has a finite variance and MA(1) error structure, as can be seen from (51).

If I do not estimate stock market discounts using the IV procedure, but instead use OLS, I find a coefficient of $\rho_{\beta^l} = 0.9986$, implying a yearly autocorrelation of 98%. Using this value of the autocorrelation implies that very small changes in stock market discount rates have very large effects on stock prices, since they are estimated to be so persistent. This means that the model estimates only small increases in stock market discount rates during the Great Recession, and this drastically lowers the estimated impact of stock discounts on employment to less the 0.3pp. However, this value of the autocorrelation is sharply at odds with the speed of the recovery of stock prices from crashes which we see in the data. Instead, the OLS procedure estimates the higher autocorrelation in order to match the longer term movements in stock prices. Additionally, the results for capital discounts are mostly unchanged, with the peak response only lowered from 1.5pp to 1.3pp and the increase by the end of the sample still above 2.5pp. Thus, the effects of adding capital, which are a key contribution of this paper, are robust to the choice of estimation strategy for stock market discounts.

In my baseline results I only estimate the AR(1) coefficient for stock market discounts using the IV procedure, but I have explored robustness to applying it to other shocks. For capital discounts, β_t^k, estimating the autocorrelation with the IV approach leads to an implausibly low value of $\rho_{\beta^k} = 0.6052$, which means a yearly autocorrelation of $0.6052^{12} = 0.0024$. This lowers the estimated effect of capital discounts on unemployment to 1.6pp during the initial peak and to 1.5pp by the end of the sample. However, given that 1) the marginal product of capital, which underlies the capital Euler equation is found to be stationary, and 2) the estimated rise in discounts during the Great Recession is found to be essentially permanent, I prefer to use the baseline results which estimate capital discounts to be more persistent.

Estimating the AR(1) coefficients for the value of home production shock, b_t, dividends, d_t, and participation rate, p_t, using the IV strategy all strengthen the results. Combined, they raise the initial peak effect on unemployment from discounts from 2.5pp to 3.4pp.

D.3 Calculation of internal rates of return

When discounts are time varying, the month-to-month discount at t, β_t, is not fully informative about, for example, the incentives to invest in capital at time t. Since installed capital will persist for many periods, until it is fully depreciated, all of the discount rates over the near future will be important for calculating the total discounted returns from the investment.

In the case of investment incentives, iterating the Euler equation (5) forwards yields the formula

$$p_t^k = \sum_{s=1}^{\infty} \left(\prod_{j=0}^{s-1} \frac{1-\delta}{1+r_{t+j}} \right) \alpha z_{t+s} k_{t+s-1}^{\alpha-1} l_{t+s-1}^{1-\alpha}, \tag{54}$$

where I replaced $\beta_t = 1/(1+r_t)$. Thus, all of $r_t, r_{t+1}, ...$ affect the incentive to invest at t since capital lasts into the infinite future, although it depreciates. The lower is the depreciation rate, δ, the more future discount rates will matter for current investment incentives.

Following Hall (2017), I use the internal rate of return to summarise the total discount rate over the life of an investment. This is calculated as the constant discount rate which sets the net present value of the investment to zero:

$$p_t^k \equiv \sum_{s=1}^{\infty} \frac{(1-\delta)^{s-1}}{(1+\hat{r}_{k,t})^s} \alpha z_{t+s} k_{t+s-1}^{\alpha-1} l_{t+s-1}^{1-\alpha} \tag{55}$$

Note that the internal rate of return at t, $\hat{r}_{k,t}$, is effectively an average of the rates of return over the life of the investment $(r_t, r_{t+1}, ...)$ weighted by the size of the associated payoffs at each date. Once the model has been solved for paths for the endogenous variables, $\hat{r}_{k,t}$ can be directly calculated period-by-period by solving the (non-linear) equation (55).

In the case of job posting incentives, iterating the FOC (6) forwards and applying the internal rate of return concept yields the formula:

$$\frac{\kappa}{q_t} \equiv \sum_{s=1}^{\infty} \frac{(1-\rho)^{s-1}}{(1+\hat{r}_{l,t})^s} \left((1-\alpha) z_{t+s} k_{t+s-1}^{\alpha} l_{t+s-1}^{-\alpha} - w_{t+s} \right) \tag{56}$$

D.4 Model implied wages

While I do not use any time series on wage data in my estimation, only the estimated correlation with labour productivity and average labour share, it is interesting to compare my model-implied wages with various measures of wages in the data. I do this in Figure 8, which shows that model-implied wages share the same broad pattern as wages measured from the national accounts over the whole sample. This reflects the relative stability of the labour share ($LS_t = w_t l_{t-1}/y_t$) over time, and the fact that the estimation procedure matches the time series for both output and employment.

By construction, one should be cautious in comparing my overall equilibrium model-implied wages with the data in any individual time period, since I add the "value of home production" shock, b_t, which drives wages in order to exactly match employment over the sample. By construction, this shock therefore moves wages as a stand in for any non-modelled shocks or frictions which drive unemployment. With that in mind, during the Great Recession, from 2008 to the end of the sample, model-implied real wages fall by more than the fall in the shorter Fed wage tracker wages, but less than wages from the national accounts. During the early phases of the crisis, real wages in the model rise, which is in response to the increase in the value of home production shock which the model uses to explain the residual rise in unemployment not explained by discounts. Real wages in the Fed wage tracker data also rise early in the crisis, but by less.

Figure 8: Real wages: model and data

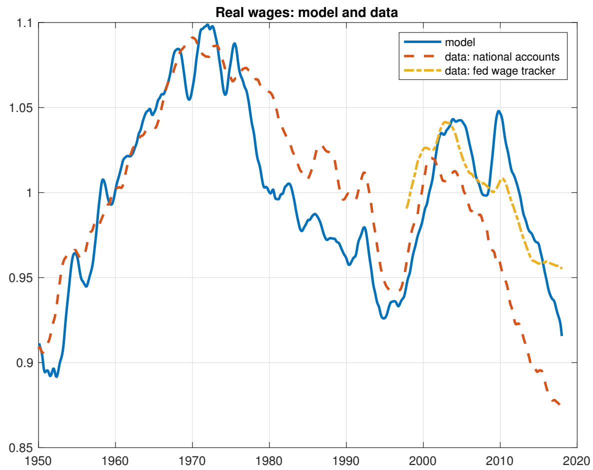

Figure compares the model-implied real wage (solid blue) with two measures of real wages from the data. Real wage data are not used in the estimation of the model, and the model-implied real wage is backed out as part of the decomposition to match the data on other variables, including employment. Dashed red line gives real earnings per person measured from the national accounts, and dash-dotted yellow gives earnings per person computed from the Atlanta Fed's wage tracker. All series are normalised to have mean 1 over the sample.

Figure 9: Performance of approximate wage function

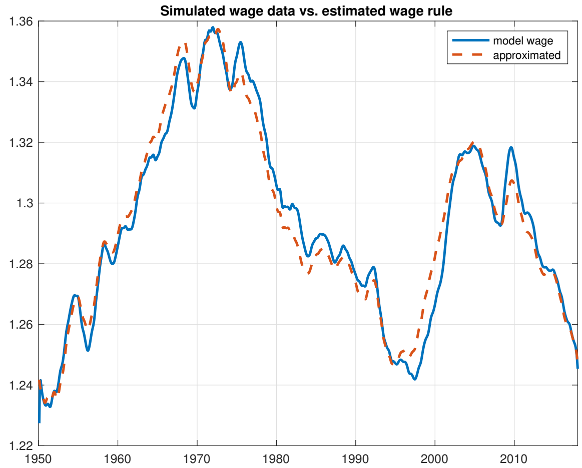

Figure plots the model-generated wage data used to calibrate the degree of wage rigidity (solid blue), which corresponds to a counterfactual simulation of the model subject to all shocks apart from the value of home production. The dashed red line gives the wage implied by the approximated wage rule used to construct rationing and frictional unemployment.

D.5 HP-filter versus linear detrending

In this section I compare my approach of linearly detrending the data to the approach of using the HP filter. In practice, pre-crisis the results of my simple detrending versus using the HP-filter with the literature's standard smoothing parameter of 10^5 are very similar, as I show in Figure 10 for real GDP. Simple detrending finds the peaks and troughs of recessions to be larger, since the flexible trend in the HP filter absorbs some of this cyclicality. However, during the crisis, the well known end-point problem of the HP-filter means that it finds that the cyclical component of GDP has fully recovered by the end of the sample, despite GDP being more than 10% below the pre-crisis trend.

Figure 10: Linear detrending versus HP-filter

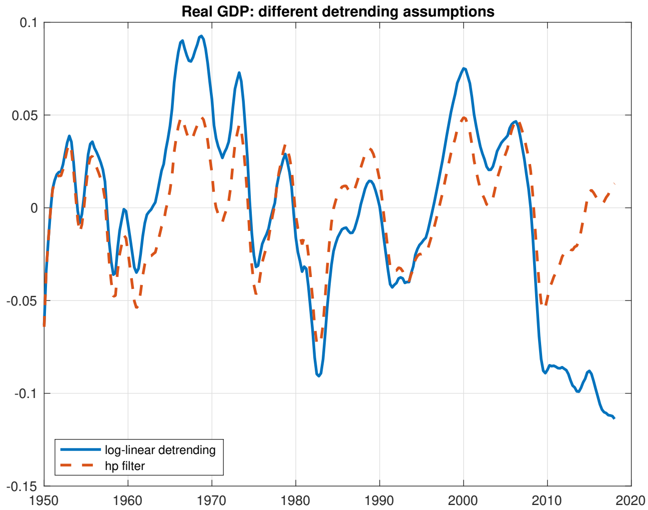

Figure plots real GDP detrended via two methods. Solid blue line gives detrending versus a constant (pre-crisis) growth rate, and dashed red gives detrending of logged GDP via the HP-filter with parameter 10^5.

D.6 Analysis over whole sample

While the focus of this paper is the Great Recession period, in this section I provide a brief analysis of the business cycle properties of the model over the whole post-war sample. Interestingly, the model shows a prominent role for discounts over the whole sample. The data and estimated shocks over the whole sample are given in Figure 11.

Figure 11: Data for full sample

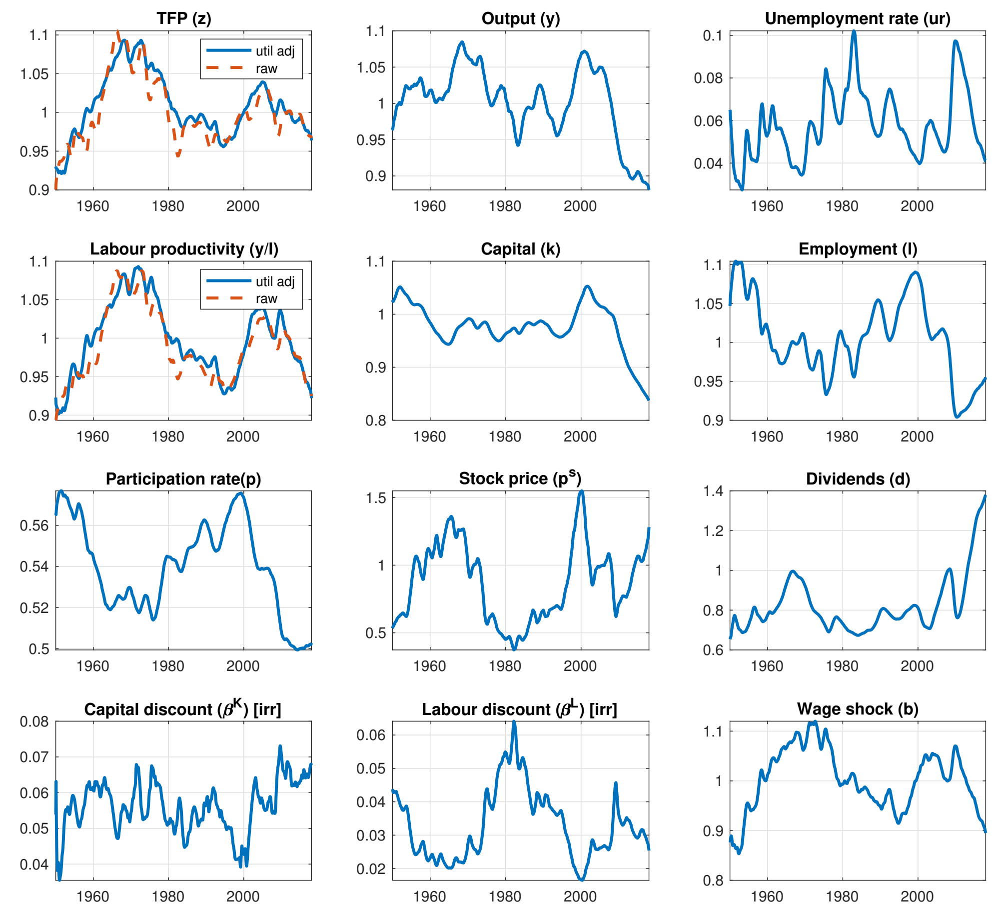

The figure plots my data and estimated shocks over the whole sample period. See Section 4.1 for further details of data construction. All data apart from the unemployment rate, participation rate, and discounts are detrended, with the trend being the average pre-crisis growth rate. Data are real and per capita, and detrended variables are normalised to equal 1 in 2008M1. Utilisation-adjusted TFP and labour productivity data are from Fernald (2014). Discounts are annualised internal rates of return.

In order to assess the role of discounts in driving unemployment over the whole sample, I construct HP-filtered business cycle statistics for unemployment and labour productivity in the spirit of Shimer (2005). Since I have a saturated model where my estimated shock series can exactly replicate the path for unemployment in the data by construction, I am not interested in whether my model can generate the volatility of unemployment seen in the data, but rather interested in *which shocks* contribute most to unemployment. Recall that, intuitively, my unemployment value shock b_t was chosen so that the model could exactly replicate the path for unemployment conditional on the other shocks. Thus, one was to assess the "success" of the model is the ability of the other shocks to generate meaningful movements in unemployment. Following Shimer (2005), I look at HP-filtered deviations of variables from trend. Specifically, I first construct quarterly data by averaging, log any variables considered, and then HP filter them with parameter 10^5. I do this in the data, and therefore also to any series generated from counterfactuals in my model.

Table 2: Business cycle moments: data versus counterfactuals

Moment	Data	β^k	β^l	β^k and β^l	z	b	p
std(u)	0.196	0.093	0.095	0.161	0.193	0.288	0.051
corr(u, u_{-1})	0.970	0.957	0.970	0.975	0.978	0.968	0.971
std(y/l)	0.017	0.004	0.002	0.004	0.01	0.005	0.004
corr($y/l, y_{-1}/l_{-1}$)	0.976	0.98	0.975	0.971	0.978	0.953	0.980
corr($u, y/l$)	0.234	-0.468	0.738	-0.018	-0.992	0.925	-0.907
corr(u, u^{data})	–	0.607	0.548	0.649	0.120	0.513	-0.651

Each column gives the reported statistic from a counterfactual simulation of the model subject only to the estimated shocks for the shock stated in the column. The first column does the same for the true data. The final row computes the correlation between the counterfactual unemployment series and the true unemployment data. All variables are logged and then HP-filtered with parameter 10^5. Data is averaged up to the quarterly level and moments are calculated on data from 1950Q2 to 2017Q3.

In Table 2 I plot various moments of the HP-filtered unemployment rate and labour productivity in the data and counterfactuals from my model, simulated in response to one estimated shock series at a time. The top row gives the standard deviation of unemployment. The first column ("Data") gives the value from the data, which is equal to 0.196, close to the number reported in Shimer (2005). Each remaining column gives the counterfactual standard deviation of unemployment when simulated only subject to one shock. This shows that discounts are able to generate almost all of the volatility of unemployment over the whole sample: individually each discount (columns "β^k" and "β^l") gives a standard deviation above 0.09, and taken together (column "β^k and β^l") they produce a standard deviation of unemployment of 0.161, which is 82% of the standard deviation in the data.

This suggests an important role for discounts in driving unemployment over the whole sample, and not just during the Great Recession. However, looking only at the standard deviation is deceiving: notice that the standard deviation of unemployment generated by the shocks to TFP (column "z") and wage shifter (column "b") are even larger, and the sum of all standard deviations greatly exceeds the standard deviation from the data. This is because the shocks (particularly z and b) are correlated in the sample, so that the negative correlation must be taken into account when summing variances.

To see this most clearly, in the final row I compute the correlation between the counterfactual unemployment series in response to this shock and the actual unemployment rate in the data. A low correlation indicates that, while a particular shock might generate volatile unemployment, it does so in a way unrelated to the unemployment we actually see in the data. Crucially, note the very low (12%) correlation between unemployment generate by TFP shocks and unemployment in the data. Thus, while TFP shocks are volatile and can move unemployment, they do not seem to be the driver of the unemployment variation we actually saw in the postwar period. Looking instead at discounts, the unemployment they generate is very correlated with the actual unemployment we see in the data: capital discounts generate a correlation of 61%, labour discounts 55%, and together a correlation of 65%. In fact, the unemployment generated from discounts is the most highly correlated with actual unemployment during the postwar period out of all the shocks considered.

Finally, discounts offer a potential channel for unemployment volatility which is uncorrelated with labour productivity. This is a well known issue with the search model, since at least Shimer (2005), since unemployment and labour productivity are relatively uncorrelated in the data, while search models which rely on productivity shocks as their driving force imply that unemployment and labour productivity are highly negatively correlated. In my data, unemployment and labour productivity are actually slightly positively correlated (23%) which may be due to using utilisation adjusted TFP. However, simulating the economy subject to only TFP shocks delivers a very counterfactual negative correlation of -99%, hence giving another issue with using TFP shocks as the driving force of the model.

Can discounts generate unemployment which is less correlated with labour productivity? Since discounts do not *directly* move labour productivity this is possible. However, my model with capital reveals an interesting subtlety here, since discounts can drive endogenous movements in labour productivity. As shown in the table, labour discounts driving unemployment actually generate a large positive (74%) correlation between unemployment and labour productivity. The intuition is that, since capital moves little in response to labour discounts due to capital adjustment costs, rising labour discounts reduces employment, raising the capital-labour ratio, and hence raising labour productivity. Thus, labour discounts alone also create a counterfactually large correlation between unemployment and labour productivity, just in the opposite direction from TFP. Capital discounts generate a smaller negative correlation (47%) than TFP, but this is still far from the value in the data. The reason is that capital discounts reduce labour productivity by reducing capital and hence the capital-labour ratio. What is the role of discounts overall? Taking both capital and labour discounts together generates unemployment which is almost completely uncorrelated from labour productivity (-2%) because the effects of the two discounts go in opposing directions. Thus together both discounts are able to capture the relatively low correlation of unemployment and labour productivity.

Overall, these results suggest an important role for discounts over the whole post-war sample. Taken together, labour and capital discounts generate 82% of the unemployment volatility in the data, with the highest correlation to true unemployment of any shock, and generate unemployment which is relatively uncorrelated from labour productivity, as in the data.

In order to illustrate the role of each shock in each recession individually, in Figure 12 I plot the HP-filtered unemployment rate in the data and counterfactuals from my model, simulated in response to one shock at a time. The figure shows how both discount shocks drive unemployment quantitatively significantly and coincidently with the actual unemployment rate. This appears to be true in most recessions, but particularly from the 1990s onwards. The figure also shows how TFP shocks drive often counterfactual unemployment movements: despite generating large movements in unemployment, these are not as highly correlated with actual unemployment from the data. The wage shock also generates volatile unemployment, especially in recessions where the other shocks are not able to fully explain the rise in unemployment. However, its role is also often exactly negatively correlated with the unemployment caused by TFP, and is thus the model's way of neutralising TFP's effect on unemployment in periods where TFP moved but unemployment did not.

Figure 12: HP-filtered unemployment: data versus counterfactuals

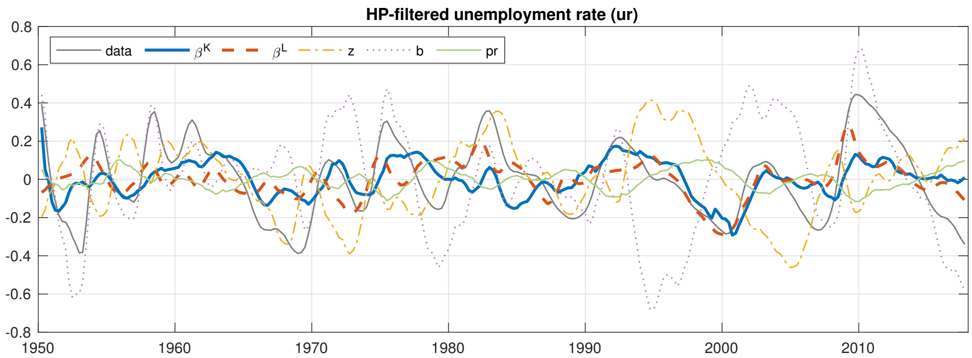

The figure plots the HP-filtered unemployment rate in the data (thin solid black) versus counterfactual unemployment when the model is simulated in response to one estimated shock series at a time. Unemployment is logged and then HP-filtered with parameter 10^5.

I finish the general model analysis by displaying impulse responses to the discounts rate shocks to illustrate how they propagate through the system. I consider negative five standard deviation shocks to the innovations to capital discounts, labour discounts, and TFP for comparison. The shocks are large, since the model is monthly, and typically negative innovations persistent for several months (the results are nearly identical for a sequence of five negative one standard deviation shocks).

The impulse responses to key variables are plotted in Figure 13. In the top left panel I plot the response of the unemployment rate to the three shocks. Capital discounts generate a persistent, hump-shaped response of unemployment. The hump shape reflects capital adjustment costs, so that the capital stock and labour productivity is slow to deteriorate following a discount shock, and the persistence reflects this force and the high estimated persistence of capital discount shocks. Labour discounts generate much less persistent movements in unemployment, which reflects that they are estimated from stock markets, whose estimated discount rate changes are not very persistent. In the seventh panel we see that capital discount generate declines in labour productivity while labour discounts generate slight rises. In the bottom left panel, capital discounts generate protracting falls in capital, while labour discounts have minimal effects on capital. Capital discounts and TFP have similar quantitative effects on unemployment, although the hump-shaped pattern is more pronounced for capital discounts. Labour discounts have a similar effect on unemployment to TFP shocks on impact, but the effects are less persistent.

Figure 13: Impulse responses to discount and TFP shocks

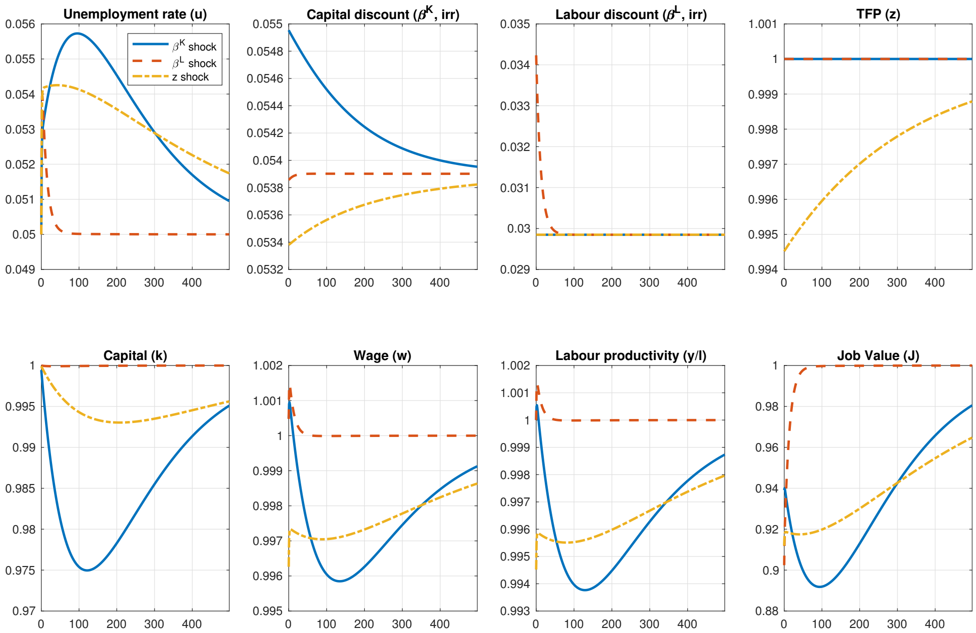

The figure plots impulse responses to shocks to the capital discount, labour discount, and TFP shock respectively. Discount rates are expressed as annualised internal rates of return. All other variables are expressed as deviations from steady state, apart from the unemployment rate.

D.7 Additional figures for Section 4

Figure 14: Full decomposition during Great Recession: additional variables

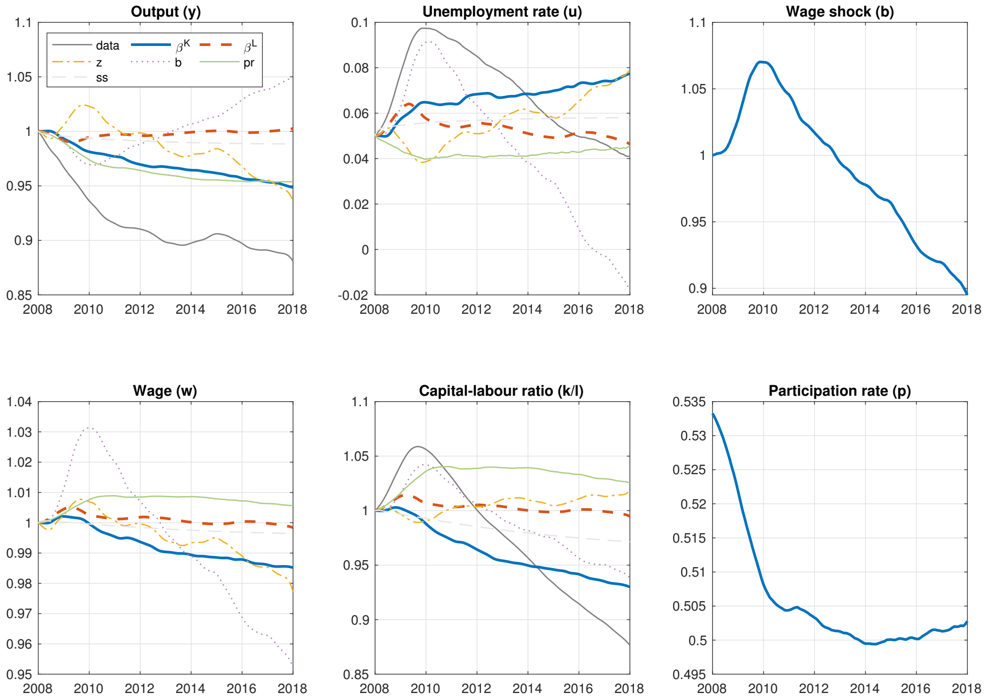

The left and centre panels decompose one variable each during the Great Recession (data: thin black line) into counterfactual contributions from different shocks. The right panels give the estimated shocks for the wage shifter and participation rate. All variables are expressed as deviations from their 2008M1 values, apart from the participation and unemployment rates. Unemployment rate goes slightly negative for the wage shock counterfactual due to approximation error from my numerical approximation method being so far from steady state. Counterfactuals are given for capital discounts (thick blue), labour discounts (dashed red), TFP (dash-dotted yellow), unemployment value (dotted purple), participation rate (solid thin green), and steady state drift (dashed grey).

E Further exercises and robustness

In this appendix I present additional numerical exercises, designed to provide robustness, and give further insights into the main results.

E.1 Nonlinear solution method

My baseline results use a first-order perturbation solution method. This has the advantage of being fast, which is helpful in the estimation of the model given the number of shocks. Nonetheless, the shocks hitting the economy during the Great Recession were large, so one might worry about the accuracy of linearisation methods far from the steady state of the model. In order to address this concern, I also solve the model and repeat my exercises using a fully nonlinear perfect-foresight method. This solves the model's exact nonlinear equations, using the approximating assumption that the path of shocks from 2008M1 onwards is deterministic, and perfectly known to the agents in the economy. The model is coded in Matlab, and I allow 1000 periods between the initial and final steady states. The final steady state is taken to be the values of all endogenous variables and shocks at the end of my data sample.

The model used is exactly the same as the baseline model, with the exception of the wage rule which is replaced with the exogenous wage rule with estimated wage flexibility of $\gamma = 0.4903$.[29] The shock series are now backed out under the assumption of perfect foresight to match the same data aggregates as in the baseline exercise.

I repeat my main exercise, and plot the estimated nonlinear response of unemployment, capital, and labour productivity to the new discount shocks in Figure 15. The results are qualitatively similar to those from the linear solution method, and with two quantitative differences. Firstly, the response of unemployment to capital discounts is now much stronger, which reflects the stronger response of capital to discounts estimated from perfect foresight, since agents now understand that the increase in these discounts is permanent. Secondly, the effect of labour discounts, estimated from the stock market, is now dampened by an initial fall in unemployment which lasts until 2008M12. This reflects a quirk of the perfect foresight assumption, which is that estimated labour discounts must actually fall, not rise, at the beginning of the crisis.[30] However, the increase from 2008M12 to the peak in 2009M7 is of the same magnitude as the total rise in unemployment caused by labour discounts estimated from the linearised model.

[29] I use this rule because it features a constant elasticity of wages with respect to labour productivity regardless of the size of the shock. The alternative-offer bargaining game features non-trivial nonlinearities which make the wage effectively stickier in response to larger shocks. While this is potentially interesting, it makes my estimated alternative offer bargaining parameters, which are estimated in the simulation of the linearised model, unsuitable for use in the nonlinear model. Using the alternative offer wage determination in the nonlinear model would make the effects on unemployment even stronger, but the high degree of wage rigidity makes the solution method unstable. The results from the solution to the linear model with the exogenous wage rule are similar to the baseline results, and can be found in Figure 18.

[30] This is because stock market prices fell quickly during 2008, but still not immediately. In a perfect foresight model, agents perfectly anticipate the future decline in stock prices, and through the Euler equation this should place almost as large downwards pressure on stock prices today. Hence to rationalise why stock market prices did not immediately fall as far in 2008M1 as they eventually do by early 2009, initial discounts must increase.

E.2 Robustness

In this section I perform robustness to several key parameters. These exercises reveal how these important parameters affect my results, and also demonstrate that the results are robust to reasonable perturbations in these parameters. For each robustness exercise I re-solve and estimate the model, and plot the responses of unemployment and key variables to the (potentially different) estimated discount series.

First, in Figure 16 I vary the degree of wage stickiness. I do this by varying the probability of bargaining collapse, ψ, and recalibrating all other parameters to match the moments targeted in the baseline calibration. In the baseline calibration the relationship between wages are labour productivity is calibrated to match Haefke et al.'s (2013) evidence, and wages are reasonably flexible. As shown in the figure, in response to both discount shocks, real wages fall by 65.4% of labour productivity during the Great Recession, which is in line with the 70% elasticity that I match during their 1979-2006 sample. I recalibrate ψ to target lower and higher elasticities of 60% and 80%, and plot the paths of unemployment and real wages during the Great Recession. The results are qualitatively similar to the baseline, with wage stickiness affecting the quantitative magnitudes in the expected way. Quantitatively, wage rigidity is more important for the response to capital discounts, and the combined peak increase in unemployment is increased to 3.5pp for stickier wages and reduced to 1.75pp for more flexible wages.

Second, in Figure 17 I vary the replacement ratio used to calibrate the steady-state flow value of unemployment to the worker. I again change this parameter, and recalibrate all other parameters to match the moments targeted in the baseline calibration. In particular, increasing the flow value of unemployment lowers the required bargaining disruption cost to the firm, χ, required for the bargaining rule to support the targeted steady state wage, as per equation (14). Since the flow value of unemployment, b_t, is assumed to be an exogenous shock, while the firm disruption cost, $\chi(1-\alpha)\frac{y_t}{l_{t-1}}$ depends on the equilibrium MPL, the relative weight of these two values does have effects on the equilibrium responses. In order to quantitatively assess the importance of the targeted value, I take higher and lower values from Hall and Milgrom (2008). They report a replacement rate of 25%, and argue that this can be interpreted as a flow value of unemployment of 71% if interpreted through preferences with curvature in utility. I thus target replacement rates of 25% and 70% as two alternative values. The results are qualitatively similar to the baseline, with the replacement rate interestingly affecting the observed degree of wage flexibility in response to these shocks, despite average wage flexibility during the whole sample continuing to be calibrated to the same value. A lower replacement rate makes wages more sticky in response to discount rate shocks. Quantitatively, this is more important for the response to capital discounts, and the combined initial peak increase in unemployment is increased to 3.4pp for the lower replacement rate and reduced to 1.65pp for the higher replacement rate.

Third, as an additional exercise I also solve the model assuming that true wages are set according to the simple exogenous wage rule used to calculate rationing unemployment. After calibrating the rule to match Haefke et al.'s (2013) data, the results are qualitatively similar to the baseline model, and are plotted in Figure 18. This reflects the behaviour of real wages: the models have the same level of calibrated wage flexibility, and accordingly the responsiveness of wages to the shocks over the whole crisis (measured from 2008 to 2018) are similar. However, under alternating offer bargaining wages initially *rise* in response to rising discounts, which places upwards pressure on initial unemployment. The initial peak responses of unemployment are therefore smaller in this version of the model, with the peak increase from labour discounts reduced from 1.5pp to 1.1pp, and the increase from capital discounts by 2010M1 from 1.5pp to 0.9pp. This result was first discussed by Hall (2016).

Fourth, I investigate robustness in how discount rates are allocated to Euler equations. In particular, the choice to estimate labour market discounts using stock market data is an identifying assumption, following Hall (2017), which I maintain in the paper. We might worry that stock and labour market discounts are not identical, and so to investigate this concern I perform various experiments where I replace the discount rates for labour and capital with various combinations of my estimated discounts. In particular, I replace my Euler equations with the more general formulation:

$$p_t^k = \beta^k \left(\frac{\beta_t^k}{\beta^k}\right)^{\alpha_{k,k}} \left(\frac{\beta_t^l}{\beta^l}\right)^{\alpha_{l,k}} \mathrm{E}_t \left[\alpha z_{t+1} k_t^{\alpha-1} l_t^{1-\alpha} + p_{t+1}^k (1-\delta)\right] \tag{57}$$

$$\frac{\kappa}{q_t} = \beta^l \left(\frac{\beta_t^k}{\beta^k}\right)^{\alpha_{k,l}} \left(\frac{\beta_t^l}{\beta^l}\right)^{\alpha_{l,l}} \mathrm{E}_t \left[(1-\alpha) z_{t+1} k_t^{\alpha} l_t^{-\alpha} - w_{t+1} + \frac{(1-\rho)\kappa}{q_{t+1}}\right] \tag{58}$$

This formulation allows each discount to affect each equation depending on the values of the four exponents. $\alpha_{i,j}$ controls the effect of discount $i = \{k, l\}$ on Euler equation $j = \{k, l\}$. The baseline model corresponds to $\alpha_{k,k} = \alpha_{l,l} = 1$ and $\alpha_{l,k} = \alpha_{k,l} = 0$. To perform robustness I re-solve the policy functions under various assumptions for the $\alpha_{i,j}$ parameters, and then feed through my originally estimated discount rate shocks to redo my main results with different assumptions about how the estimated discounts affect the incentives to invest in capital and labour. The results are plotted in Figure 22, where each line corresponds to a different experiment, with the baseline results in solid blue. "50/50" (red dashed) corresponds to $\alpha_{i,j} = 0.5$ for all i, j, meaning that the discount applied to each Euler is the geometric average of the stock and capital discounts. "Flipped" (dash-dotted yellow) corresponds to reversing which discount is applied to which Euler ($\alpha_{k,k} = \alpha_{l,l} = 0$ and $\alpha_{l,k} = \alpha_{k,l} = 1$), "$\beta^K$ both" (dotted purple) applies the capital discount to both Eulers ($\alpha_{k,k} = \alpha_{k,l} = 1$ and $\alpha_{l,k} = \alpha_{l,l} = 0$), and "$\beta^L$ both" (thin green) applies the stock discount to both Eulers ($\alpha_{k,k} = \alpha_{k,l} = 0$ and $\alpha_{l,k} = \alpha_{l,l} = 1$). The left and centre panels plot the response of the unemployment rate, capital, labour productivity, and job value to both discounts for each configuration. The top left panel plots the response of unemployment, and while the is some variation across configurations, in all configurations apart from "β^K both" the total effect of discounts on unemployment is actually strengthened. The "β^K both" applies the capital discounts to both Euler equations, and therefore completely disregards the large spike in discounts estimated from stock markets at the beginning of the crisis. This leads the model to generate a dampened initial spike up of unemployment in 2009, since capital discounts initially rise less strongly than labour discounts during that period. Nonetheless, even in this most pessimistic scenario the quantitative effect of discounts on unemployment are similar from 2010 onwards, and the total effect of discounts on unemployment appears very robust across configurations.

Fifth, I investigate the role of capital adjustment costs in both measuring capital discounts, and in controlling their effects on unemployment. Adjustment costs are important because they control how much the price of capital responds to changes in investment, and hence affect measured discounts according to the Euler equation (5). I follow the literature in setting adjustment costs so that the elasticity of capital prices to investment is equal to 25%. In Figure 21 I redo my main exercises when I instead target an elasticity of 12.5% and 50%. This shows that higher adjustment costs make the estimated increase in capital discounts during the Great Recession larger, since capital prices are estimated to have fallen further and hence a larger fall in discounts must be required to rationalise why agents would not purchase this capital at a cheaper price. Accordingly, higher adjustment costs also lead to a larger effect of estimated capital discounts on unemployment. Estimated effects are still large even for adjustment costs equal to half of my baseline value, with unemployment only rising 0.3pp less by the end of the sample.

Sixth, I add time-varying job separation shocks to the model. Job separations increased sharply at the beginning of the crisis, and allowing them to increase is important in properly accounting for unemployment in the model. I follow Shimer (2005) in measuring the separation rate using data on those unemployed for less than five weeks, and estimate an AR(1) process for separations, adding it to the list of shocks and recalibrating the model. The results are given in Figure 20. While separation shocks are a driver of unemployment, this does not quantitatively affect the estimated importance of my two discounts measures on unemployment, which remain essentially unchanged. Instead, rising separations simply reduce the role of the unemployment value shock in driving the residual unexplained unemployment in the decomposition.

Seventh, I consider robustness to using raw, rather than utilisation-adjusted, TFP as my measure of TFP in the data. I re-estimate the model and shocks using raw TFP, and repeat my main exercise. The results are given in Figure 19, which shows that the results are very similar. The two TFP measures broadly agree from 2012 onwards, with the chief difference between the two being the behaviour in the early phases of the crisis. The raw TFP measure falls sharply, by 3%, between 2008M1 and 2009M1, before recovering back to its original value by 2011M8. Using this measure of TFP implies that the peak effect of labour discounts on unemployment is only slightly reduced, while the early response to capital discounts is reduced more, from 1.5pp to 1pp by 2009M9. This is because the model now needs a less severe early rise in capital discounts to rationalise why investment fell in the early phases of the crisis, since declining TFP reduces the incentive to invest. However, this reduction in TFP likely represents temporary factor underutilisation, as evidenced by the temporary nature of the decline, and the fact that utilisation-adjusted TFP does not decline during this period, which is why I prefer the utilisation-adjusted TFP measure. Given the similarly of the two TFP measures in the later phases of the crisis the longer term effects of capital discounts on unemployment are affected less.

Finally, I investigate the robustness of my results to changes in participation. In my baseline decomposition I exactly matched the path for the unemployment rate and employment, and backed out the time-series for the participation rate required to square these two series. For robustness, I also take the opposite approach and input data on participation into the model, and back out the implied rate of unemployment. Doing so his minimal effects on the results. To illustrate this, I make an extreme alternative participation assumption: Using demographic data from the CPS on the participation rate by age, I construct a counterfactual participation rate series which holds the participation rates by age constant at their 2008 values, and calculates the decline in participation caused solely by the ageing of the population. Due to the coincidental ageing of the baby-boomer generation into retirement age around the Great Recession, this effect is large, and ageing alone accounts for a 4% decline in the 16+ participation rate between 2008 and 2018, which is 80% of the total observed decline. I feed this participation rate, scaled to generate a 5% steady state unemployment rate, into the the model along with employment data and construct an implied unemployment rate series. This notion of participation therefore assumes that any measured decline in participation *within age groups* during the Great Recession is actually involuntary, and such workers are in fact still in the labour force and should be counted as unemployed. As I show in

Figure 23, doing so massively increases the measured increase in the unemployment rate, with peak unemployment now at nearly 14%. However, using this alternative notion of unemployment has minimal impact on the estimated effects of discounts on unemployment. For more details on data construction for this exercise, see Appendix C.2.2.

Figure 15: Robustness: Results from non-linear solution method

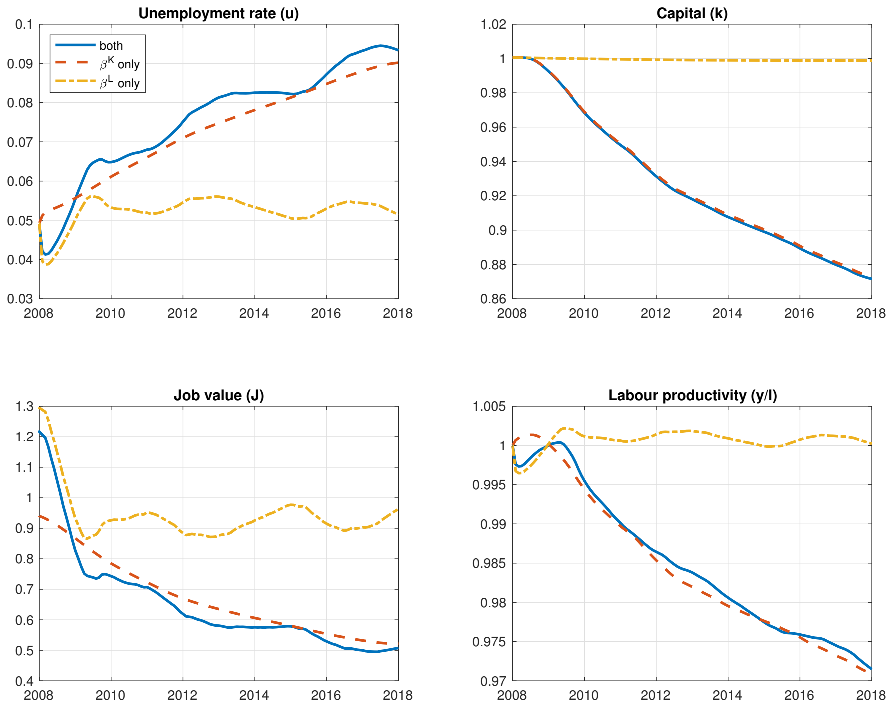

Response of economy to estimated discount shocks when model is solved nonlinearly using perfect foresight methods. Variables are given as deviations from 2008M1 values, apart from the unemployment rate. Solid blue line gives response to both discount shocks combined. Dashed red and dash-dotted yellow decompose this into capital discount and labour discount shocks respectively.

Figure 16: Robustness: wage rigidity

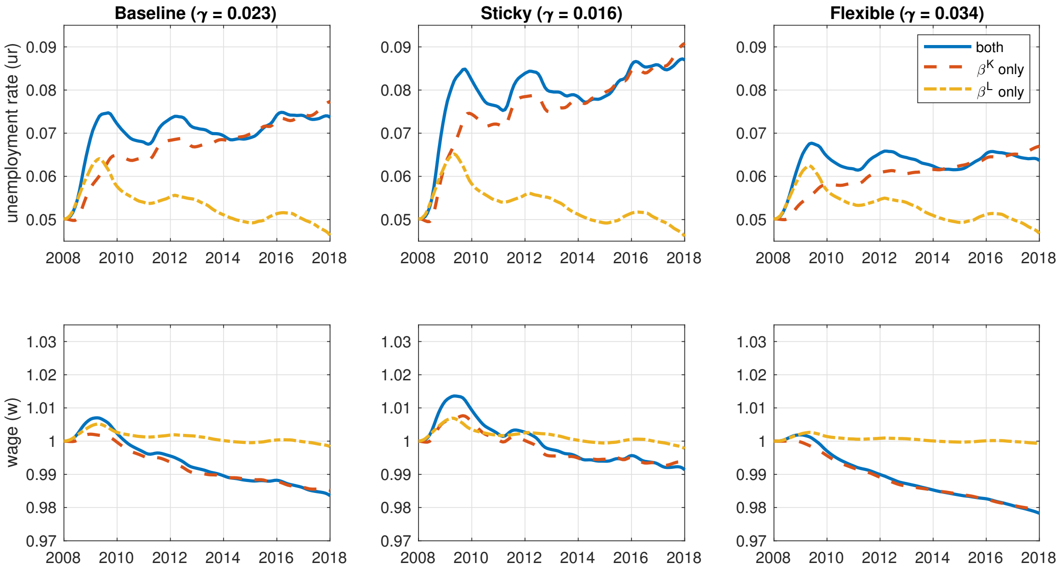

Robustness to different values of wage rigidity. Figure repeats baseline exercise (left panels) for higher (middle) and lower (right) wage rigidity. Top row gives counterfactual unemployment in response to estimated discount rate shocks, and bottom row gives estimated real wage.

Figure 17: Robustness: replacement rate

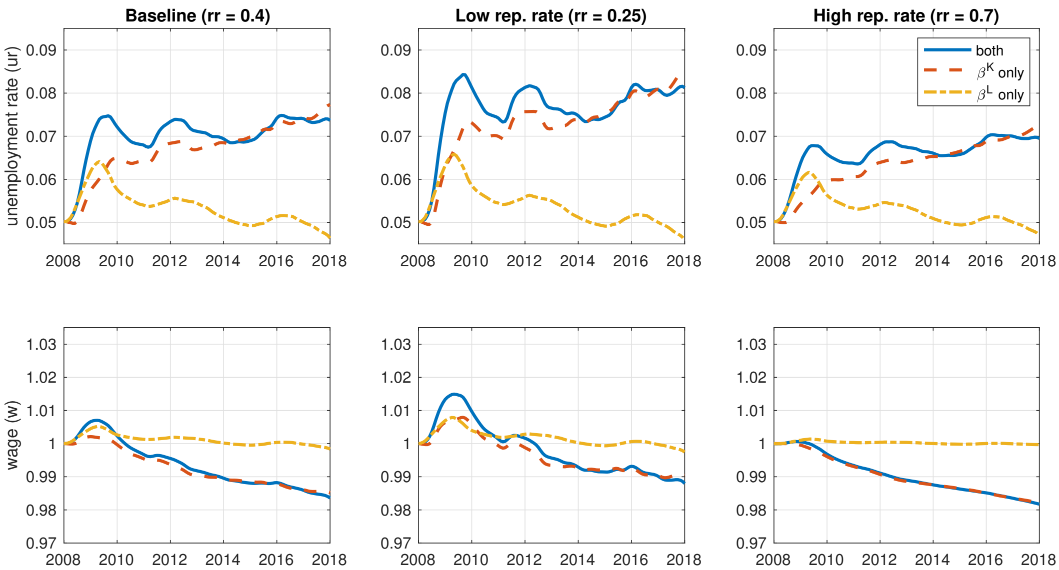

Robustness to different values of wage rigidity. Figure repeats baseline exercise (left panels) for higher (middle) and lower (right) wage rigidity. Top row gives counterfactual unemployment in response to estimated discount rate shocks, and bottom row gives estimated real wage.

Figure 18: Robustness: Results with exogenous wage rule

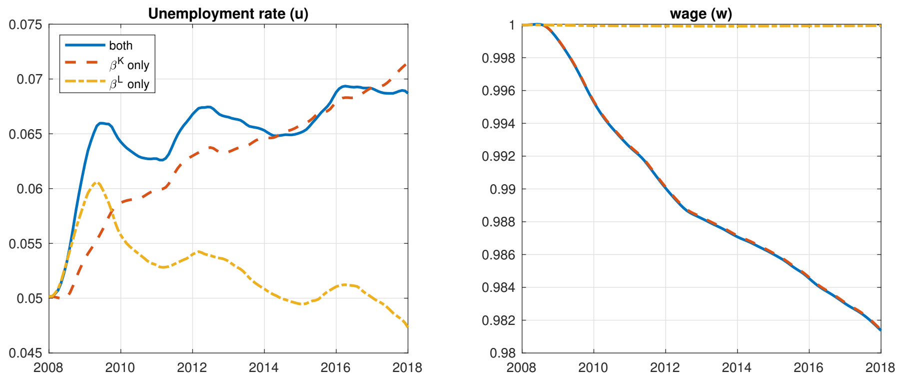

Figure recalculates the counterfactual responses of unemployment and wages to discount rate shocks when wages are instead determined by the calibrated exogenous wage rule (15). Left panel gives counterfactual unemployment, and right panel real wages.

Figure 19: Robustness: Results with raw (not utilisation adjusted) TFP

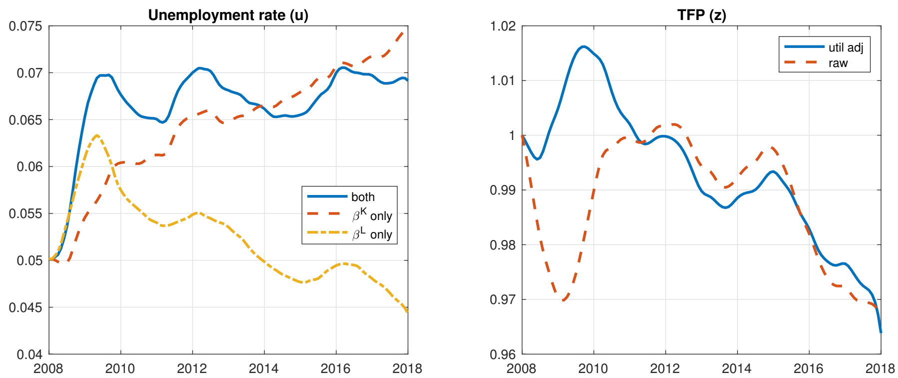

Figure recalculates the counterfactual responses of unemployment to discount rate shocks when TFP data are replaced with the non-utilisation adjusted TFP. Left panel gives counterfactual unemployment, and right panel both adjusted and non-adjusted TFP.

Figure 20: Robustness: Results with job separation shocks

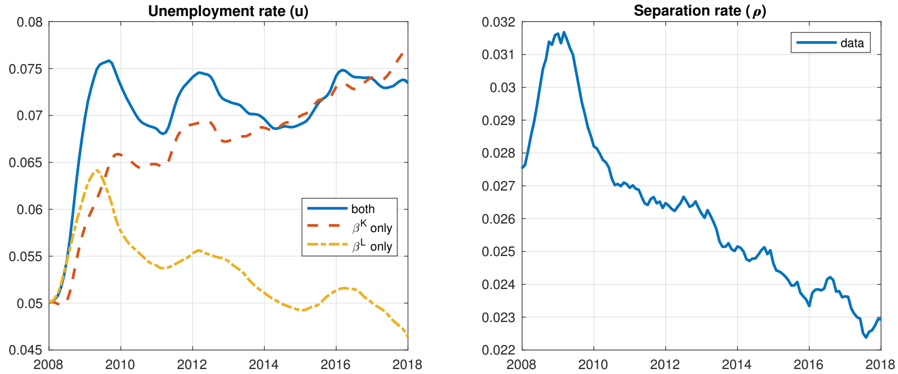

Figure recalculates the counterfactual responses of unemployment to discount rate shocks when the model is enhanced to include job separation shocks. Left panel gives counterfactual unemployment, and right panel separation rate data used in the decomposition.

Figure 21: Robustness: Estimated discounts and decomposition for different adjustment costs

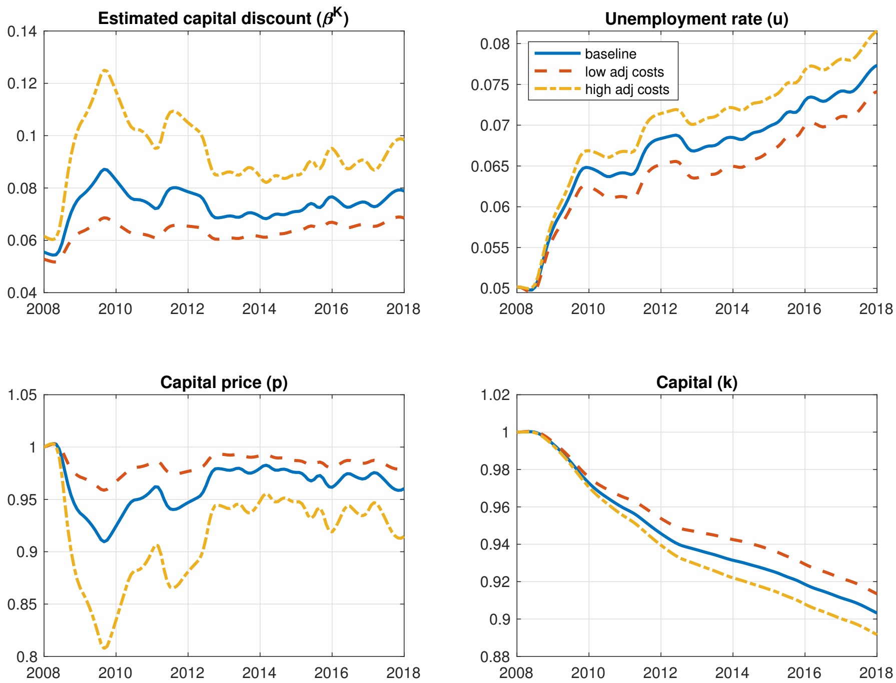

Robustness to higher and lower capital adjustment costs, corresponding to double (dash-dotted yellow) and half (dashed red) the baseline (solid blue) elasticity of capital prices to investment. Top left panel gives the re-estimated discount rate series, and remaining panels give counterfactual unemployment, capital prices, and capital.

Figure 22: Robustness: Allocation of discounts to Euler equations

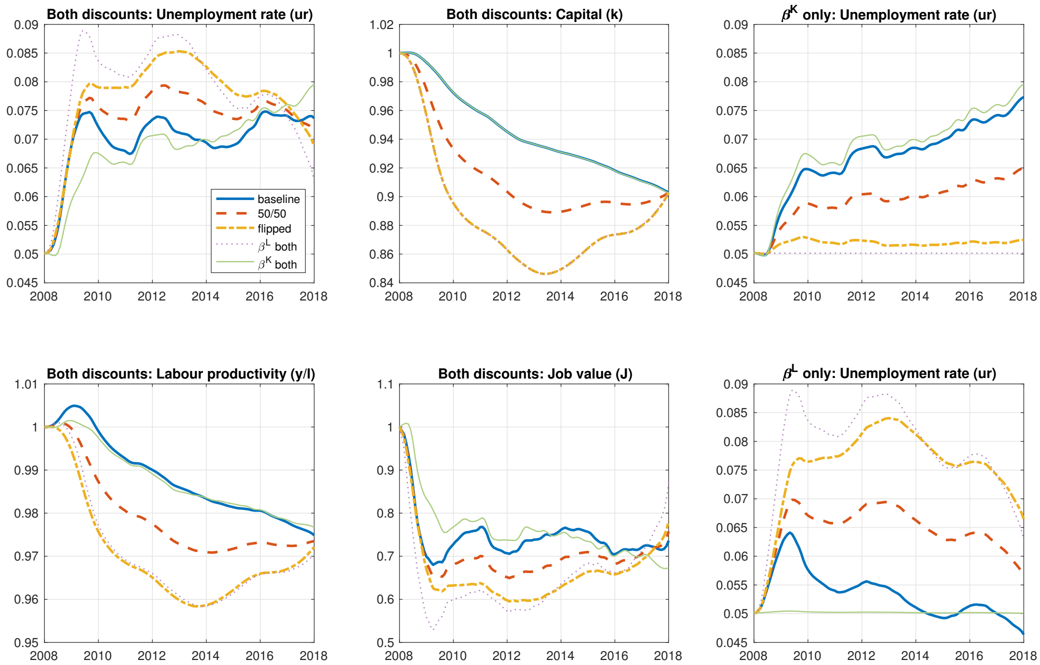

Robustness to different ways of allocating estimated discounts to Euler equations. The left and centre columns plot the counterfactual response to both discounts for unemployment rate, capital labour productivity, and job value. The right column gives the counterfactual for unemployment in response to each discount individually. Each line corresponds to a different allocation of discount rates to the capital and employment Euler equations, as described in the text.

Figure 23: Robustness: controlling for demographic changes in participation

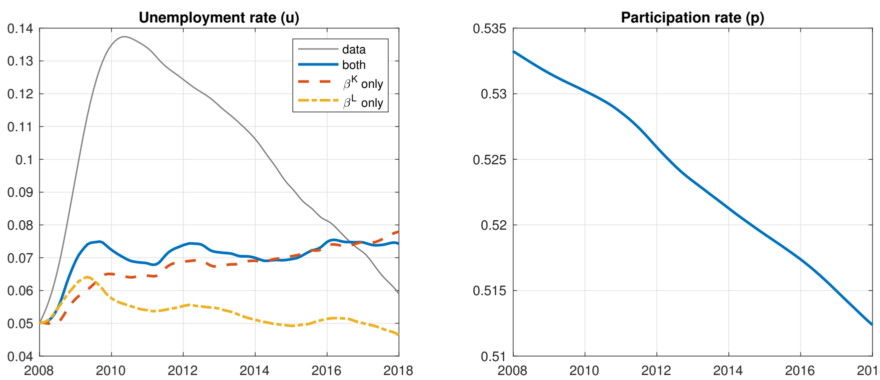

Figure recalculates the counterfactual responses of unemployment to discount rate shocks when participation rate data is replaced with counterfactual participation driven only by demographic change. Left panel gives counterfactual unemployment, including the data in thin black, and right panel the separation rate data.