Firms' precautionary savings and employment during a credit crisis*

Davide Melcangi
University College London†

March 7, 2016

Abstract

Can the macroeconomic effects of credit supply shocks be large even in an economy in which the share of credit-constrained firms is small? I address this question using a model with firm heterogeneity, in which the interaction between real and financial frictions gives rise to precautionary cash holdings. Using UK firm-level balance sheet data, I show that firms hoarded cash relative to their assets during the last recession, and cash-intensive firms cut their workforces by less. A quantitative version of the model, disciplined by these data, generates similar dynamics in response to a tightening of firms' credit conditions. The simulated economy experiences a sizeable fall in aggregate employment and prolonged substitution from capital to cash. Most of the aggregate dynamics are driven by unconstrained firms, pre-emptively responding to changes in credit conditions, in anticipation of future idiosyncratic productivity shocks. The model's ability to generate predictions in line with the data crucially relies on this precautionary channel.

Keywords: financial frictions, precautionary savings, employment, heterogeneous firms
JEL classification: E44, L25, G01, G32

*I am grateful to Vincent Sterk and Marco Bassetto for invaluable advice. For helpful comments, I also thank Morten Ravn, Wei Cui, Mariacristina De Nardi, Wendy Carlin, and seminar participants at University College London and Universitat Autonoma de Barcelona. Thanks to the ESRC for financial support.

†Department of Economics and Centre for Macroeconomics, address: Drayton House, 30 Gordon Street, London, WC1H 0AX, United Kingdom. E-mail: *d.melcangi.12@ucl.ac.uk*

1 Introduction

The recent economic crisis has renewed interest in financial and labour markets, and the potential interconnections that may link them. One of the Great Recession narratives suggests that firm credit tightening is at the root of the increase in unemployment. However, the extent to which financial frictions affect firm's decision-making, and hiring decisions in particular, is controversial. At the aggregate level, firms have large savings and generate internal funds substantially in excess of what they need to finance operations, as documented by Shourideh et al. (2012) for the US and the UK. Moreover, empirical proxies[1] suggest that only a moderate fraction of firms is credit-constrained.[2] These observations might lead to conclude that firm-level credit constraints play a limited role for the cyclical behaviour of aggregate employment.

In this paper, I show that the macroeconomic effects of a credit tightening can be large even in economies in which the share of credit-constrained firms is small. A persistent tightening of credit constraints affects not only the decisions of currently constrained firms, but also those of firms that are currently not credit-constrained but which face some probability of becoming constrained in the future. In the wake of a shock that tightens the credit constraints, these firms may cut investment in capital and hiring for *precautionary* reasons, as this allows them to build up larger cash holdings. More cash alleviates the impact of a credit tightening and reduces the probability of hitting the constraint.

The first contribution of this paper is to build a quantitative model to investigate this precautionary mechanism. I develop a partial equilibrium model with shocks to firms' idiosyncratic productivity and aggregate uncertainty, where precautionary savings in cash arise endogenously from the interaction between real and financial frictions, and affect the transmission mechanism of credit supply shocks onto labour demand. Firms have to finance their wage bill in advance of production and can do so through accumulated cash holdings and an intraperiod loan. Such loans are collateralised with capital and subject to aggregate shocks. Firms face a tradeoff; on the one hand, more cash reduces the probability of being credit-constrained. On the other hand, saving in cash may require cutting back on capital investment and hiring, which in turn reduces production. Firms have incentives to hoard cash because capital is partially irreversible, equity can be issued at a cost and there are hiring and firing costs.

The second contribution of this paper is to use balance sheet data from UK firms to motivate and discipline the model. I use the FAME (Financial Analysis Made Easy)

[1] Silva and Carreira (2012) offer an excellent review of the literature on the identification of financial constraints.

[2] Campello et al. (2010) report that 20% of US CFOs identified their firms as very affected by financial constraints. Angelini and Generale (2008) use Italian matched data and estimate a share of 5%.

dataset, a large panel of UK firms between 2004 and 2013. This sample is more representative than other alternatives often used in the literature, as it mainly includes private firms, whereas US Compustat is limited to publicly listed firms. This feature makes FAME particularly suitable for the study of financial frictions, because private firms are often small and young. They may rely more heavily on external finance, as documented by Shourideh et al. (2012), and have a more limited access to credit (Spaliara, 2009). Finally, FAME balance sheet data also include cash, which I use as a proxy for precautionary savings.

In the FAME data, I document two stylised facts. First, the average cash to assets ratio increases when aggregate employment falls.[3,4] With a simple back of the envelope calculation, I show that the increase in aggregate cash between 2008 and 2009 would have been more than enough to keep the net job creation at pre-crisis levels, if used to hire workers at the average wage. Even if only a share of this excess cash was allocated to the wage bill, the 2009 increase in unemployment rate would have been a third of the one observed in reality. Moreover, the increase in cash ratios in 2009 is common to firms with different fundamentals. Second, I show that cash-intensive firms cut their workforces by less when aggregate employment falls.

The model is calibrated to the UK economy using both aggregate and firm-level moments. Most importantly, the real frictions on capital, labour and dividend payouts are disciplined using FAME firm-level data. Among others, the model matches the cross-sectional distribution of cash ratios. The calibrated model also performs well in approximating additional microeconomic features of the sample, not explicitly targeted. For instance, it correctly predicts that small and more labour intensive firms will hold relatively more cash.

I evaluate the model's ability to explain macroeconomic and firm-level outcomes during the aftermath of the financial crisis, simulating an exogenous tightening of the credit conditions.[5] I show that the precautionary channel allows the model to explain the joint evolution of three key variables: (i) the decline in aggregate employment, (ii) the increase in the aggregate cash-to-assets ratio and (iii) the initial increase and subsequent decline in the cross-sectional correlation between the firm-level cash-to-asset ratio and employment growth rate. The predicted decline in aggregate employment is as large as in the data, despite the fact that in the model the share of credit-constrained firms never exceeds 16%. In fact, I show that unconstrained firms that act for precautionary reasons account

[3]The finding is robust to the median and the aggregate cash ratio.

[4]Bacchetta et al. (2015) document the same fact for US aggregate data and for US publicly quoted firms.

[5]The model displays aggregate uncertainty, which means that firms know the possibility of a credit shock and attach a conditional probability to this event. This is in sharp contrast with other papers evaluating the effects of an unexpected aggregate credit supply shock, as Buera et al. (2015).

for more than 60% of the fall in aggregate employment.

Each of the real and financial frictions present in my model can be found in earlier literature. The financial friction is closely related to Jermann and Quadrini (2012). Adjustment costs in labour and capital can be found, for example, in Bloom (2009). Costs that limit the speed at which firms can raise additional equity are often implemented in the corporate finance literature.[6] I show that, for the precautionary channel to arise in full, these frictions need to be included simultaneously in the model. Indeed, they all play a complementary role: they make it costly for firms to quickly circumvent the effects of a binding credit constraint by either selling capital, firing workers or raising additional equity. I show quantitatively that removing one of the frictions substantially weakens the precautionary channel. In contrast with the data, these versions of the model predict a rise in aggregate capital after a credit supply shock. They also fail to generate many microeconomic features of the data, as the high average cash ratio or the fact that smaller and more labour intensive firms are more cash intensive.

This paper is organized as follows. After briefly reviewing the literature, in section 2 I document empirical stylised facts on cash ratio and employment, which motivate the model, introduced thereafter. Section 3.5 provides intuition for the key model mechanisms. The quantitative analysis starts with the description of the calibration strategy and the data used. I then turn to the steady state performance of the model, before investigating the aggregate effects of a credit tightening and its microeconomic drivers. In section 4.5 I use the model to shed light on the identification of financial constraints. My findings suggest that simple proxies, as those typically used in the empirical literature, do not identify well financial constraints even when we use a structural model calibrated to the data. Finally, I show in section 5 how versions of the benchmark model without some, or all, real frictions would fail to match key empirical predictions.

Related literature This paper is related to various strands of literature. First, it fits into the literature studying the effects of firm credit tightening on aggregate employment.[7] Among others, Jermann and Quadrini (2012) consider a representative agent model where the wage bill needs to be paid in advance of production and is financed through loans secured by a collateral. They show that financial frictions show up as a labour wedge. Representative firm models, however, are not suitable to study economies with a small share of constrained firms and a large amount of savings, as typically observed in the data.

[6]Falato et al. (2012) and Hennessy and Whited (2007) are some examples. Moreover, the cost is also used in the macroeconomic literature, as Jermann and Quadrini (2012).

[7]Examples of empirical papers of this sort are Chodorow-Reich (2014) and Benmelech et al. (2011).

This paper, instead, combines heterogenous firms and financial constraints as done, for instance, by Khan and Thomas (2013) and Hennessy and Whited (2007).[8] Buera et al. (2015) use a model with heterogeneous entrepreneurial productivity and search frictions to argue that a credit crunch can translate onto a protracted increase in unemployment. I bring both firm's labour demand and precautionary cash holdings to the forefront to argue the importance of precautionary responses to changes to credit conditions. Moreover, I back the quantitative model with firm-level data.

Finally, my work also contributes to the literature studying firms' cash holdings. The determinants of cash holdings have been widely studies in corporate finance, both empirically – Bates et al. (2009) – and theoretically – Riddick and Whited (2009) structurally estimate a model with accumulation of liquid assets.[9] Differently from these works, my paper focuses on the amplication effects of precautionary cash holdings. Bacchetta et al. (2015) combine the analysis of firms' cash holdings and employment in face of liquidity shocks. The key distinctive element of my paper is that I allow for occasionally binding constraints both over time and across firms. This implies that not only the model generates a distribution of firms affected to a different extent by the financial friction, but also that firms not facing a currently binding constraint have precautionary reasons to respond to changes in credit conditions. I show that this channel is not only consistent with empirical stylised facts, but also quantitatively important for the aggregate response to a credit supply shock.

2 Stylised facts

This section documents empirical stylised facts on firms' precautionary savings and employment in the UK, during the Great Recession. I show that, during the recent crisis, the average firm started hoarding cash while simultaneously cutting employment, and that cash-intensive firms reduced their workforces by less. I will use these findings to motivate a model in which firms have precautionary reasons to respond to changes in credit conditions.

The primary data source used in this paper is the FAME (Financial Analysis Made Easy) dataset. It comprises panel data observations for a large number of UK firms, for the period 2004-2013.[10] The key advantage with respect to US Compustat consists in the

[8]Producer heterogeneity is also featured, among others, by Bassetto et al. (2015).

[9]Han and Qiu (2007) stress the importance of precautionary motives for corporate cash holdings. Liquid assets are less frequent in macro models and their analysis is mostly motivated by the empirically documented secular increase in aggregate cash holdings in the US; Falato et al. (2013) propose a structural model where rising intangibility of capital accounts for this trend. Kiyotaki and Moore (2012) and Macera (2013) propose models where firms have incentives to *save for the sunny days*.

[10]This is the time frame I consider in this paper. FAME data are available only 10 years at a time,

Figure 1: Aggregate increase in cash ratio and aggregate net job creation in the UK

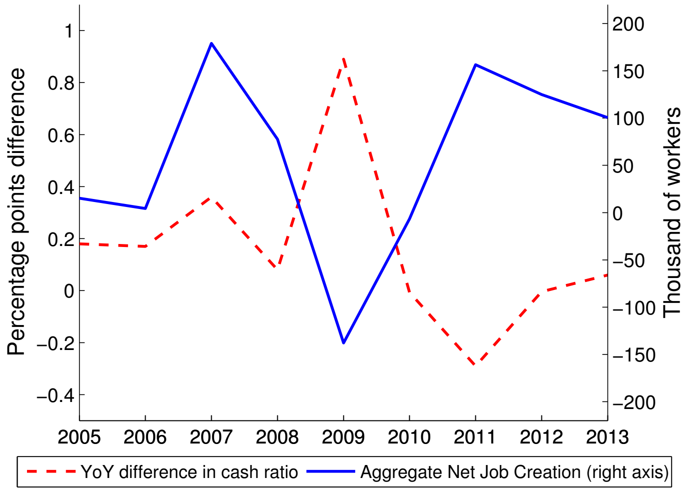

Notes: Firm-level net job creation is the difference in number of employees for a given firm from one year to the other (Moscarini and Postel-Vinay, 2012). The cash ratio is the sum of total cash & equivalents over total assets. The dashed line shows the year-on-year differences in the cross-sectional average cash ratio. The sample considers firms that have weakly positive observations for employment, cash and total assets for all the years 2004 - 2013. Results hold for an unbalanced panel too. Firms with UK SIC code referring to "Financial and insurance activities" are excluded.

ownership structure of the firms. FAME does not only include publicly quoted firms; on the contrary, they are mainly private. This allows to have a representative sample, where the size and age distribution of firms is close to the UK universe of firms. The presence of young and small firms makes FAME particularly suitable for the analysis of financial frictions, since those firms are likely to rely more heavily on external finance and face more difficulties in accessing credit. The dataset contains firm-level information on both the asset and employment structure of the firms. It also includes data on cash holdings, recorded in firm's balance sheets as *Bank deposits*. Appendix A provides additional information on the data.

I document that, in the sample considered, the average cash ratio[11] increases when aggregate employment falls.[12] Figure 1 shows the first differences of average cash ratio and aggregate employment. In 2009, the year of the trough in aggregate employment of the UK dataset, firms hoarded a large amount of cash with respect to their assets.

The increase in cash ratio in 2009 is quantitatively sizeable and a simple back of the envelope calculation can show this. Suppose the cash ratio remained constant after 2008, and the cash in excess of this counterfactual cash was used to hire workers at the average

with older years being replaced as more recent become available.

[11]The cash ratio is defined as the share of cash holdings over total assets. The results are very similar for aggregate cash ratio and median cash ratio across firms.

[12]Bacchetta et al. (2015) find that, in the US, deviations of aggregate cash ratio and employment are negatively correlated. The correlation for my data is -0.66.

Figure 2: Cash and unemployment: a back of the envelope calculation

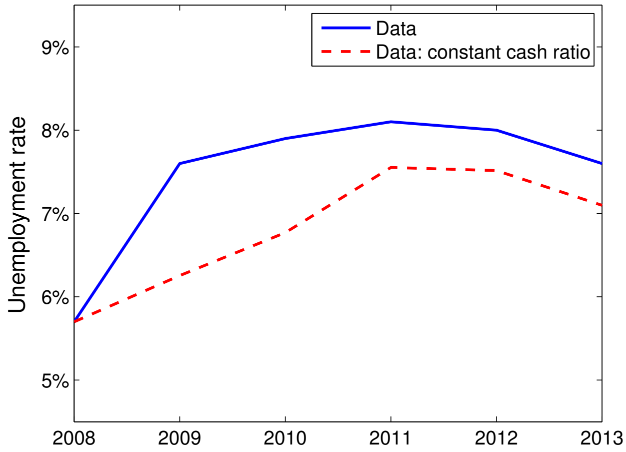

Note: The dashed line is calculated as follows. At each period, the excess cash is computed as $\tilde{C}_t = \left(C_t - \hat{C}_t\right) - \left(C_{t-1} - \hat{C}_{t-1}\right)$, where C_t is the empirically observed aggregate cash and $\hat{C}_t$ the aggregate cash required to keep the cash ratio constant at 5.38%, its value in 2008, conditional on observed aggregate total assets. Every period, 27% of $\tilde{C}_t$ is used to hire additional workers at the cross-firm average wage, and the remaining excess cash is used for other purposes and gone. The additional workers are scaled up by a factor 4.99, which is the ratio between ONS and FAME aggregate employment, and added to the time series of unemployed workers (ONS). Finally, the unemployment rate is computed dividing the counterfactual unemployment by the ONS labour force.

wage.[13] In this scenario, 317,089 additional workers would have been hired, more than offsetting the fall in aggregate employment observed in FAME.

It may be argued that salaries are not the only expenses that a firm faces. In the FAME data I find that the wage bill accounts, on average, for 27% of operating and capital expenditures.[14] Allocating only this share of additional cash to the wage bill, firms in FAME could have hired 85,249 additional workers in 2009, unwinding more than half of the negative job creation. Figure 2 shows how important this is in the aggregate. I scale up FAME additional net job creation to the UK economy and compute the counterfactual unemployment rate, shown by the dashed line.[15] Under this scenario, the increase in unemployment rate in 2009 is less than a third of the one observed in reality, because part of the excess cash is used to hire workers. By doing so, the firms in the counterfactual scenario have a lower stock of cash in the following years, and thus unemployment rate increases faster between 2010 and 2011 than in reality.

[13]The counterfactual cash is the aggregate cash required to keep the aggregate cash ratio constant, taking aggregate total assets as given. The average wage in 2008, in the FAME sample, is £31,377.

[14]Operating expenditures are all the expenses before the EBITDA.

[15]In 2008, 29.6 million workers were employed in the UK, five times the aggregate employment in my FAME sample. This factor remains constant over time, confirming that FAME and ONS data display similar employment dynamics.

The increase in cash ratio does not seem to be driven only by constrained firms. I classify the firms in 2005 by quartiles of size distribution. Size is a popular proxy for financial constraint in the empirical literature.[16] The largest firms experience the largest increase in cash ratio between 2009 and 2008, slightly above 12%. Small firms, by contrast, increase their cash ratios by 1% only.

Besides documenting the cyclical patterns of cash ratios, the data can shed some light on the role played by precautionary cash holdings in the transmission mechanism of credit shocks onto the labour market. Figure 3 shows, at different years, the cross-firm correlations between lagged cash ratio and employment growth. The correlation more than doubles in 2009, while then turning negative in 2011.[17] I propose a possible explanation to this behaviour. Consider a credit supply shock that dries up external liquidity available to firms. Firms with higher cash ratios are better equipped to cope with the crisis, and thus they have to cut their workforces by less. During the credit tightening period, most firms hoard cash to partially counteract the scarce external funding, and smooth employment growth throughout the recession. A share of firms, however, is likely to be so disrupted by the credit supply shock that is not able to internally generate liquidity. This cash-scarce group of firms leads the recovery when credit conditions are restored, driving the cross-firm correlation negative.

I will show how my model will be able to generate similar dynamics. An increase in cash has a twofold effect: on one hand, it takes resources away from production, amplifying the negative effects of credit shocks. On the other hand, in the following periods, it helps the adjustment to tighter credit conditions and, coupled with labour adjustment costs, smoothes the recovery.

There could be alternative potential explanations of the dynamic evolution of the cross-firm correlation. Firms that manage to grow even during the crisis could receive more revenues, which would translate in higher cash. As suggestive evidence against this explanation, I regress employment growth on lagged cash ratio at each year, including cash flow as a control: the dynamic evolution of the coefficients on lagged cash ratio tracks the correlations shown in figure 3. Similarly, this should account for the possibility of an unexpected negative productivity or demand shock that induces firms to lay off workers and generates more cash flow. If that was the case, we should see a drop in the correlation between cash ratio and employment growth in 2009, which does not happen instead. Finally, the results are robust to the inclusion of sector fixed effects.

[16] Repeating the same analysis using age as a proxy delivers the same qualitative results.

[17] To get a sense of the magnitude of these fluctuations, I regress employment growth on lagged cash ratio at every year and then compute the marginal elasticities. A 1% increase in cash ratio in 2008 is associated with 0.63% higher employment growth in 2009. Similar results hold with a panel regression with interacted year dummies.

Figure 3: Cross-firm correlations between lagged cash ratio and employment growth - UK data

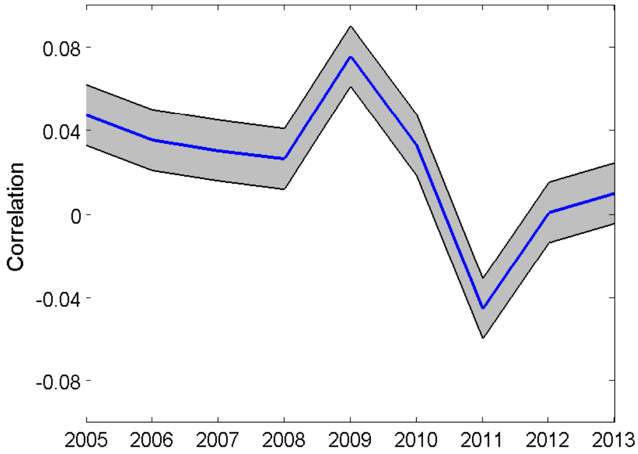

Notes: UK FAME firm-level data, YoY changes. The sample considers firms that have weakly positive observations for employment, cash and total assets for all the years 2004 - 2013. Results hold for an unbalanced panel too. Firms with UK SIC code referring to "Financial and insurance activities" are excluded. Employment growth for a firm j at year t is calculated as $\Delta n_{j,t} = \frac{n_{j,t} - n_{j,t-1}}{\alpha n_{j,t} + (1-\alpha) n_{j,t-1}}$, with $\alpha = 0.5$; Moscarini and Postel-Vinay (2012) explain the advantages of this symmetric approach. Shaded grey bands indicate 95% confidence intervals.

3 The model

I consider a partial equilibrium model that investigates firms' behaviour. The economy is populated by heterogeneous firms that are subject to idiosyncratic productivity shocks and aggregate credit shocks. Firms can invest in physical capital, used for production together with labour, or in liquid assets. They face a liquidity need originated by the payment of the wage bill, which can be covered either by external intra-period loans or by cash holdings. This assumption generates interactions between employment and portfolio choices. Short-term borrowing is collateralised by tangible assets, in the form of capital, and subject to persistent credit shocks, which restrict the amount of loans for a given level of collateral. Firms incur capital and labour adjustment costs and can issue equity, at an increasing and convex cost. These elements give rise to firms' precautionary behaviour, further exacerbated during tight credit periods. I will start presenting the main features of the model and the firm's value function. Section 3.5 will shed further light on the key mechanisms of the model.

3.1 Technology

The economy is populated by a very large number of infinitely-lived heterogeneous firms that use capital k and labour n to produce a final good. I assume that each firm operates a diminishing returns to scale production function with capital and labour as the variable inputs. A firm produces output y according to:

$$y_t = z_{j,t} k_t^{\nu} n_t^{\omega} \ , \ \nu + \omega < 1 \tag{1}$$

where $z_{j,t}$ is a stochastic and persistent idiosyncratic productivity[18] that follows a Markov chain: $z \in Z \equiv z_1, .., z_{N_z}$, where $Pr(z_{t+1} = z_i | z_t = z_j) = \pi_{ji}^z \geq 0$ and $\sum_{i=1}^{N_z} \pi_{ji}^z = 1$.

3.2 Working capital constraint

I assume that firms need to pay their wage bills in advance of production. As in Jermann and Quadrini (2012), this stems from the cash-flow mismatch between the payments made at the beginning of the period and the realization of revenues.[19] Fernandez-Corugedo et al. (2011) analyse UK firms working capital positions over the business cycle and find that firms have typically a funding gap between the payments of the costs of the inputs to production and the sales revenues, which typically come much later.

The timing goes as follows. At the beginning of the period, after the realization of the idiosyncratic and aggregate shocks, firms choose the stock of workers that will be productive in the same period. They have to pay the associated wage bill $\overline{w} n_t$ out of accumulated cash m_t, before the realization of revenues, which come at the end of the period. If the wage bill exceeds the accumulated cash, firms can obtain external funds at the beginning of the period and repay at the end of it. This form of intra-period loan entails no interest, as in Jermann and Quadrini (2012), and cannot be larger than a stochastic fraction ϕ of collateral, that is, the liquidation value of capital. The following equation describes how, according to the collateral constraint, the financing funds need to be greater or equal than the financing needs:

$$\underbrace{\phi_{s,t}(1-\vartheta)(1-\delta_k)k_t + m_t}_{\text{financing funds}} \geq \underbrace{\overline{w} n_t}_{\text{financing needs}} \tag{2}$$

The ability to borrow intra-temporally is bounded by the limited enforceability of debt contracts, as firms can default on their obligations. The decision to default arises after the realization of revenues. There is no default at equilibrium. Since liquidity can be easily diverted, the only asset available for the liquidation is physical capital k_t, as

[18] Since the model is in partial equilibrium, the production function can be seen as a revenue function where z combines productivity and demand terms into one index, as in Bloom (2009).

[19] Jermann and Quadrini (2012), Christiano et al. (2010) and Mendoza (2010) consider advance financing of the wage bill in a model with financial frictions. This assumption is employed also in cash-in-advance models as Candless (2006) and monetary policy models (Christiano et al. (1996)).

in Kiyotaki and Moore (1997). In particular, this will be the non-depreciated fraction of capital; moreover, lenders incorporate the fact that, in case of default, they will sell the seized capital at a lower, resale price $(1-\vartheta)$. This form of partial irreversibility will be described in the following subsection.

The collateral fraction $\phi \in \phi_1, .., \phi_{N_\phi}$ is assumed to be common to all firms and will be referred to as *credit tightness*. It is assumed to follow a Markov chain, with $Pr(\phi_{t+1} = \phi_m | \phi_t = \phi_s) = \pi^{\phi}_{sm} \geq 0$ and $\sum_{m=1}^{N_\phi} \pi^{\phi}_{sm} = 1$. This variable can be interpreted in many ways. It could reflect the efficiency of the economy's financial sector, as in Khan and Thomas (2013), or capture the variations over time in the degree of credit market tightness (Finocchiaro and Mendicino, 2013). Eisfeldt and Rampini (2006) provide some evidence about the cyclicality of ϕ. The quantitative analysis in section 4 will consider a drop in ϕ, resembling an exogenous reduction in the amount of available external funds.

Similarly to Svensson (1985), cash holdings decisions are made before the realization of the shocks, which gives rise to precautionary incentives to accumulate cash. The latter will be softened by the possibility to top up the wage bill payment through external financing. As will be explained later, the presence of real frictions act in the opposite direction, amplifying the incentives to behave pre-emptively. Finally, the wage is assumed to be fully rigid and common across firms. The relevance of this assumption for the model dynamics and the quantitative results will be investigated in Appendix E.

3.3 Real frictions

Besides the working capital constraint, firms face three real frictions. As will be made clearer later on, the interaction between real and financial frictions implies that precautionary cash holdings arise endogenously in the model.

Firms face linear and symmetric hiring and firing costs, as in Bloom et al. (2012). The firm begins the period t with a pre-determined employment stock n_{t-1}, a fraction δ_n of whom immediately separates. Firms choose the new stock of workers, pay the wage bill and use pre-determined capital and the newly available labour to produce. The labour adjustment costs can be summarized as follows:

$$AL(n_{t-1}, n_t) = \chi | \left(n_t - (1-\delta_n)\, n_{t-1} \right) | \tag{3}$$

Consistent with the typical timing convention, capital k_t is chosen at time $t-1$ and predetermined at time t. It evolves according to $k_{t+1} = (1-\delta_k)k_t + i_t$ where i_t is investment and δ_k is the depreciation rate. The only capital adjustment cost faced by the firm is a partial irreversibility. Firms buy capital at a unitary price, as in a neoclassical growth model, but, for each unit of used assets, only $(1-\vartheta)$ fraction is useful for other buyers.

ϑ represents the reallocation costs, the partial irreversibility of the capital stock due to capital specificity or adverse selection problems. If the firm changes its quantity of capital from k_t to k_{t+1}, the cost of doing so is:

$$AC(k_t, k_{t+1}) = \begin{cases} k_{t+1} - (1-\delta_k)k_t & \text{if } k_{t+1} \geq (1-\delta_k)k_t \\ -(1-\vartheta)[(1-\delta_k)k_t - k_{t+1}] & \text{if } k_{t+1} < (1-\delta_k)k_t \end{cases} \tag{4}$$

Finally, I assume that firms incur in a quadratic cost[20] when they deviate from a target level of dividends, given by: $\xi(d_t) = \kappa(d_t - \overline{d})^2$, as in Jermann and Quadrini (2012) and Finocchiaro and Mendicino (2013). This cost is in line with empirical evidence that underwriting fees display increasing marginal cost in the size of the offering (Altinkilic and Hansen, 2000). It is also a reduced form to capture the fact that managers are concerned with smoothing dividends over time. The seminal work by Lintner (1956), repeatedly confirmed by more recent studies, found that approximately 90% of firms smooth their dividend payments with respect to their earnings. It may be argued that a single dividend target is a restrictive assumption in a heterogeneous firms model. I will show in the following sections that this helps the model getting closer to the low volatility of dividends over time observed in the data. With a productivity-specific dividend target, or even in the absence of a dividend cost, dividends would fluctuate more freely in response to idiosyncratic productivity shocks.

3.4 The Firm's value function

I denote $V(m_t, k_t, n_{t-1}, z_{j,t}; \phi_{s,t})$ the value function of a firm. The 5 state variables are given by (1) the firm's cash stock m_t, (2) the firm's capital stock k_t, (3) the firm's stock of workers n_{t-1}, (4) the firm's idiosyncratic productivity $z_{j,t}$ and (5) the aggregate credit tightness $\phi_{s,t}$. The dynamic programming problem of the firm consists of choosing dividends, labour, capital next period and cash next period to maximise the present discounted value of future dividends:

[20] It is possible to consider different adjustment costs, for instance an asymmetric equity issuance cost as in Corbae and D'Erasmo (2014). The majority of macroeconomic models with heterogeneous firms restrictively assumes no equity issuance.

$$V(m_t, k_t, n_{t-1}, z_{j,t}; \phi_{s,t}) = \max_{d_t, m_{t+1}, k_{t+1}, l_t, n_t} \left\{ d_t - \xi(d_t) + \beta \sum_{m=1}^{N_\phi} \pi^\phi_{sm} \sum_{i=1}^{N_z} \pi^z_{ji} V(m_{t+1}, k_{t+1}, n_t, z_{i,t+1}; \phi_{m,t+1}) \right\}$$

subject to:

$$\begin{aligned}
z_{j,t} k_t^\nu n_t^\omega - l_t &= AC(k_t, k_{t+1}) + AL(n_{t-1}, n_t) + m_{t+1} + d_t && (5)\\
l_t &\geq \overline{w} n_t - m_t && (6)\\
l_t &\leq \phi_{s,t}(1-\vartheta)(1-\delta_k)k_t && (7)\\
m_{t+1} &\geq 0 && (8)\\
\xi(d_t) &= \kappa(d_t - \overline{d})^2 && (9)
\end{aligned}$$

and (3)-(4). I denote with l_t the intra-period loan that the firm receives at the beginning of the period and repays at the end of it. In equilibrium, (6) is always binding, since firms borrow exactly the amount of wage bill that exceeds the accumulated cash. In contrast, (7) binds occasionally over time and across firms. All firms face the same aggregate credit tightness $\phi_{s,t}$. The wage $\overline{w}$ is exogenous, time-invariant, common across firms and taken as given by each firm.[21] Let $k^*(m_t, k_t, n_{t-1}, z_{j,t}; \phi_{s,t})$, $m^*(m_t, k_t, n_{t-1}, z_{j,t}; \phi_{s,t})$, $n^*(m_t, k_t, n_{t-1}, z_{j,t}; \phi_{s,t})$ and $d^*(m_t, k_t, n_{t-1}, z_{j,t}; \phi_{s,t})$ represent the optimal choices of next-period capital and cash, labour and dividends respectively, made by the firm with current idiosyncratic productivity $z_{j,t}$ and under aggregate credit tightness $\phi_{s,t}$. I characterize these decision rules in section 3.5.

3.5 Firm's behaviour

Before turning to the quantitative part of the paper, it is useful to shed some light over the main mechanisms generated by the model. I will start by showing the trade-off between capital and cash, and how this is affected by the real frictions. I will then turn to the hiring decision and finally show how the precautionary mechanism falls apart when each of the real frictions is removed.

Firms face a trade-off between very liquid but unproductive assets, denoted as cash, and productive but partly liquid and partly collateralizable assets, capital. This is shown in equations 10 and 11, which show the first order conditions of a firm with idiosyncratic productivity $z_{j,t}$ for capital and cash respectively:

[21] Appendix E discusses this assumption and the robustness of the main results to general equilibrium.

$$\beta E_t \left[\underbrace{\lambda_{t+1} \nu z_{t+1} k_{t+1}^{\nu-1} n_{t+1}^{\omega}}_{\mathrm{I}} - \underbrace{\lambda_{t+1} AC_k \left(k_{t+1}, k_{t+2}\right)}_{\mathrm{II}} + \underbrace{\phi_{t+1}(1-\vartheta)(1-\delta_k)\mu_{t+1}}_{\mathrm{III}} \right] = \lambda_t AC_k \left(k_t, k_{t+1}\right) \tag{10}$$

$$\beta E_t \left[\lambda_{t+1} + \mu_{t+1}\right] = \lambda_t - \psi_t \tag{11}$$

where λ_t, μ_t and ψ_t are the Lagrange multipliers associated to (5), (7) and (8) respectively.[22] For illustrative purposes, I leave the derivatives of labour and capital adjustment costs unspecified.[23] The left hand side of equations (10) and (11) shows the marginal benefit of holding an additional unit of capital and cash respectively. For (10), this can be decomposed in three parts: (I) the expected marginal product of capital, (II) the expected marginal benefit that an additional unit of capital brings tomorrow in terms of reduced adjustment costs and, finally, (III) the expected marginal benefit of holding capital as collateral. Capital is only partly collateralizable and hence its *financing return* $\beta E_t \phi_{t+1} \mu_{t+1}$ is scaled down by a factor $(1-\vartheta)(1-\delta_k)$ smaller than 1. The portfolio allocation between capital and cash is forward-looking, since decisions taken this period affect the financing conditions in the following. In other words, employment growth can be sustained by different allocations of internal and external financing.

The real frictions play a crucial role for the endogenous accumulation of cash. The partial irreversibility makes capital less liquid and, in turn, shifts the portfolio allocation towards cash. First, it scales down the *financing return* by a factor $(1-\vartheta)$, as shown in III. Moreover, it both affects the marginal benefit and cost of holding capital. Appendix C deals in detail with how much firms value their capital, especially when their investment decision is inaction. Intuitively, firms incorporate the fact that a negative shock may require to sell off capital, at a lower resale price. Hence, they act pre-emptively and hoard more cash instead.

The dividend cost also adds to firms' precautionary incentives. It implies that the

[22] Complementary slackness conditions for λ, μ and ψ have been omitted. For all the equations of this section, the expectation operator E_t is used as a reduced form for $\sum_{m=1}^{N_\phi} \pi_{sm}^{\phi} \sum_{i=1}^{N_z} \pi_{ji}^{z}$. The non-convexities in labour and capital adjustment costs introduce kinks in the value function. As shown by Cui (2014), the differentiability of the value function can be proved using methods from Clausen and Strub (2012). In Appendix C I show how this works and deal with the envelope conditions in detail. In any case, the numerical solution presented in appendix B does not use the optimality of the first order conditions.

[23] Appendix C deals with this issue and shows how the FOCs can be rewritten in terms of marginal values of capital and labour that satisfy the envelope condition, without loss of generality.

shadow value of wealth can be different than 1, as shown in the FOC for dividends:

$$1 - 2\kappa\left(d_t - \overline{d}\right) = \lambda_t \tag{12}$$

On one hand, it limits the room for issuing equity in face of negative shocks, therefore inducing firms to accumulate cash instead. On the other hand, it also induces firms to retain cash instead of distributing it as dividends, after a positive productivity shock. Both effects go in the same direction, implying that cash is used as a tool to move resources from one period to the other. This feature has been documented empirically by Dittmar and Duchin (2011). In the quantitative part of the paper I will show that, in absence of inter-temporal substitution in the savings decision, relevant empirical moments of UK firm-level data would be missed by the model. Intertemporal dividend decisions spill over also to capital and labour optimal decisions. The dividend cost implies that the Tobin's Q for capital fluctuates around 1 even without partial irreversibility. Moreover, firms cannot finance freely additional capital through equity issuance, and thus easily circumvent the financial constraint. This possibility would generate a counter-factual increase in investment during credit crunches, as shown in section 5.

Finally, labour adjustment costs have two opposite effects, which show up in the optimal decision for labour:

$$\lambda_t \omega z_{j,t} k_t^{\nu} n_t^{\omega-1} - \beta E_t \lambda_{t+1} AL_n\left(n_t, n_{t+1}\right) - \overline{w}\mu_t = \lambda_t\left[\overline{w} + AL_n\left(n_{t-1}, n_t\right)\right] \tag{13}$$

On one hand, hiring costs induce labour hoarding and make it less likely for booming firms to face a binding collateral constraint, because they increase the marginal cost of labour. On the other hand, firing costs imply that firms "on the way down", those that face negative shocks, may fear to be at the binding collateral constraint. Intuitively, the firing cost reduces the possibility of cutting labour. This affects the expected marginal benefit of having an additional worker next period, $E_t \lambda_{t+1} AL_n\left(n_t, n_{t+1}\right)$. The dividend cost affects labour optimal choices through the budget constraint; in turn, currently unconstrained firms may have incentives to adjust labour to changes to credit conditions.

The precautionary mechanism shows up in full if all the real frictions are included simultaneously. To show this, I use a version of the model with no real frictions and then discuss the role played by each of them. Equations 14-16 show the optimal decisions for capital, cash and labour in this case:

$$\beta E_t \left[\nu z_{t+1} k_{t+1}^{\nu-1} n_{t+1}^{\omega} + 1 - \delta_k\right] + \beta E_t \phi_{t+1}(1-\delta_k)\mu_{t+1} = 1 \tag{14}$$

$$\beta E_t \left[1 + \mu_{t+1}\right] = 1 - \psi_t \tag{15}$$

$$\omega z_{j,t} k_t^{\nu} n_t^{\omega-1} = \overline{w}\left[1 + \mu_t\right] \tag{16}$$

The timing assumption of the working capital constraint potentially gives rise to precautionary cash holdings by itself. Even in this version of the model, optimal capital and cash decisions still depend on the expectation of a binding financial constraint next period. Nevertheless, these expectations do not feed into the hiring decision through the budget constraint, because the absence of dividend costs implies that $\lambda_t = 1$. In other words, labour decisions of credit-unconstrained firms - for which $\mu_t = 0$ - are not affected by aggregate credit conditions when there are no real frictions. The dividend cost by itself is not enough to induce forward-looking hiring decisions: the combination of labour adjustment costs and dividend rigidity is required. Further introducing capital adjustment costs induces firms to tilt their portfolio allocation towards cash. Whether there will be firms, in a model without real frictions, for which the marginal benefit of holding cash is greater than the marginal return on capital is a quantitative question. I will show in section 5 that, for reasonable calibrations, the absence of real frictions implies that firms will strictly prefer to hold zero cash reserves. In this setting, firms hold cash only when ϕ approaches 0 and, effectively, the financial constraint resembles a cash-in-advance constraint.[24]

4 Quantitative exploration

This section considers a quantitative version of the theoretical framework, in order to investigate the effect of an exogenous tightening of the collateral constraint. The model is calibrated to the UK economy, using aggregate and microeconomic moments. In the steady state, the model is able to match a set of additional moments not explicitly targeted. I then show the effects of a credit supply shock, in the form of a drop in ϕ, both in terms of aggregate dynamics and microeconomic forces driving them.

[24]The same result could also be obtained with strong assumptions as full depreciation of capital. This, however, would not allow to match the calibration targets outlined in section 4 and 5.

4.1 Calibration

The model is parameterised so that the stationary equilibrium matches relevant aggregate and firm-level moments in the UK.[25] I set the time period to a quarter. Eight parameters are calibrated simultaneously using a simulated method of moments. The upper panel of Table 1 summarizes the parameter values that result from the calibration, whereas Table 2 compares empirical and model-generated moments. Most of the empirical moments are obtained using UK firm-level balance sheets from the FAME dataset, and are averages of the pre-crisis period 2004-2006.[26] The aggregate cash to assets ratio is defined as the aggregate sum of cash holdings over the aggregate sum of total assets. This is much smaller than the average cash ratio across firms, because of the relative skewness of the two distributions. Since precautionary cash holdings are a key feature of the model, I also target the cross-sectional standard deviation of firm-level cash ratios. In section 5 I will show how versions of the model that do not entail the full precautionary mechanism are not able to generate the skewness and the right tail of the cash ratio distribution. The dividend cost is disciplined by matching the average dividend payout ratio across firms and the cross-firm correlation between dividends and size. In absence of it, the correlation would be much higher. In the following subsection I explore other possibilities to calibrate this friction. To my knowledge, there is no reference for a direct comparison of this parameter. In a representative firm setting, Jermann and Quadrini (2012) calibrate κ to 0.146. Firms' discount factor is an important parameter since it affects the extent to which the collateral constraint is binding; a low enough value ensures that at least some firms do not accumulate enough assets to self-finance. As previously mentioned, this is disciplined by the amount of savings in the economy. Firms' capital intensity is disciplined by the sales to tangible assets ratio, an accounting measure of the ability of a firm to generate sales from its existing tangible assets.[27] Finally, the labour exponent of the production function is calibrated to match the aggregate labour share. The resulting decreasing returns to scale are close to the value calibrated by Khan and Thomas (2013).

Idiosyncratic productivity is assumed to follow a AR(1), which is then discretized using Tauchen and Hussey (1991) method to obtain $\left(\pi^z_{ij}\right)_{i,j=1}^{N_z}$. Its persistence is mainly identified by the autocorrelation of investment ratios. Khan and Thomas (2013) use the same strategy and estimate a similar value. It is relevant to mention that this autocorrelation would be negative in absence of partial irreversibility of capital.

[25] The stationary equilibrium in the steady state is defined as follows: I solve the firm's problem allowing for aggregate uncertainty and obtain the individual's time-independent decision rules. Then I set $\phi = \phi_H$ and compute the invariant distribution of firms over (m, k, n, z).

[26] See Appendix A for a description of the data and Appendix B for the details of the calibration strategy and the numerical method.

[27] The direct counterpart in the model is given by the production output to capital ratio.

Table 1: Parameter values

Calibrated values		
β	0.98	Firm discount factor
$\overline{d}$	0.465	Dividend target
ω	0.762	Exponent on labour in production function
ν	0.14	Exponent on capital in production function
κ	0.05	Dividend rigidity cost
ρ_z	0.928	Quarterly persistence of idiosyncratic productivity
σ_z	0.08	Quarterly standard deviation of innovations to idiosyncratic productivity
ϕ_H	0.55	Steady state credit tightness
Pre-defined values		
$\overline{w}$	1	Wage (normalisation)
δ_k	0.0375	15% annual depreciation of capital stock (Riddick and Whited, 2009)
χ	0.072	Per worker hiring/firing cost in % of annual wage bill (Bloom 2009)
δ_n	0.025	UK (ONS) average quarterly voluntary job separation rate 1996-2007
ϑ	34%	Resale loss of capital in % (Bloom, 2009)
Aggregate credit shock		
π^{ϕ}_{HH}	0.976	Quarterly transition probability of remaining in high ϕ
π^{ϕ}_{LH}	0.214	Quarterly transition probability from low to high ϕ
ϕ_L	0.495	Tight credit conditions to match drop in short-term loans (UK FAME)

Table 2: Model fit

Targeted Moments	Model	Data	Source
Aggegate labour share	0.74	0.70	Fernandez-Corugedo et al. (2012)
Aggregate cash-to-assets ratio	0.05	0.05	FAME
Average cash-to-assets ratio	0.17	0.18	FAME
Average tangible assets to sales ratio	0.75	0.77	FAME
Average dividend payout ratio	0.09	0.08	FAME
Cross-firm standard deviation of cash ratios	0.21	0.22	FAME
Autocorrelation of investment ratios	0.21	0.17	FAME
Cross-firm correlation between dividends and size	0.47	0.43	FAME

Non-targeted Moments	Model	Data	Source
Cross-firm correlation between cash ratio and capital-labour ratio	-0.15	-0.15	FAME
Autocorrelation of cash ratios	0.80	0.87	FAME
Scaled average volatility of dividends	1.63	0.44	FAME
Average cash ratio by quartile of size:			
≤25th percentile	0.48	0.31	FAME
>25th & ≤50th percentile	0.05	0.18	FAME
>50th & ≤75th percentile	0.04	0.14	FAME
>75th percentile	0.04	0.11	FAME

Notes: FAME data are averages for the period 2004-2006. Same data trimming as Figure 1 applies. The panel is balanced, for consistency with the model that does not account for entry and exit. Empirical payout ratio is calculated as the ratio between dividends and turnover, whereas in the model it is dividends over output. Investment ratio for a firm j at time t is defined as $\frac{CAPEX_{j,t}}{\alpha k_{j,t}+(1-\alpha)k_{j,t-1}}$, with $\alpha = 0.5$, where CAPEX are capital expenditures as recorded in the Income statement, and k is fixed assets as recorded at balance sheet. Autocorrelation is one year. Size is measured in number of employees in the data and labour n in the model. Empirical capital to labour is the ratio of tangible assets over number of employees. The scaled average volatility of dividends is calculated in the model as follows: I simulate the stochastic steady state, with ϕ_H, for a large number of quarters. Then dividends are annualised, and for each firm is calculated the standard deviation over time. I take the cross-sectional average of these standard deviations, and scale it by the cross-sectional average level of dividends. In the data, I follow the same approach, before 2006. Extending the sample to 2013 increases the moment to 0.90.

The model entails aggregate uncertainty with respect to the credit tightness ϕ, whose stochastic process is discretized using a 2-states Markov chain. The aggregate cash ratio is the most informative moment for the calibration of the ordinary credit tightness ϕ_H. As will be discussed in the following sections, a credit shock consists of a fall in ϕ. Agents form their expectations over future credit conditions according to the transition probability matrix. I define a credit shock as a 10% drop in ϕ. This allows the model to match the HP-detrended reduction in aggregate short-term loans held by UK firms in FAME data, between 2007 and 2009. The issues related with the estimation of the credit shock process are particularly relevant for the UK, where data on aggregate financial conditions[28] start in Q1 2007. In terms of timing, both survey and lending measures suggest that firms' credit conditions start to deteriorate only at the beginning of 2009.[29] I set the transition probabilities such that the credit tightening lasts on average 13 months and occurs every 10 years. The length of the credit tightening is in line with a various range of financial indicators in the UK. The frequency may seem high, especially if we interpret it as a financial crisis. Nevertheless, the financial friction in my model restrains the ability to borrow within the period, resembling what is normally referred to as line of credit, as in Bacchetta et al. (2015). Shocks affecting the supply of this form of liquidity are likely to happen much more frequently than a full-blown financial crisis. Moreover, relatively frequent credit tightening episodes should dampen the disruptive effects of a credit supply shock, since firms have an additional incentive to save ahead of the crisis and therefore are better equipped to face episodes of credit crunch.

The exogenous job separation rate δ_n is directly taken from ONS data on UK average quarterly voluntary job separation rate between 1996 and 2007. The wage is normalised to 1. Per worker hiring and firing costs are symmetric and equal to 1.8% of the annual wage bill. This estimate, as well as the capital adjustment cost, is taken from Bloom (2009). Ideally, these values should be estimated to the UK data as well. In Appendix D I explore the effects of using different estimates available in the literature, and show that Bloom (2009) values allow to match over-identifying UK empirical moments.

4.2 Steady state

The stationary distribution of firms in the stochastic steady state, with $\phi = \phi_H$, is in line with a number of UK empirical moments that are not explicitly targeted in the calibration. The lower panel of Table 2 shows some of them. The negative cross-firm correlation between cash ratio and capital to labour ratio informs us about important

[28] Bank of England Credit Conditions Survey.

[29] Among others, Ivashina and Scharfstein (2010) document that US firms drew down on their lines of credit during 2008. Bacchetta et al. (2015) build a liquidity measure consistent with a fall in firm's access to liquidity starting at the beginning of 2009.

model dynamics. Consider a firm with a low capital to labour ratio: this implies that it needs to finance a large wage bill, especially relative to the available collateral. Hence, the firm has incentives to shift towards a more cash-intensive portfolio.

The very positive autocorrelation of cash ratio is, in the model, a by-product of the dividend cost, which makes savings decisions intertemporal. This implies a pronounced stickiness of cash holdings, which would be otherwise missed in models that do not feature the dividend cost. As previously mentioned, the dividend cost breaks down the relationship between dividends and size, allowing the model to match the data. Moreover, it reduces the volatility of dividends over time. Following a positive productivity shock, for instance, a firm would like to distribute more dividends to the shareholders. The dividend cost limits the amount of additional dividends the firm will pay out. With multiple dividend targets, this effect would be milder, because the target would adjust together with the productivity shock, making the cost of deviation smaller and leading to larger fluctuations in dividends. I show that, even with a single dividend target, the model overshoots the volatility of dividends over time compared to the data. I compute the standard deviation of dividends over time for each firm, then I take the average of these volatilities across firms and scale it by the cross-sectional average of dividends, to allow for comparability between the model and the data.[30] As will be shown in Table 4, alternative versions of the model without the dividend cost predict a much higher volatility of dividends.

Finally, smaller firms have relatively higher cash ratios. In the data, the firm distribution of employees before 2007 has been categorized in four different quartiles, and for each of them the average cash ratio has been computed. The negative correlation is present also in the model, although this is overly driven by very small firms. Intuitively, small firms have little collateral and are likely to face an increasing schedule in employment growth. Hence, they face a tradeoff between expanding their capital stock together with labour or accumulating internal cash as a financing tool. While they are small, the low pledgeability of capital makes them to lean towards cash. As they grow, they get more easily access to external credit and thus start reducing their stock of cash.

4.3 The aggregate response to a credit shock

I simulate the aggregate dynamics of the model following a tightening of the collateral constraint. This takes the form of a 10% drop in ϕ, as explained in section 4.1, and can be interpreted as an increase in the haircut applied on collateral.

Figures 4a-4d show the aggregate dynamics of a credit tightening that lasts 5 quar-

[30] An alternative approach involves computing the coefficient of variation at the firm-level and then taking the cross-sectional average. This delivers similar results.

Figure 4: The effects of a credit supply shock

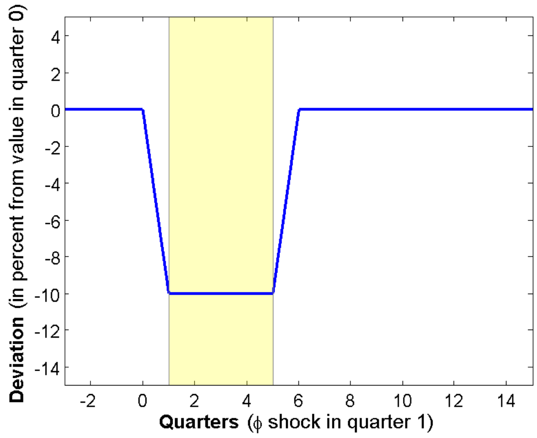

(a) Credit shock

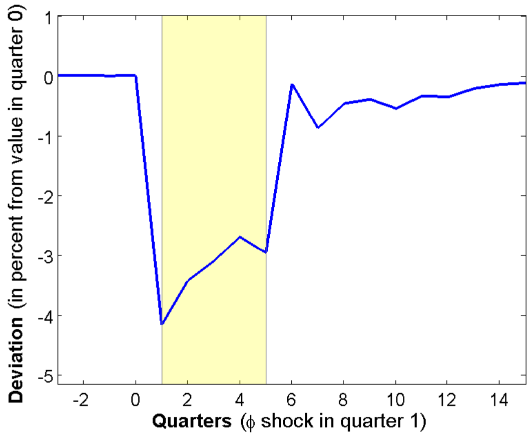

(b) Aggregate Employment

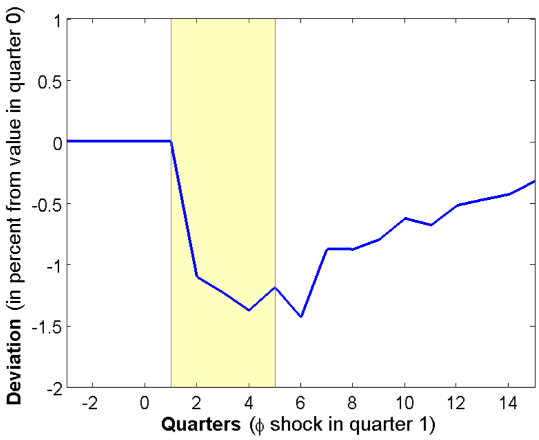

(c) Aggregate Capital

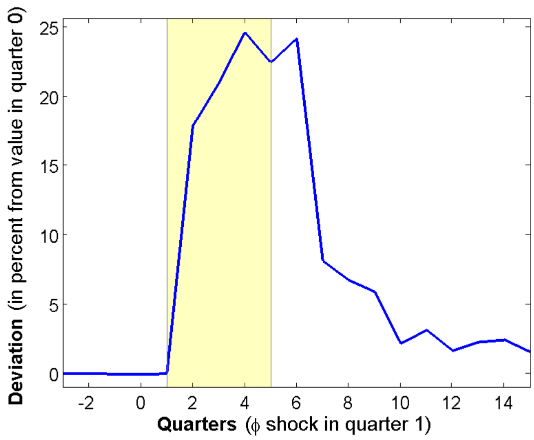

(d) Aggregate Cash

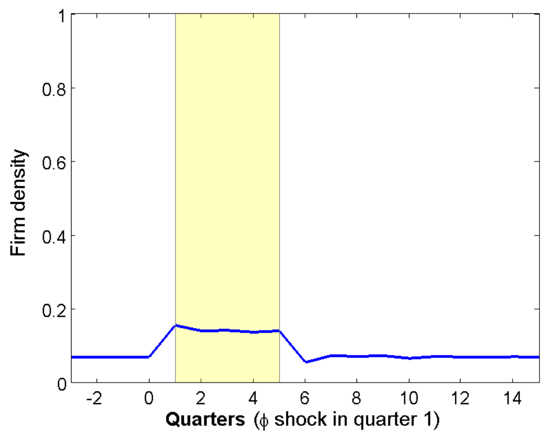

(e) Share of constrained firms

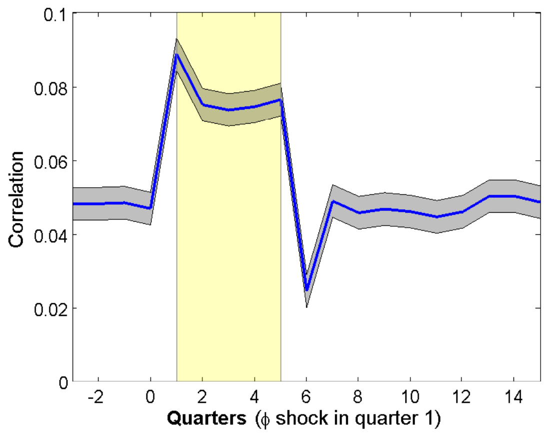

(f) Time-series of cross-firm correlations between lagged cash ratio and employment growth

Note: Simulations over 200,000 firms for 700 quarters, discarding the first 100 quarters. Simulations have been repeated 100 times. The IRFs show the mean response (across simulations) of a credit shock lasting 5 quarters, generated by a Monte Carlo simulation over the combined transition matrix for aggregate and idiosyncratic shocks. Yellow areas indicate credit tightening periods. Shaded grey bands indicate 95% confidence intervals. Constrained firms are those for which equation (7) binds.

ters.[31] A credit crunch has two main effects. First of all, it causes a substitution from capital to cash. During credit tightening periods, capital is less worth in terms of collateral. For this reason, firms switch from external financing to internally generated liquidity. Credit shocks enlarge capital inaction regions typical of non-convex adjustment costs; this implies a very sluggish recovery in capital and, conversely, aggregate cash levels well above steady state conditions for many quarters after the end of the credit crunch.

The second effect is a sizeable fall in aggregate employment, which remains subdued throughout the credit tightening period. Less available credit implies that some firms are not able to finance the same wage bill as before, and thus they have to reduce their workforces. Aggregate employment rebounds as soon as ordinary credit conditions are restored, although it takes many quarters to completely come back to pre-crisis levels.

As a crucial result in this paper, most of the aggregate dynamics are driven by unconstrained firms, which respond to changes in credit conditions by anticipating the possibility of future constraints. In the next subsection I will show the quantitative importance of this mechanism, and how precautionary cash holdings affect the transmission mechanism of credit shocks to the labour market.

Firm heterogeneity and the interaction between financial and real frictions lie behind this result. The financial constraint binds occasionally both over time and across firms. As shown in Figure 4e, the share of firms facing a currently binding credit constraint increases during credit tightening periods, but never exceeds 16%. As a direct consequence of this, in the aggregate, the available financing funds are largely in excess of firms' financing needs. This result is in line with empirical findings by Shourideh et al. (2012), as mentioned in the introduction.

Figures 5a-5e show that the effects of a credit shock in the model replicate fairly well key aggregate dynamics in the UK.[32] This is especially true in the first year after the shock. The recovery of aggregate employment and output is faster in the model than in the data; this is probably due to the absence of demand effects that may propagate the effects of a credit shock. I also show that a model without real frictions is unable to generate the empirically observed aggregate dynamics. This model does not entail any precautionary mechanism; hence, no firm holds cash and the fall in aggregate employment is negligible. Moreover, this alternative version of the model is not able to predict a fall in investment, whereas the benchmark model can explain half of the empirically observed drop. In section 5 I investigate in further depth the quantitative importance of the

[31] An alternative approach consists of simulating many different economies where we allow the aggregate credit shock to evolve naturally after an initial one-period drop. The aggregate dynamics for this case are shown in Appendix E.

[32] Model-generated quarterly data have been annualised, for direct comparison with FAME data. Appendix B deals with time aggregation. The shock hits the economy in the first quarter of 2009, in line with data on firms' credit conditions and short-term lending to UK businesses.

Figure 5: Model performance to the UK economy

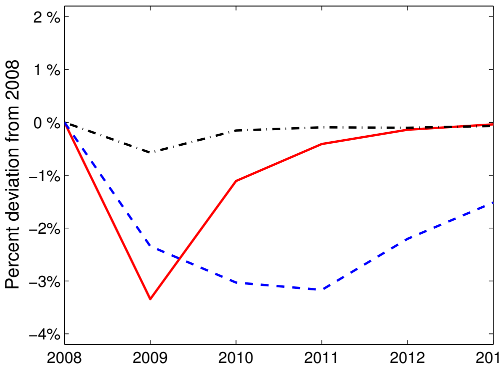

(a) Aggregate employment

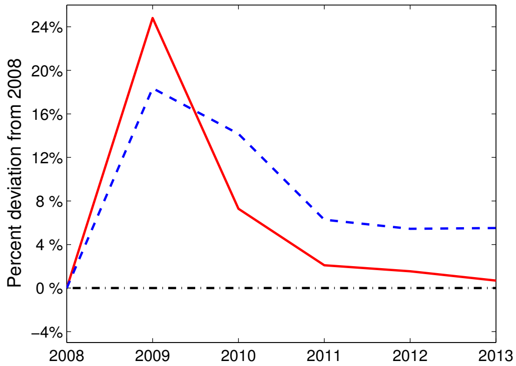

(b) Aggregate cash to assets ratio

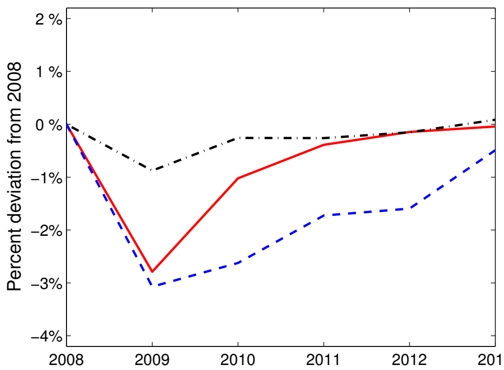

(c) Aggregate output

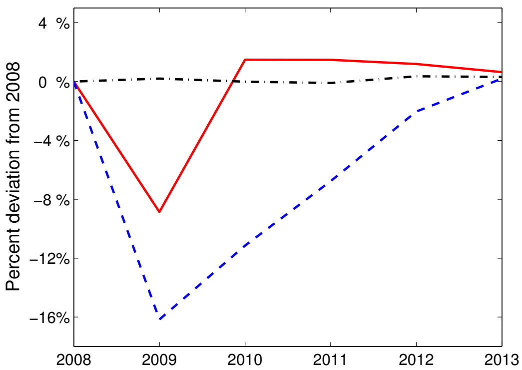

(d) Aggregate investment

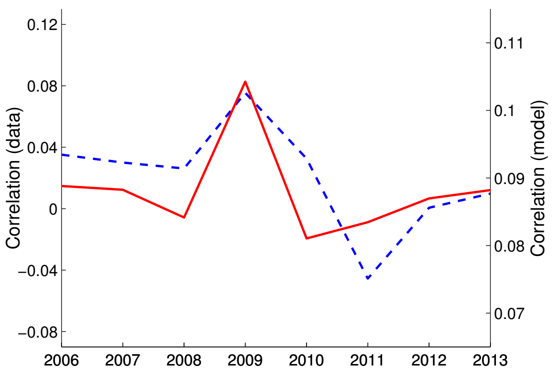

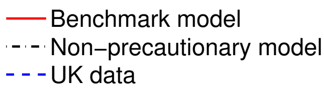

(e) Time-series of cross-firm correlations between lagged cash ratio and employment growth

Note: Model-generated data refer to the same simulation outlined in Figure 4. Quarterly simulated data have been annualised using standard accounting techniques, as explained in Appendix B. The credit shock hits the economy in the first quarter of 2009. In the "Non-precautionary model", χ, ϑ and κ have been set to 0, while the remaining parameters are the same as in Table 1. UK data for aggregate employment refer to employment rate (aged 16-64) from the Office for National Statistics (ONS). Data for panel b are from the FAME dataset, where previous notes apply. UK GDP in panel c is chained volume, from the ONS. Data has been detrended with HP filter, over the period 1950-2014. HP-filtered data are scaled by GDP in 2008. UK aggregate investment is Business Investment (ONS). In panel e, it is not possible to compute the correlation for the non-precautionary model, since the cash ratio is always equal to 0.

precautionary channel, through various counterfactual analyses. A version of the model with only partial irreversibility of capital also generates aggregate investment dynamics at odds with the data, besides implying wrong predictions at the micro-level.

4.4 Micro-level drivers of a credit shock

Besides correctly predicting the evolution of key aggregate variables, the model is able to replicate the dynamic evolution of the cross-firm correlation between lagged cash ratio and employment growth, as shown in Figure 4f. The model mechanisms help rationalise this result. Firms face a tradeoff when choosing their cash ratio. On one hand, they may need to reduce their hiring to finance an increase in cash, because cutting on dividends or capital investment is limited by the associated frictions. Moreover, reducing the capital stock exerts further negative pressure on labour through the production function. On the other hand, more cash alleviates the extent to which the credit constraint binds and helps sustaining employment growth. In normal times, these two effects generally offset each other, as shown in Figure 4f. After a credit supply shock that dries up liquidity, however, cash becomes a relatively more valuable source of financing. Cash-intensive firms are associated with higher employment growth throughout the crisis.

Figure 5e compares the annualised simulation for the benchmark model and the data.[33] The dynamic pattern is qualitatively similar, although time aggregation negatively affects the quantitative performance of the results. As mentioned for the aggregate time series, the effects of the crisis are more short-lived than in the data, and thus correlation remains high only for one year. During the recovery, model and data diverge the most. As in the data, the correlation drops below pre-crisis level, suggesting that the expansion in employment is driven by firms that were not able to increase cash during the credit tightening periods and benefit from the restored credit conditions. Differently from the data, however, the model-generated correlation remains positive. This may be due to the fact that the model does not generate a sufficient portion of very disrupted firms, which are restrained in the possibility of hoarding cash during the credit tightening period.

The key mechanism that characterizes this paper is a precautionary behaviour in anticipation of future financial constraints. In order to evaluate its quantitative importance, I consider the simulation described in the previous section and classify the firms in four groups, depending on whether they faced a currently binding credit constraint in this period and/or in the previous. Firms that do not face a binding constraint are loosely labelled as unconstrained, although they will be clearly affected by financial frictions and more so the closer they are to the binding constraint. Figure 6 shows, for each period,

[33] In a model with no real frictions, no firm holds cash, hence we cannot compute the cross-firm correlation.

Figure 6: The importance of unconstrained firms

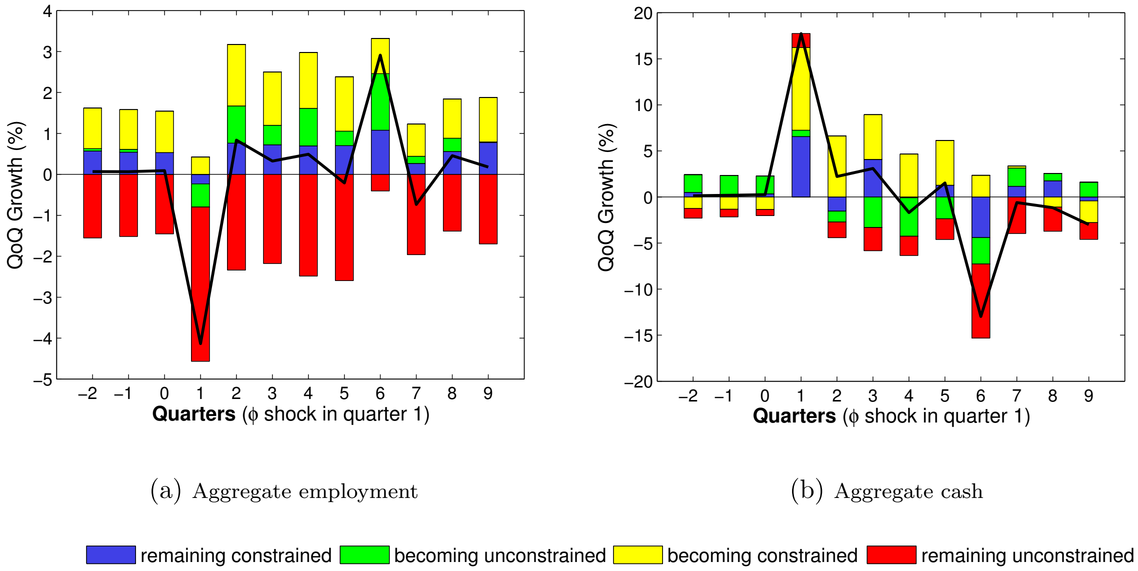

(a) Aggregate employment

(b) Aggregate cash

Note: Each period t, firms are classified in 4 groups: "remaining constrained" if they faced a binding financial constraint in t and $t-1$, "becoming unconstrained" if they moved from constrained to unconstrained, "becoming constrained" from unconstrained to constrained, "remaining unconstrained" if they remained unconstrained. Each bar denotes the contribution of the group to growth. The economy starts at the stationary distribution and experiences a credit tightening in quarter 1, lasting 5 quarters. Previous notes on the simulation apply.

the contribution of each group to the growth of aggregate employment and cash. In normal times constrained firms are booming and hoard cash to sustain their employment growth. When a credit shock hits the economy, all the four groups are affected. Unconstrained firms react to changes in credit conditions by cutting employment more than in normal times and accumulating cash. They account for more than 60% of the total fall in aggregate employment.

4.5 The credit-constrained firms

The identification of credit constraints is still a debated topic in the empirical literature, and it is not clear whether it is possible to find observables that unambigously help the identification of credit-constrained firms. The structural model outlined in this paper predicts that credit-constrained firms are generally productive, large and illiquid. The first two characteristics directly stem from the working capital constraint, which by construction is more binding for firms that have a large wage bill to finance. The relationship between cash ratio and the probability of being constrained is less clear ex ante. On one hand, higher cash implies that firms are less likely to face a binding financial constraint. On the other hand, this is normally associated with small firms, which do have little wage bill to finance but also little collateral to pledge. As mentioned in the previous sections

Table 3: The performance of financial constraint proxies in the model

	% incorrectly classified as	
Proxy	constrained	unconstrained
Size (employees): below median	96.8%	11.2%
Size (total assets): below median	96.2%	10.4%
Leverage ratio: below median	85.7%	1.4%
No dividend payment	88.0%	6.1%
Size (employees): bottom tercile	97.6%	9.7%
Size (total assets): bottom tercile	97.9%	9.6%
Leverage ratio: upper tercile	81.9%	1.4%
All firms	92.9%	7.1%

Note: Stochastic steady state of the model, with steady state credit tightness ϕ_H. The first column lists some proxies for financial constraints used in the empirical literature. Size as a proxy for financial constraints has been used in a wide range of empirical papers and the cutoff is generally set to the median of the distribution (Silva and Carreira, 2012). Fazzari et al. (1988) suggest that firms not paying dividends are financially constrained. The results reported in the second column identify the share of firms that would be constrained according to the proxy, but which are actually unconstrained. Conversely, in the third column, the share of firms that would be unconstrained in the proxy, but which are actually constrained in the model.

and shown in the quantitative analysis, being constrained is not a clear-cut concept in the model. For this reason, there are often non-linear and time-varying correlations between firm's characteristics and the distance from the binding constraint.

Coming up with a simple proxy which identifies financial constraints in the model is difficult. To show this, I run the following experiment. I take different proxies for financial constraints normally used in the empirical literature and ask: if we were to classify firms according to each of these proxies, how many would be correctly identified as financially constrained in the model? Table 3 shows the performance of these proxies.[34] Only 3.2% of the firms that would be categorized as constrained according to their size, are actually facing a binding collateral constraint in the model, while 96.8% are not. In other words, categorising firms by size would identify financial constraints more poorly then just looking at all firms. The validity of size as a good measure of financial constraint has been challenged by recent studies, which tend to prefer age as a better predictor of financial frictions. Given the precautionary mechanism outlined in the previous sections, the fact that large firms are generally more credit-constrained in the model does not

[34]Clearly, this exercise is inherent to the specific features of the model. Nevertheless, the objective is to show that even in a structural model calibrated to firm-level data, identification of financial constraints is difficult.

necessarily mean that they will be more affected by a credit supply shock. Indeed, In appendix E I show that large firms reduce employment only marginally more than small firms in the wake of a credit shock. Fazzari et al. (1988) suggest that firms not paying dividends are financially constrained. This proxy performs slightly better in the model, although dividend smoothing may be associated to non-financial factors. Leverage ratio is also a better proxy for financial constraints in the model. Finally, not only simple proxies of financial constraints perform poorly in the model, but even multi-dimensional proxies cannot identify credit-constrained firms. A probit regression for the probability of being constrained on all the firms' state variables delivers a pseudo-R^2 of only 81%.[35]

5 The importance of the precautionary mechanism

Previous sections have already shown how the precautionary channel allows the benchmark model to generate predictions in line with the data. In this final section, I formalize this statement by considering two variations of the benchmark model shown in section 3.

I start by shutting down all the real frictions at the same time. In this version of the model, firms generally find it optimal to invest in capital and set their cash reserves to 0, as shown in Figure 5. The aggregate cash ratio is matched only if ϕ_H is driven to 0.32. At such low value, external credit is so limited that some firms prefer to resort to internally generated liquidity.[36] The second version of the model entails only partial irreversibility of capital, besides the working capital constraint.

Figure 7 shows the same quantitative exercise of section 4.3, where a credit tightening shock hits the economy for five quarters, for the benchmark model and three variations. When the economy does not face any real friction, the fall in aggregate employment is very mild and no firm holds cash, as shown by the dash-dotted line. When ϕ_H is recalibrated to a sufficiently low value, this version of the model generates a large drop in employment and a spike in cash. Capital, however, remains fairly stable.

Two counteracting forces are at work: on one hand, absent any equity issuance cost, firms can respond to changes in credit conditions by issuing equity to fund capital investment, thus circumventing the financial friction. On the other hand, constrained firms

[35] McFadden's pseudo-R^2 does not exactly mean the proportion of variance of the dependent variable explained by the regressors, as in OLS. Hence, interpretation should be taken with caution. Nevertheless, this illustrative example still makes the point that the identification of financial constraints is hard when non-linearities and non-convexities matter.

[36] Extreme calibrations of the exponents of production function and the discount factor can generate the empirically observed aggregate cash ratio but miss other relevant targets as the sales to tangibles ratio and the aggregate labour share. An alternative approach departs from the assumption that cash earns no interest; in this setting, cash is held via a risk-free bond as in Riddick and Whited (2009). Even with an annual interest rate on cash of 12%, however, the frictionless model does not match the aggregate cash ratio.

Figure 7: Impulse response functions to a credit supply shock - alternative versions of the model

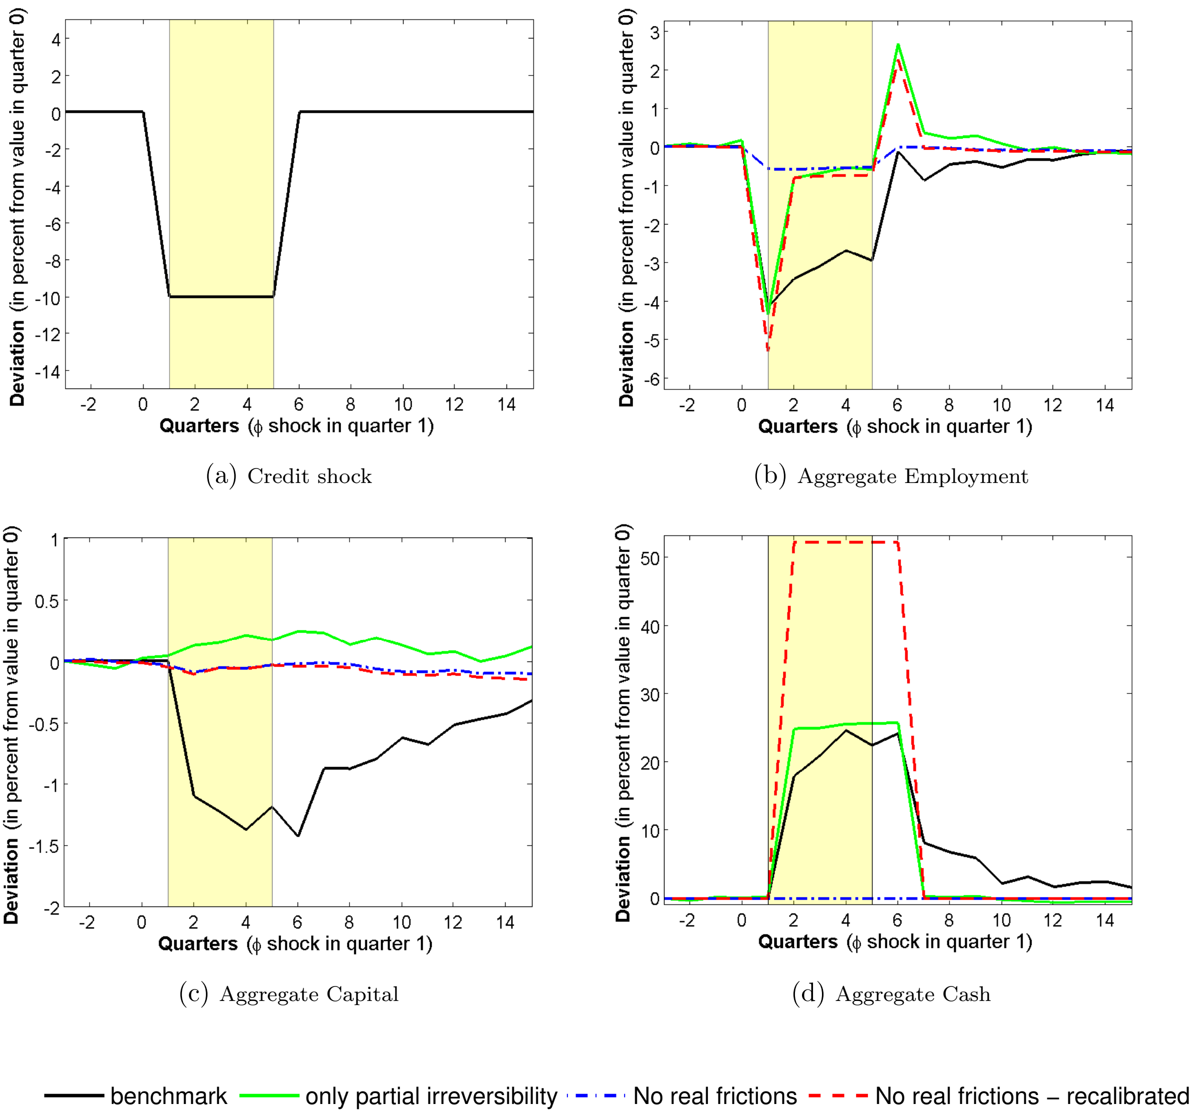

(a) Credit shock

(b) Aggregate Employment

(c) Aggregate Capital

(d) Aggregate Cash

Note: The economy starts at the stationary distribution and experiences a credit tightening in quarter 1, lasting 5 quarters. Previous notes on the simulation apply. χ and κ have been set to 0 in all versions of the model except the benchmark. In the version of the model with only partial irreversibility, the quarterly interest rate on cash has been recalibrated from 0% to 0.68% to match the aggregate cash ratio. In the models without real frictions, also ϑ has been set to 0. The recalibrated version of the model without real frictions requires $\phi_H = 0.32$ to match the aggregate cash ratio.

Figure 8: Model with only partial irreversibility of capital

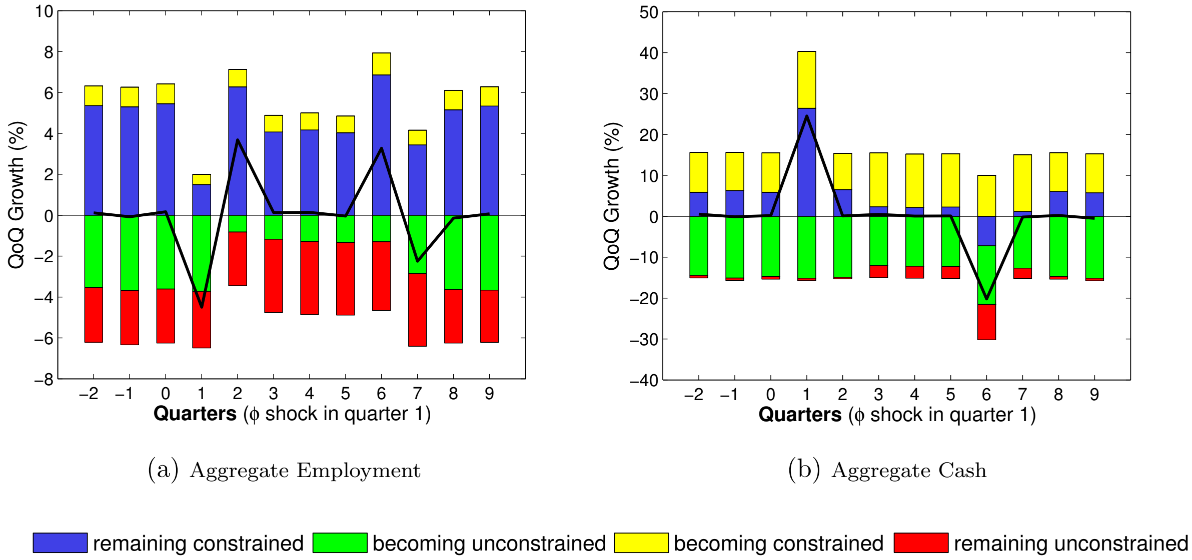

(a) Aggregate Employment

(b) Aggregate Cash

Note: Each period t, firms are classified in 4 groups: "remaining constrained" if they faced a binding financial constraint in t and $t-1$, "becoming unconstrained" if they moved from constrained to unconstrained, "becoming constrained" from unconstrained to constrained, "remaining unconstrained" if they remained unconstrained. Each bar denotes the contribution of the group to growth. The economy starts at the stationary distribution and experiences a credit tightening in quarter 1, lasting 5 quarters. Previous notes on the simulation apply. χ and κ set to 0 and quarterly interest rate on cash recalibrated to 0.68%.

reduce labour upon impact, and this exerts a negative pressure on capital through the production function. Which effect dominates is a quantitative question: the benchmark calibration suggests that the two effects offset each other. A model with only partial irreversibility of capital generates similar aggregate dynamics. In fact, capital even slightly increases, because more firms refrain from selling capital at a lower resale price. In both alternative versions of the model, the response of aggregate employment is very short-lived. Without labour and dividend costs, firms can adjust rapidly to the change in credit conditions, mainly because labour is not a state variable. Most of the aggregate dynamics are driven by constrained firms, as shown in figure 8 for the model with only partial irreversibility of capital.

The two alternative versions of the model also generate microeconomic predictions that are at odds with the data, as shown in Table 4. First of all, they are not able to generate the empirically observed right tail of the cash ratio distribution. In the model with only partial irreversibility, for instance, the average cash ratio is only 4.1%, far away from 18% observed in the data. This version of the model also predicts that larger firms are more cash intensive, against what observed in the data. Finally, both alternative versions of the model fall short in terms of autocorrelation of cash ratios. The dividend cost present in the benchmark model helps in this sense, since it makes cash a useful tool

Table 4: Fit of alternative versions of the model

Moments	Data	Benchmark model	No real frictions, recalibrated ϕ	only partial irreversibility
Average cash-to-assets ratio	0.18	0.17	0.06	0.04
Cross-firm standard deviation of cash ratios	0.22	0.21	0.06	0.07
Cross-firm correlation between cash ratio and size	-0.04	-0.39	-0.20	0.30
Cross-firm correlation between cash ratio and capital-labour ratio	-0.15	-0.15	0.32	-0.42
Autocorrelation of cash ratios	0.87	0.80	0.38	0.44
Scaled average volatility of dividends	0.44	1.63	3.35	2.27

Note: Stochastic steady state of the model, with credit tightness ϕ_H. Size is measured in number of employees in the data and labour n in the model. Notes on figure 7 and table 2 apply.

to move resources between periods. Moreover, labour adjustment costs are quantitatively important for the trade-off between liquidity and size. Finally, including a dividend cost with a single dividend target reduces the volatility of dividends over time. Versions of the model without this friction predict much larger fluctuations.

6 Conclusions

The quantitative role of firm-level financial frictions and credit supply shocks in affecting the real economy has been recently object of debate in the literature. In this paper I argue that a precautionary mechanism, which induces firms to respond to changes in credit conditions in anticipation of future idiosyncratic shocks, plays a quantitatively important role. This paper entails two main contributions. Using UK firm-level data, I document that firms accumulated cash during the last recession and cash-intensive firms decreased their workforces by less. Motivated by these facts, I build a heterogeneous-firm model where precautionary cash holdings arise endogenously from the interaction between real and financial frictions. The model is disciplined by UK data and its ability to replicate empirical facts relies on the precautionary channel. Credit tightening implies a sizeable fall in aggregate employment and a substitution from capital to cash. As a key result, these aggregate dynamics are mainly driven by firms not facing a currently binding constraint, who behave pre-emptively in anticipation of future shocks.

References

[1] Altinkilic, O., and R. S., Hansen, 2000, "Are there economies of scale in underwriting fees? Evidence of rising external financing costs", *Review of Financial Studies*, 13(1): 191–218.

[2] Angelini, P. and A. Generale, 2008, "On the evolution of firm size distributions", *American Economic Review*, 98(1): 426–438.

[3] Bacchetta, P., Benhima, K., and C. Poilly, 2015, "Corporate Cash and Employment", University of Lausanne.

[4] Bassetto, M., Cagetti, M., and De Nardi, M., 2015, "Credit crunches and credit allocation in a model of entrepreneurship", *Review of Economic Dynamics*, 18(1): 53–76.

[5] Bates, T. W., Kahle, K. M., and Stulz, R. M., 2009 "Why do US firms hold so much more cash than they used to?", *The Journal of Finance*, 64(5): 1985–2021.

[6] Benmelech, E., Bergman, N. K., Seru, A., 2011, "Financing labor", *NBER Working Papers 17144*.

[7] Bloom, N., 2009, "The impact of uncertainty shocks", *Econometrica*, 77(3): 623–685.

[8] Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I. and S. J. Terry, 2012, "Really Uncertain Business Cycles" *NBER Working Papers 18245*.

[9] Buera, F. J., Jaef, R. N. F., Shin, Y., 2015, "Anatomy of a credit crunch: From capital to labor markets", *Review of Economic Dynamics*, 18(1): 101–117.

[10] Campello, M., Graham, J. and C. Harvey, 2010, "The real effects of financial constraints: evidence from a financial crisis", *Journal of Financial Economics*, 97(3): 470–487.

[11] Chodorow-Reich, G., 2014, "The employment effects of credit market disruptions: Firm-level evidence from the 2008-9 financial crisis", *The Quarterly Journal of Economics*, 129(1): 1–59.

[12] Corbae, D. and P. D'Erasmo, 2014, "Reorganization or Liquidation: Bankruptcy Choice and Firm Dynamics", presented at UCL Macro seminar, April 2014.

[13] Cui, W., 2014, "Delayed Capital Reallocation", Available at SSRN 2482431.

[14] Dittmar, A. K. and Duchin, R., 2011, "The Dynamics of Cash", Ross School of Business Paper No. 1138.

[15] Eisfeldt, A. L. and A. A. Rampini, 2006, "Capital Reallocation and Liquidity", *Journal of Monetary Economics*, 53(3): 369–399.

[16] Falato, A., Kadyrzhanova, D., and Sim, J. W., 2013, "Rising intangible capital, shrinking debt capacity, and the US corporate savings glut"' *Board of Governors of the Federal Reserve System (US)*, No. 2013–67.

[17] Fazzari, S. M., Hubbard, G. R. and B. C. Petersen, 1988, "Financing Constraints and Corporate Investment," *Brookings Papers on Economic Activity*, Economic Studies Program, The Brookings Institution, 19(1): 141–206.

[18] Fernandez-Corugedo , E., McMahon, M., Millard, S. and L. Rachel, 2011, "Understanding the macroeconomic effects of working capital in the United Kingdom", *BoE Working Paper 422*, Bank of England.

[19] Finocchiaro, D. and C. Mendicino, 2013, "Debt, Equity and Monetary policy".

[20] Han, S. and J., Qiu, 2007, "Corporate precautionary cash holdings", *Journal of Corporate Finance*, 13 (2007): 43–57.

[21] Hennessy, C. A. and T. M. Whited, 2007, "How costly is external financing? Evidence from a structural estimation", *The Journal of Finance*, 62(4): 1705–1745.

[22] Jermann, U., and V. Quadrini, 2012, "Macroeconomic effects of Financial shocks" *American Economic Review* 102(1): 238–271.

[23] Khan, A. and J. K. Thomas, 2013, "Credit shocks and aggregate fluctuations in an economy with production heterogeneity", *Journal of Political Economy*, 121(6): 1055–1107.

[24] Kiyotaki, N., and J. H. Moore, 1997, "Credit cycles" *Journal of political economy*, 105(2): 211–248.

[25] Kiyotaki, N., and Moore, J., 2012, "Liquidity, business cycles, and monetary policy", *NBER Working Papers 17934.*

[26] Lintner, J., 1956, "Distribution of Incomes of Corporations Among Dividends, Retained Earnings, and Taxes", *American Economic Review*, 46: 97–113.

[27] Macera, M., 2013, "Credit Crises and Private Liquidity: The Role of Household Debt"

[28] Mendoza, E., 2010, "Sudden Stops, Financial Crises and Leverage", *American Economic Review*, 100(5): 1941–1966.

[29] Moscarini, G. and F. Postel-Vinay, 2012, "The contribution of large and small employers to job creation in times of high and low unemployment", *American Economic Review*, 102(6): 2509–2539.

[30] Quadrini, V., 2011, "Financial Frictions in Macroeconomic Fluctuations", *Richmond Fed Economic Quarterly*, 97(3): 209–254.

[31] Riddick , L. A. and T. M. Whited, 2009, "The Corporate Propensity to Save," *The Journal of Finance*, 64(4): 1729–1766.

[32] Sharpe, S. A., 1994, "Financial market imperfections, firm leverage and the cyclicality of employment", *American Economic Review*, 84(4): 1060–74.

[33] Shourideh, A. and A. Zetlin-Jones, 2012, "External Financing and the Role of Financial Frictions over the Business Cycle: Measurement and Theory" In 2012 Meeting Papers (No. 321). *Society for Economic Dynamics.*

[34] Silva, F., and C. Carreira, 2012, "Measuring firms financial constraints: a rough guide." *Notas Economicas*, 36: 23–46.

[35] Spaliara, M. E., 2009, "Do financial factors affect the capital-labour ratio? Evidence from UK firm-level data", *Journal of banking and finance*, 33(10): 1932–1947.

[36] Svensson, L. E., 1985, "Money and asset prices in a cash-in-advance economy" *The Journal of Political Economy*, 93(5): 919-944.

A The data

The primary data source used in this paper is FAME dataset, gathered by Bureau van Dijk. It contains information on over 9 million companies in UK and Ireland, 2 million of which are in detailed format, over the period 2004-2013.[37] I restrict the dataset to UK only. A standard company report includes a balance sheet, profit and loss account, turnover, employees and industry codes.[38] In contrast to other datasets as US Compustat, 93% of the firms contained in the FAME dataset are non-publicly traded. This implies that there is a large number of small and medium-sized companies.[39] Since the model does not feature life cycle, I restrict the sample to a balanced panel; firms that have weakly positive observations for employment, cash and total assets are kept in the sample. Following the standard procedure employed in similar studies, I exclude from the sample firms with UK SIC code referring to "Financial and insurance activities". The final sample consists of 17,762 firms each year. The balanced panel may certainly bring along selection bias; for this reason, I check that the stylised facts shown in section 2 hold also with an unbalanced panel of firms. Although the reporting requirements slightly bias the sample towards large firms, the size distribution is much closer to the UK universe than a dataset with only publicly quoted firms. For instance, the median firm in 2006 had 77 employees. The sample is representative also in terms of aggregate dynamics. The evolution of aggregate employment, for instance, closely resembles the one for Non-financial corporations published by the UK Office for National Statistics. Cash is recorded in firm's balance sheets as *Bank Deposits*. This is probably a stricter definition than the one employed in Compustat, because it excludes short-term liquid assets as Money Market Mutual funds holdings. Nevertheless, it suits better with the model assumption that cash earns no interest. As I explain in Appendix E, results are robust to this extension. It may be argued that firms may adjust their stock of cash because of changes in taxation or repatriation motives (Foley et al., 2007). Eliminating the upper first percentile of the cash distribution affects only mildly the dynamics of the aggregate stock of cash and leaves all the micro-economic moments unaffected. Net job creation is defined as the difference in number of employees for a given firm from one year to the other. I define employment growth for a firm j at year t as $\Delta n_{j,t} = \frac{n_{j,t}-n_{j,t-1}}{\alpha n_{j,t}+(1-\alpha)n_{j,t-1}}$, with $\alpha = 0.5$. Moscarini and Postel-Vinay (2012) explain the advantages of this symmetric

[37] A maximum of 10 years data history can be downloaded at once. Companies are registered at Companies House in the UK.

[38] Some firms report also the Cash Flow statement. Moreover, the data includes detailed ownership and subsidiary information.

[39] Unlike in the US, UK firms have to disclose their accounts even when not traded on the stock market. Following the UK Companies Act 1985, large firms have to report detailed accounts, whereas medium-size companies do not have to disclose turnover details and small firms are required to submit only an abridged balance sheet.

approach, which bounds employment growth between -2 and 2. Moreover, it allows to make inference on employment growth even with an unbalanced panel, with firm entry and exit. As in Sharpe (1994), robustness checks involve excluding observations associated with net job creation more than 300% of previous year employment. The dividend payout ratio is defined as the ratio between dividends and turnover, as reported in the Income Statement. This strategy allows to have a direct comparison with the model. To avoid being affected by outliers, I discard payout ratios observations above 500%; this accounts for less than 0.3% of all firms. Finally, I define investment ratio using the same strategy as for employment growth. Investment ratio for a firm j at time t is defined as $\frac{CAPEX_{j,t}}{\alpha k_{j,t}+(1-\alpha)k_{j,t-1}}$, with $\alpha = 0.5$, where CAPEX are capital expenditures as recorded in the Income statement, and k is fixed assets as recorded at balance sheet.

B Numerical Method

The firm's problem is solved with Value function iteration. The state dimension for (n_{t-1}, m_t, k_t, z_t) is discretized over a equidistant meshed grid. The choice dimension is instead specified over a much finer grid. The choice grid for capital always comprises the inaction decision $k_{t+1} = (1 - \delta_k)k_t$; this is quantitatively important given the partial irreversibility of capital. The choice grid for labour exploits the features of the financial constraint and thus has a (N_k, N_m, N_ϕ) dimension which depends on the dimension of the state grid for capital, cash and ϕ. This also allows to account for a binding financial constraint exactly. Stochastic shocks are discretized as explained in section 4.1. Having defined the value function, I iterate on the Bellman equation until convergence. At each round of iteration, the value function is interpolated using linear interpolation techniques, to accommodate the discrepancy in the number of grid points between states and choices. Linear interpolation has the advantage of preserving the shape of the policy functions and the kinks arising from the constraints that characterize the model. Some robustness checks have been run, namely testing different solution techniques and grid choices, with the results being always confirmed.

Since the model does not entail endogenous prices, the cross-sectional distribution of firms' states does not constitute an additional state variable.

For the scope of the calibration and the quantitative results, as set out in section 4, the model is simulated over a large number of firms.[40] The transition back from fine choices to coarser states is implemented using a nearest neighbour approach; the simulation keeps

[40] For robustness, I also simulate the model over a continuum of firms, with unaltered results. In that case, I improve Young (2010) algorithm, adapting it to a three-dimensional setting. Firms whose choices are in-between grid points for current states are split up over the two grid points, with weights determined according to the distance of their choice to the nodes.

track of sequential inaction choices and adjusts the policy functions accordingly.

Eight parameters are calibrated simultaneously so that the stochastic steady of the model matches eight empirical moments. I use a multi-start approach with parallel local search in the spirit of Guvenen (2011). First, I generate, 12 vectors of parameter guesses. For each guess vector, I solve the dynamic program allowing for aggregate uncertainty, I then fix the policy functions to the steady state aggregate credit tightness ϕ_H and find the stationary distribution over (m, k, n, z). I compute the model moments and compare to the data moments. I update the guesses, using local search.[41] The algorithm stops when the vector minimizes the sum of squared difference between data and model moments. The 12 local searches are run in parallel on a HPC processor. Then, I compare the 12 local minima and pick the global minimum. This approach is particularly suitable for models with non-linearities and non-convexities, where the global minimum is likely to be surrounded by a large number of local minima.[42]

The simulation results shown in the paper refer to a credit tightening shock lasting 5 periods. I simulate the economy for 200,000 firms over 700 quarters, discarding the first 100 quarters. I then consider a credit shock such that ϕ drops to its low value, stays there for 5 quarters and rebounds back to its steady state. I repeat this simulation for 100 economies, different only with respect to the realisations of the idiosyncratic productivity shock. For the aggregate results, I take the average of aggregate quantities across the economies. An alternative approach lets the aggregate credit shock evolve naturally after a one-period shock. I show in Appendix E the impulse response functions for this case.

The time period in the model is a quarter, and the results shown in the paper follow this frequency. Little information on the frequency of the decision making at firm level is known (Bloom, 2009). I decide to strike a balance between monthly frequency of board meetings in public firms and the annual balance sheet data. When required, model-generated quarterly data is converted into annual figures using standard accounting techniques. Flow figures from the Income Statement are added across the quarters of the year, stock figures from the Balance sheet are taken from the year end values. As reported in FAME company reports, the number of employees is the average over the accounting year.

[41]The results shown in the paper use pattern search. The calibration delivers very similar results using a Nelder-Mead downhill simplex algorithm.

[42]I also implement the simulated annealing algorithm as in Bloom (2009). Although very robust to the existence of multiple local minima, this algorithm is very slow to converge.

C The firm's problem

As mentioned in section 3.5, non-convex labour and capital adjustment costs raise potential concerns with respect to the differentiability of the value function. As shown by Cui (2014), the value function $V(m_t, k_t, n_{t-1}, z_t; \phi_t)$ is differentiable at $k_t > 0$ and satisfies the envelope condition.[43]

The first order conditions that pin down the optimal decisions for dividends, labour, capital and cash respectively, of a firm with idiosyncratic productivity $z_{j,t}$, are shown below. Equation (6) is always binding, which allows to combine it with (7) in (2). With a slight abuse of notation, V' is a compact form for $V\left(m_{t+1}, k_{t+1}, n_t, z_{t+1}; \phi_{t+1}\right)$.

$$\begin{aligned} 1 - 2\kappa\left(d_t - \overline{d}\right) &= \lambda_t && (17) \\ \lambda_t \omega z_{j,t} k_t^{\nu} n_t^{\omega-1} + \beta \frac{\partial EV'}{\partial n_t} - \overline{w}\mu_t &= \lambda_t \left[\overline{w} + AL_{n_t}\left(n_{t-1}, n_t\right)\right] && (18) \\ \beta \frac{\partial EV'}{\partial k_{t+1}} &= \lambda_t AC_{k_{t+1}}\left(k_t, k_{t+1}\right) && (19) \\ \beta \frac{\partial EV'}{\partial m_{t+1}} &= \lambda_t - \psi_t && (20) \end{aligned}$$

And the envelope conditions for labour, capital and cash are:

$$V_{n_{t-1}}\left(m_t, k_t, n_{t-1}, z_{j,t}; \phi_t\right) = -\lambda_t AL_{n_{t-1}}\left(n_{t-1}, n_t\right) \quad (21)$$

$$V_{k_t}\left(m_t, k_t, n_{t-1}, z_{j,t}; \phi_t\right) = \lambda_t \left[\nu z_{j,t} k_t^{\nu-1} n_{t-1}^{\omega} - AC_{k_t}\left(k_t, k_{t+1}\right)\right] + \phi_t(1-\vartheta)(1-\delta_k)\mu_t \quad (22)$$

$$V_{m_t}\left(m_t, k_t, n_{t-1}, z_{j,t}; \phi_t\right) = \lambda_t + \mu_t \quad (23)$$

Combining equations (17-23) gives the first order conditions (10-13) shown in section 3.5. Following Cui (2014), it is possible to further decompose the derivatives with respect to labour and capital adjustment costs. For instance, let $q\left(m_t, k_t, n_{t-1}, z_{j,t}; \phi_t\right)$ be the marginal value of capital that satisfies the envelope condition, which we shall refer to as q_t thereafter. Then, equation (22) can be rewritten as:

$$V_{k_t}\left(m_t, k_t, n_{t-1}, z_{j,t}; \phi_t\right) = \lambda_t \left[\nu z_{j,t} k_t^{\nu-1} n_{t-1}^{\omega} + q_t(1-\delta_k)\right] + \phi_t(1-\vartheta)(1-\delta_k)\mu_t \quad (24)$$

Intuitively, q_t is the marginal reward of adjusting capital. When it reaches 1, a firm

[43] The differentiability of $V(m_t, k_t, n_{t-1}, z_t; \phi_t)$ when $k_{t+1} \neq (1-\delta_k)k_t$ and $n_t \neq (1-\delta_n)n_{t-1}$ is standard, as proved by Benveniste and Scheinkman (1979). The differentiability at $k_{t+1} = (1-\delta_k)k_t$ and $n_t = (1-\delta_n)n_{t-1}$ can be proved using methods from Clausen and Strub (2012), as shown by Cui (2014). The intuition is that the value function is super-differentiable, but also sub-differentiable, given the potential downward kink stemming from the adjustment costs. Being both super-differentiable and sub-differentiable implies the differentiability of the value function.

buys capital. The lower bound of q_t is instead $1 - \vartheta$; selling capital is associated to this marginal reward to decrease capital. When the firm is inactive in its capital investment decision, q_t is less than 1 and greater than $1 - \vartheta$. Inside the inaction region, q_t is the option value of remaining inactive.

D Alternative parameterisations of the adjustment costs

In the calibration exercise outlined in section 4.1, labour and capital adjustment costs are not calibrated to the UK economy, but rather inherited from Bloom (2009). This may pose some concerns about the quantitative performance of the model, since partial irreversibility of capital and hiring/firing costs are important drivers of the model mechanisms. In this section, I evaluate the model performance using different estimates available in the literature. Table D1 shows how models with different labour and capital adjustment costs would imply a worse fit of the over-identifying moment restrictions.

Table D1: Model performance with alternative adjustment costs

		Correlation between cash ratio and capital-labour ratio	Autocorrelation of cash ratios
χ	ϑ		
0.072	34%	-0.15	0.80
0.072	50%	-0.34	0.75
0.32	50%	-0.33	0.70
0.072	2.5%	-0.20	0.39
0.32	2.5%	-0.11	0.68
0.32	34%	-0.3	0.73
Data		-0.15	0.87

Note: Labour (χ) and Capital (ϑ) adjustment costs taken from other estimates often used in the literature. The first line refers to the benchmark calibration used in this paper. Nickell (1986) estimates partial irreversibility of labour at 8% of annual wage (second, fourth and fifth row). Ramey and Shapiro (2001) estimate partial irreversibility of capital between 40 and 80%, whereas Cooper and Haltiwanger (2006) at 2.5%. Moments reported in the table are calculated from the stationary distribution of the stochastic steady state of the model, when credit conditions are at ϕ_H.

E Additional results and pseudo-GE

Table E1 shows how the shape of the firm size distribution in the model resembles the one for the UK economy, although it overestimates the medium-sized firms at the expenses of the right tail of the distribution.

Table E1: Firm size distribution

	Firm share		Employment share	
Firm size	Data	Model	Data	Model
1-9	0.81	0.53	0.16	0.07
10-49	0.16	0.39	0.19	0.57
50+	0.03	0.08	0.65	0.36

Note: Data refer to the number of businesses in the private sector and their associated employment, UK, 2013, Business Population estimates, collected by the Department for Business, Innovation Skills.

The simulation results shown in this paper all refer to a credit supply shock that lasts 5 quarters. Figure E1 shows an alternative approach. 100 Economies are simulated over 200,000 firms for 200 quarters, discarding the first 100 quarters. I impose a credit shock in the same quarter for all economies. ϕ is allowed to evolve naturally in each economy from this quarter onwards. Then I compute the average and median aggregate levels across economies and plot the percent deviation from the pre-shock quarter.

The model outlined in this paper is in partial equilibrium, which implies that prices do not change over time. The interest rate on cash has been set to 0, to help the interpretation of cash holdings in the model.[44] The other price in the economy is the wage. This is assumed to be fully rigid, in the spirit of Shimer (2012) downward nominal wage inflexibility. This may be more controversial in the UK, where real wages have been falling since 2010. On the other hand, the literature still debates over the link between credit shocks and wages.[45] In Figure E2, I redo the simulation exercise outlined in section 4, assuming that wages fall by 1% during credit tightening periods. Compared to the benchmark model, the response in cash is only mildly affected, whereas aggregate employment and capital fall to a smaller extent and recover more quickly. Intuitively, the fall in wages relaxes the financial constraint and counterweights the effects of a credit tightening. Since

[44] Corporate finance models often assume that corporate cash holdings earn a risk-free interest rate (Riddick and Whited, 2009). This is generally net of cash-carry cost, or subject to a tax disadvantage. In such cases, cash holdings comprise a general class of liquid assets, as treasury securities or savings deposits. In my model, I abstract from this possibility. Robustness checks show that this has no impact on the quantitative results.

[45] "An examination of falling real wages, 2010 to 2013" (Office for National statistics, 2014) explores possible reasons behind the fall in UK real wages. Among others, Michelacci and Quadrini (2005) theoretically explore the interactions between real wages and financial markets.

Figure E1: Aggregate responses to a credit supply shock

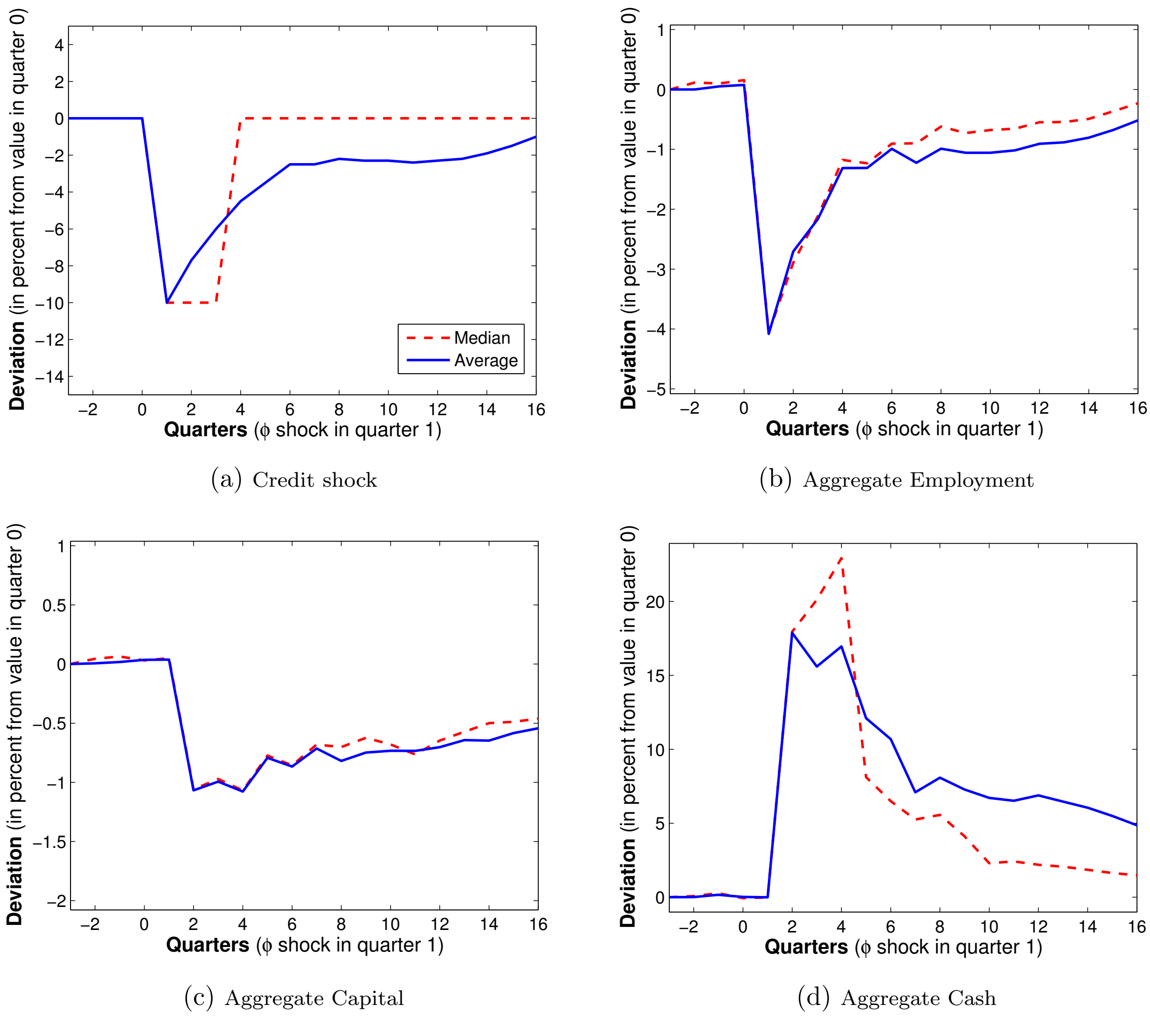

(a) Credit shock

(b) Aggregate Employment

(c) Aggregate Capital

(d) Aggregate Cash

Note: Simulations over 200,000 firms for 200 quarters, discarding the first 100 quarters. Simulations have been repeated for 100 economies. I impose a credit shock in quarter 1, allowing ϕ to evolve naturally thereafter. Then I compute the average and median aggregate levels across economies and plot the percent deviation from quarter 0.

most of the dynamics are driven by unconstrained firms behaving pre-emptively, the anticipation of wage drops during tight credit periods softens the precautionary incentives.

Figure E2: The effects of a credit shock when prices adjust

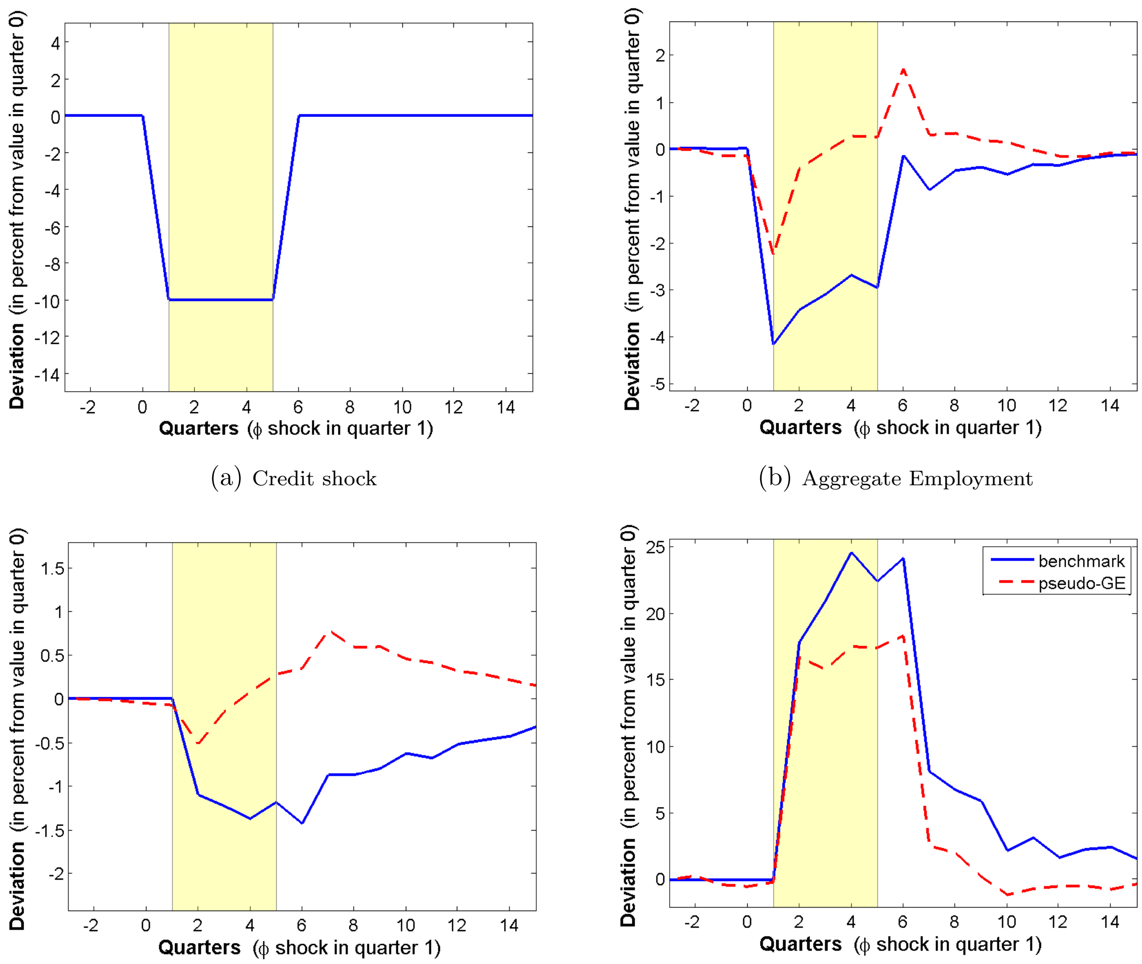

(a) Credit shock

(b) Aggregate Employment

(c) Aggregate Capital

(d) Aggregate Cash

Note: Simulations over 200,000 firms for 700 quarters, discarding the first 100 quarters. Simulations have been repeated 100 times. The IRFs show the average response (across simulations) of a credit shock lasting 5 quarters, generated by a Monte Carlo simulation over the combined transition matrix for aggregate and idiosyncratic shocks. Yellow areas indicate credit tightening periods. In the pseudo-GE economies, the wage falls by 1% together with ϕ. This is known by the firms.

Finally, there has been a recent debate in the literature on the relative response of small and large firms to cyclical shocks. As shown in section 4.5, credit-constrained firms in the model are generally large. However, the presence of precautionary mechanism implies that also unconstrained firms will react to changes to credit conditions. Figure E3 shows that employment at large firms is marginally more responsive to credit shocks than at small firms, although also the latter group is affected. This finding is in line with empirical results by Moscarini and Postel-Vinay (2012). Following their strategy, I classify firms at every quarter, according to their beginning-of-period employment size. I then compute the average employment growth for each group. On average, large firms reduce their employment growth 0.5 percentage points more than small firms. The finding is robust to alternative classification cutoffs and growth measures.[46]

Figure E3: The effects of a credit shock on small and large firms

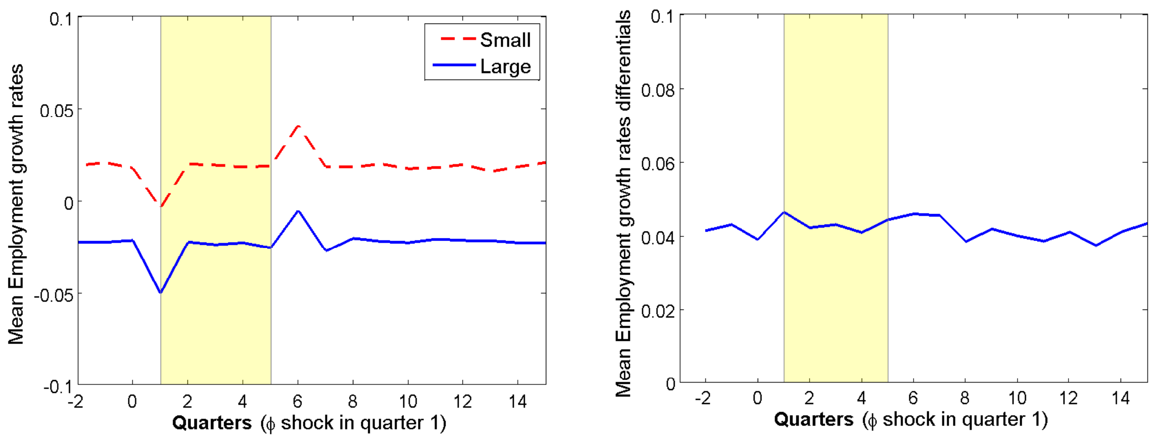

(a) Average employment growth rates

(b) Average employment growth rates differential

Note: Simulations over 200,000 firms for 700 quarters, discarding the first 100 quarters. Simulations have been repeated 100 times. Yellow areas indicate credit tightening periods. Firms are classified each period according to their beginning-of-period size. Large firms have employment above median. Employment growth for a firm j at quarter t is calculated as $\Delta n_{i,t} = \frac{n_{j,t}-n_{j,t-1}}{\alpha n_{j,t}+(1-\alpha)n_{j,t-1}}$, with $\alpha = 0.5$. At each period, the group-specific average employment growth is computed and shown on the left panel. The right panel shows the differential between small and large average employment growth rates.

[46]For instance, it is possible to classify firms according to quantiles of the employment size distribution. Instead of average employment growth rate, Fort et al. (2013) use net employment growth rate, defined as the total stock of net employment changes for a given group over the average stock of employees for that same group.