The London School of Economics and Political Science

Essays in Macroeconomics and Finance

Alex Clymo

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, July 2015

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 54331 words.

Essays in Macroeconomics and Finance

Alex Clymo

July 2015

Abstract

I present a thesis in three chapters on the topics of Macroeconomics and Finance. In the first chapter, I study the ex ante effects of the fear of future financial crises. Crises are modelled through multiple equilibria driven by a self-fulfilling fall in asset prices. I study the effects of allowing agents to anticipate such an event. In a financial crisis, capital is pushed away from experts and towards less productive households, worsening the allocation of capital. Anticipation of this lowers asset prices, investment, and growth today, even if experts are currently well enough capitalised to survive a crisis. The possibility of future crises also creates a state-dependent "financial crisis accelerator" which can amplify business-cycle shocks. In the model, prudential policy can simultaneously increase growth and stabilise the economy, in contrast with common arguments that prudential policy should decrease growth.

In the second chapter, I present evidence that countries which experienced greater declines in total factor productivity (TFP) during the Great Recession experienced milder contractions in hours worked. Thus I show that there is a tension between the crisis manifesting itself either as a problem with productivity or with labour markets. Additionally, countries with larger falls in real wages tend to be those with TFP, and not labour market, problems. Inspired by these facts, I build a model of sticky wages, and prove that wage adjustment determines the extent to which a financial crisis leads to declines in TFP or hours worked. Larger falls in real wages protect labour markets from reductions in hours. However, lower real wages reduce the incentive to reallocate resources across firms during the crisis, leading to larger declines in productivity.

In the final chapter, I introduce financial frictions into the labour market matching model, and study interactions between the two frictions. I demonstrate a feedback between asset and labour markets which amplifies the model's response to exogenous shocks. Shocks which increase equity holders' net worth allow them to fund more

vacancies, raising market tightness and lowering the ease with which firms can hire workers. This increases the value of being an existing firm, causing stock prices to appreciate. This increases experts' net worth further, amplifying the initial shock in a mechanism akin to the traditional financial accelerator. I derive an arbitrage equation in my model similar to the standard free entry condition. I show that any matching model which possesses this arbitrage equation, including the standard matching model, is able to match 82% of the volatility in US market tightness if calibrated to match the volatility in asset prices.

Acknowledgements

I am well aware by now that I owe a lot of things to a lot of people. It seems like I am always getting lucky, but really I am just surrounded by a better bunch of people than I could possibly deserve. I will try to thank some of them here.

I wish to thank my supervisor, Wouter den Haan, for providing such amazing support during the whole of the PhD. Not only did he guide my work and provide moral support, he also taught me an awful lot of useful things. If I ever have the pleasure of being a supervisor I hope I can do half as good a job.

More than anything, this PhD would not have been the same without my friends at LSE, who have provided endless hours of fun. Esther Ann Bøler, Thomas Carr, Sergio de Ferra, Réka Juhász, Andrea Lanteri, Luis Martinez, Jonathan Pinder, and Munir Squires, among many others, have all made the whole thing worth it. Special mention goes to Réka, who I met on the way in to my first lecture at LSE all the way back in 2009. It's been quite the journey.

Outside of LSE are many people who have provided worthwhile distractions from studying. Chief amongst these is Jim Barne, who is still trying to get me to quit academia and write music with him. I'll let you know if it doesn't work out. It is hard to disentangle the PhD and London in my head. I have been an LSE student for the entirety of my adult life in London, and so I would also like to thank everyone who has made living in London such a joy. I am going to miss you all terribly.

I am more indebted than I know how to express to my parents, Jo-Anne and David Clymo. The older I get, the more I realise how lucky I am to have such loving and supportive parents. If I have achieved anything, they deserve most of the credit. Also, my siblings, Anna, Ashley, and Adam. I am lucky to have such a large and happy family.

Finally, my wonderful girlfriend, Katy Lee. Thank you for putting up with me refusing to tell you what my PhD was about for approximately five years. You are an endless source of joy and comfort, and I don't know what I'd do without you.

This thesis is jointly dedicated to Jo-Anne Clymo, David Clymo, and Katy Lee. The first two for raising me, and the third for growing up with me.

Contents

List of Figures

List of Tables

Chapter 1

Growth, Business Cycles, and the Fear of Financial Crises

1.1 Introduction

Recent years have served as a painful reminder that modern economies are not safe from financial crises. While the eventual source of financial crises is often overlooked, looking forward there is a widespread perception that future crises are possible. A casual search for “next crisis” on Google News yields a long list of recent articles on the topic. Whether because changes in regulation in response to the last crisis were inadequate, or even laid the foundations for the next crisis,[1] or just because crises seem to happen every seven years,[2] there is no shortage of potential future crises. In this paper, I ask what the ex ante effects are of such “crisis fear”.

I build a model featuring endogenous financial crises due to multiple equilibria, and study the general equilibrium effects of expectations of the possibility of future

[1] “In the U.S. and Europe, the private sector’s dependence on government support is fostering behaviors – excessive risk-taking, distortions in capital markets and maybe even inflationary pressures – that could lay the foundations for the next crisis.” – WSJ (2012)

[2] “Financial crises come round every seven years on average. There was the stock market crash of 1987, the emerging market meltdown in the mid-1990s, the popping of the dotcom bubble in 2001 and the collapse of Lehman Brothers in 2008. If history is any guide, the next crisis should be coming along some time soon.” – Guardian (2014)

crises. The model is an extension of Brunnermeier & Sannikov's (2014a, henceforth BrS) continuous time model of an economy with a financial sector. Their model features multiplicity of equilibria in some regions of the state space: if leverage is high enough, a self-fulfilling fall in asset prices which bankrupts the financial sector is possible. Brunnermeier & Sannikov (2014b) point out this multiplicity, and the conditions under which it can arise. However, they do not model it formally in agents' expectations: ex ante, agents place zero probability on a crisis happening. My contribution is to treat crises as a sunspot event, allowing agents to understand the probability distribution over future crises. I then show that the fear of future crises has negative effects on both growth and business cycles, and implications for the tradeoffs of prudential policies.

The model features two classes of agents, productive "experts" and less productive "households". Experts borrow from households in order to buy capital to take advantage of their superior production technology. However, in some states of the world a self-fulfilling fall in asset prices is possible, since a fall in asset prices may bankrupt the expert sector, forcing them to liquidate their capital to unproductive households at fire-sale prices.

The first result is the effect of crisis fear on growth. During a crisis, the liquidation of capital to the less productive household sector reduces the average productivity of capital, and hence the possibility of future crises lowers the expected return on capital. This lowers asset prices and hence investment and growth today, even if experts are currently well-enough capitalised to survive a crisis. Thus I demonstrate that the potential future misallocation of capital across agents has implications for growth. The model features endogenous growth via an "AK" structure based on Romer (1986), so changes in the investment rate have permanent growth effects based on a capital externality in production.[3]

[3] If we considered a standard neoclassical growth model, the reduction in investment would instead slow down the transition to the balanced growth path and change the steady state level of capital. Modelling endogenous growth allows me to study the interaction between financial crises and long-run growth, and is computationally convenient because the equilibrium is linear in capital.

The second result is the existence of a state-dependent "financial crisis accelerator": negative shocks which push the economy closer to the region where crises are possible reduce asset prices, which reduces expert net worth, bringing us even closer to the crisis region, reducing asset prices further in a vicious cycle. This operates on top of the traditional financial accelerator, making the overall size of the accelerator both state-dependent and asymmetric: positive shocks near the steady state have smaller effects than negative shocks. In terms of policy, these results suggest that there are both growth *and* stability benefits to dealing with financial crises, which are felt even in times when banks are well capitalised.

Thirdly, I show that while experts may deleverage in response to crisis fear in partial equilibrium, they may also actually take on *more* leverage in general equilibrium. This is because of general equilibrium price effects and limited liability. The fear of a crisis pushes down asset prices, which actually increases expected returns *conditional* on there not being a crisis in the near future. Since experts don't personally suffer all of the losses incurred during a crisis, the increase in this conditional return encourages experts to leverage up. On the other hand, experts are encouraged to deleverage because this reduces their borrowing costs by reducing the incidence of exogenous default costs. If the model features no exogenous default costs, then leverage is everywhere higher when agents fear crises. This adds an interesting interpretation to the link between leverage and crises: we typically think of high leverage causing crisis risk, but I show that the effect also works in the opposite direction, with crisis risk causing high leverage.

Finally, I show that several of BrS' results that are driven by exogenous shocks can be replicated using sunspots. In particular, their model features a bimodal distribution with serious but rare financial disasters driven by bad sequences of exogenous shocks. I show that several of their results hold if we instead consider rare crises driven by multiple equilibria. For example, they show that decreasing the volatility of exogenous shocks in their model increases the amount of endogenous volatility created in general equilibrium. I show that decreasing the volatility of exogenous shocks can increase the probability of experiencing a financial crisis, since lower volatility encourages experts to increase leverage, exposing themselves to crisis risk.

The model also highlights a crucial question: why would the financial sector expose itself to such costly crisis risks? I show that in my model, in the absence of the endogenous-growth externality, an expert who is prudent and takes on low enough leverage to allow herself to survive a crisis would earn infinite value when the crisis hits. The reason for this is intuitive: it is wonderful to be the only expert in town. Prices are low and you operate a better technology than everyone else. In this case, crises cannot exist in equilibrium, since any expert would deviate and take on low leverage if others were taking on high enough leverage to put the economy at risk.

I show that, in the presence of my endogenous-growth externality, crises become possible in equilibrium. The structure I adopt has the productivity of an individual expert being dependent on the total capital managed by the aggregate expert sector. Thus, during a crisis, when other experts are liquidating their capital, your productivity is lowered. This removes the benefit of being the only surviving expert, and allows experts to coordinate on a high leverage, crisis-inducing equilibrium.[4]

My model has implications for the effects of prudential policy, which is often cast as a tradeoff between stability and growth. It is often argued that leverage constraints reduce volatility, but reduce growth by reducing the ability of banks to intermediate capital on average. British Chancellor of the Exchequer George Osborne is quoted as conceding that regulators have to strike a trade-off between risk and economic growth in the aftermath of the recent crisis: "We dont want the financial stability of the graveyard" (FT, 2011). This is a view widely held in the popular press and within businesses. KPMG argue that regulation "... reduces the returns to investors

[4]As I discuss, there are empirical reasons to think the interpretation of it being wonderful to be the only expert in town is probably flawed. My assumption is meant to capture the idea that disruptions to financial markets during a crisis would make it hard to profit from low asset prices. An alternative explanation for experts coordinating on a high leverage equilibrium could be the expectation of receiving bailouts.

in financial institutions. And it reduces economic growth. This has been seen most powerfully and immediately in the downward spiral of bank deleveraging and weak or negative economic growth in Europe" (p. 2, KPMG, 2013). Think tanks have even estimated the negative effect of the recent Dodd-Frank act on growth (AAF, 2015). While the forces which lead regulation to cause low growth are still operational in my model, the fear of crises introduces a counteracting effect, and well-designed prudential policies can simultaneously reduce volatility *and* increase growth.

The intuition for this result builds on the previously discussed growth and volatility results. The fear of future crises reduces growth, and thus prudential policies which reduce the probability (and hence fear) of a future crisis will tend to increase growth. In other words, leverage constraints could promote growth by making the system safer. This positive effect of policy on growth competes with the usual negative effects from reducing the ability of banks to raise funds. Thus whether prudential policy increases or decreases growth is ultimately a quantitative question of which effect dominates.

I consider a policy which forces experts to reduce leverage just enough to rule out crises at all times, which I call the "minimally active" leverage constraint. This policy is countercyclical in asset prices, requiring experts to hold lower leverage when asset prices are high. In a calibrated model, I show that this policy increases growth in equilibrium for a wide range of calibrations for the frequency of financial crises. Since experts' debt is reduced by the leverage constraint, the increase in growth must be financed by an increase in equity. In the model, experts are unable to issue equity, but equity increases because experts retain more earnings by consuming less, which can be interpreted as paying lower dividends. They retain more earnings because eliminating crises increases the value of internal net worth, since it is less likely to be lost during a crisis. If the frequency of crises is high enough, implementing this policy is welfare-improving regardless of the current state of the economy, while it is not welfare-improving in the same model if we ignore crisis risk.

Of course, whether such a policy is implementable in practice is an important question. The optimal degree of countercyclicality, or average level of the minimally active leverage constraint, requires knowledge of the structure of the economy that a policymaker might not possess. I show that the minimally active leverage constraint increases welfare, and investigate the effects of policies which are too tight or loose relative to this benchmark. Policy which is too tight can ultimately lead to the misallocation costs dominating, and hence lead to lower welfare. This suggests the existence of an "*inverse-U*" relationship between leverage policy and welfare: excessively tight leverage constraints will reduce welfare, excessively loose constraints will do nothing, and only intermediate policies can increase welfare. The degree of flexibility a policymaker has in trading these two effects off is a quantitative question, but the range of welfare-improving policies is large in the baseline calibration.

Overall, the aim of the above policy discussion is to highlight that the common stability-vs-growth narrative may be overly simplistic. Even if hoping for an increase in growth from prudential policies might be too much to ask in a practical sense, the positive effects highlighted in this paper could mitigate the negative effects and ultimately reduce the costliness of such policies.

I also discuss bailout policies and market-based solutions. As in Diamond & Dybvig (1983), bailouts can completely rule out the bad equilibrium and hence incur no distortions, since they are never used in equilibrium. However, this is due to the simple structure of the model, and bailouts which are used in equilibrium may lead to distortions. Finally, I show how a simple market-based solution, offering experts insurance which pays off during a crisis, does not rule out crises. Indeed, the same forces which lead experts to be willing to take on high enough leverage to let themselves go bankrupt during a crisis are precisely those which lead them to adopt no insurance against the event.

The rest of the paper is organised as follows. In the next section, I review related literature. In section 1.3 I set up the model and describe some key features. In section

1.4, I solve a version of the model where agents place no weight in their expectations on crises happening, and present preliminary results. In section 1.5, I solve the full model where agents anticipate the possibility of future crises, and present the paper's main results. In section 1.6, I present policy results, and in section 1.7 I conclude.

1.2 Related literature

My paper builds most on the ideas and framework of Brunnermeier & Sannikov (2014a). While they don't explicitly mention the possibility of crises in this paper, they discuss it in an international economics framework in Brunnermeier & Sannikov (2014b). In some states of the world, a second "bad equilibrium" exists in their model whereby experts can go bankrupt in a self-fullfilling crisis. While they discuss the existence of such an equilibrium, they do not allow agents to anticipate that a crisis might happen, and hence cannot discuss how the fear of a potential switch to the bad equilibrium affects behaviour ex ante.

Another recent paper which discusses financial crises in a general equilibrium framework is Gertler & Kiyotaki (2013). Their model features crises via the same mechanism (a fall in asset prices bankrupting banks), and additionally they discuss the effects of anticipated crises. My contribution relative to their paper is that I solve my model globally and nonlinearly, allowing me to study the state dependence of the effects of crisis fear, whereas they linearise around the non-stochastic steady state. Additionally I am able to study growth effects, whereas they assume a fixed stock of capital. Ennis & Keister (2003) also model the effect that the expectation of financial crises has on growth. They use an overlapping generations framework, and model financial crises in a way closer to to the original Diamond & Dybvig (1983) model. In their model, a crisis destroys a fraction of the capital stock, whereas in my model it worsens the allocation of capital across agents.

In an international context, Perri & Quadrini (2014) solve a two-country model with financial crises due to multiple equilibria. Their focus is on explaining how the multiplicity implies that financial crises should be correlated across countries. Crises in their model are also anticipated events, with a known sunspot probability attached to them. Other papers also examine the link between financial frictions and multiple equilibria, and it is well known that financial frictions can lead to multiplicity. For example, Martin & Ventura (2014) and Kocherlakota (2009) present models with collateral constraints and multiple equilibria. The key idea is that higher asset prices relax collateral constraints, increasing the demand for assets and thus justifying the higher asset prices. Martin & Ventura (2014) also point out that the expectation that you might change equilibrium in the future affects today's equilibrium, which is a key theme of my paper.

The literature above and my paper could be viewed as a way to endogenise exogenous financial shocks, which have been shown to generate reasonable macroeconomic features in a recently emerging literature. Eggertson & Krugman (2012) and Jermann & Quadrini (2012) are two notable examples with models close to the representative agent framework, while Khan & Thomas (2013) and Guerrieri & Lorenzoni (2011) study heterogenous firm and consumer models respectively. Theoretically, one issue I am abstracting from is the ability of banks to issue equity. Admati & Hellwig (2013) emphasise that fears of the costs of leverage requirements could be overstated because banks can substitute equity for debt. In a Modigliani-Miller world this is exactly true, and in a world with frictions the costs of reducing debt depend on the relative size of the frictions in equity and debt issuance. My paper complements this idea by showing that even in the extreme case where banks cannot raise any equity the costs of leverage constraints might be overstated, since improvements to the stability of the system can encourage banks to retain earnings by paying less dividends.

Other papers provide evidence supporting or related to my model. Reinhart & Rogoff (2009) document that the historical average duration of recessions surrounding banking crises to be 1.9 years, which they describe as being unusually long compared to normal recessions, which last on average less than a year. They find that the recoveries in unemployment tend to be even more protracted. Claessens & Kose (2013) find similar results, and additionally find that severe financial disruptions have slower recoveries than less severe ones.[5] My model can be interpreted as one rationalisation of this fact: if the fear of future crises rises (rationally or irrationally) following a financial crisis, then investment and growth will be unusually slow in the aftermath.

My model also revolves around the idea that financial crises are anticipated, in the sense that agents worry about future crises, understanding the states of the world in which they occur, and with what probabilities. This is consistent with evidence from financial markets that agents price in future "run risk" for individual financial institutions. Schroth, Suarez & Taylor (2014) present evidence from the asset-backed commercial paper (ABCP) crisis which started in July 2007, which has been widely interpreted as a run due to the short-term nature and widespread withdrawal of financing. They show empirically that spreads on ABCP forecast future runs, consistently with run risk being priced in to lending decisions. Additionally, the economic mechanism behind crises stressed in my paper is that a large fall in asset prices severely reduces the net worth of the financial sector. In historical data, Reinhart & Rogoff (2009) show that equity prices fall on average by 55.9% during banking crises, and that housing prices fall by an average of 35.5%. Additionally, Claessens & Kole (2013) show that asset prices also fall faster during financial crises than they do in normal recessions, which is supportive of their central role.[6]

Finally, my result that policy can simultaneously improve growth and reduce

[5] Some authors have argued that financial crises tend to have faster recoveries than normal recessions. Reinhart & Rogoff (2012) provide a summary of this argument, and evidence supporting their original work.

[6] The large fall in asset prices in the US during the recent crisis is well documented and broad-based, whether we look at equity or land aggregates, or more exotic financial securities. For example, the S&P500 index lost 45% of its value between September 2008 and March 2009, and over 50% from its peak in October 2007. The S&P/Case-Shiller 20-City Composite Home Price Index started falling much earlier, and dropped over 40% of its value between April 2006 and May 2009.

volatility relates to empirical work on the relationship between economic growth and volatility. Ramey & Ramey (1995) and Imbs (2007) present cross-country evidence that economies with lower volatility tend to have higher growth rates on average, which could be interpreted as broadly supportive of my proposed mechanism.

1.3 Model

The model is an extension of BrS' model to allow agents to expect crises, and the underlying framework is very similar. The derivations are thus very similar to the derivations in the original model, with the exception that I need to introduce a sunspot jump variable to allow agents to rationally take into account the possibility that the economy can experience a crisis. Additionally, to formalise the idea of endogenous growth I derive their linear production functions as a special case of Romer's (1986) growth model.

1.3.1 Technology

There are two types of agent, each with unit mass. *Households* are relatively inefficient at production compared to *experts*, but experts will be financially constrained and hence limited in their ability to accumulate capital. Production is carried out using capital using constant returns to scale production functions[7]: Each household has production function $\underline{y}_t = \underline{a}\underline{k}_t$, where under-bars are used throughout to denote household variables as opposed to expert variables. Each expert has production function $y_t = ak_t$, where $a > \underline{a}$. Productivity is constant over time for each class of agent, and the exogenous shock to the economy will instead be a capital quality shock. Capital is accumulated by converting the consumption good into capital. There are adjustment costs, so the price of capital, denoted by q_t, is not equal to unity. Household capital

[7]There is no labour supply in the original BrS model. The linear production functions here should be considered linear *after* labour has been optimally chosen.

accumulation is given by:

$$d\underline{k}_t = \left(\Phi(\underline{\iota}_t) - \underline{\delta}\right)\underline{k}_t dt + \sigma \underline{k}_t dZ_t \tag{1.1}$$

Households' capital depreciates at the rate $\underline{\delta}$, and if they invest at rate $\underline{\iota}_t$ (where this is investment per unit of installed capital) they generate new capital at rate $\Phi(\underline{\iota}_t)$. Finally, their capital stock is subject to an aggregate Brownian shock, dZ_t, which has an effect proportional to their installed capital. This is the aggregate capital quality shock, whose variance is controlled by the parameter σ. Expert capital accumulates according to:

$$dk_t = \left(\Phi(\iota_t) - \delta\right) k_t dt + \sigma k_t dZ_t \tag{1.2}$$

Households have the same exposure to the shock as experts, and the same adjustment cost function, but have lower productivity and a potentially higher depreciation rate, $\underline{\delta} \geq \delta$. In the appendix I prove that these linear production functions can be viewed as the reduced form of a modification of Romer's (1986) endogenous growth model.

1.3.2 Sunspot & price process

In solving this kind of model it is typical at this point to conjecture that prices evolve as drift-diffusion process, with unknown drift and loading on the exogenous diffusion, dZ_t. However, given that this economy features a multiplicity of equilibria we can introduce a sunspot variable and conjecture that prices evolve as a function of this variable too. This obviously has to be verified in equilibrium. I guess that the capital price evolves according to:

$$dq_t = \mu_t^q q_t dt + \sigma_t^q q_t dZ_t + (\underline{q} - q_t) df_t \tag{1.3}$$

Where df_t is the change between t and $t + dt$ in a counting variable which increases by one in that interval with probability $\rho_{e,t} dt$, and doesn't increase with probability

$1 - \rho_{e,t}dt$. Thus the price evolves as a combination of time varying drift, diffusion from the exogenous capital quality shock, and a jump component from the sunspot. The final term is the sunspot term which says that the price could jump down from q_t to $\underline{q}$ with some probability. This is what happens in a crisis. The cause of the fall in price will be the bankruptcy of all of the experts in the economy. Thus the jump in prices will also coincide with default and their joint consequences for the other variables of the economy.

The jump intensity, $\rho_{e,t}$ has both endogenous and exogenous components, and hence has to be determined in equilibrium. For example, in regions where banks are well enough capitalised to survive a crisis this will be endogenously zero. However, in regions where a crisis is possible, the modeller has freedom to choose how likely it is that a crisis occurs. As usual with models of multiple equilibria, there is nothing intrinsic in the model which tells us when we should switch equilibria. This means that I must look outside the model for equilibrium selection, which is why sunspots are required.

The crisis price, $\underline{q}$, is also endogenous and at this point undetermined. Intuitively, this is the price at which capital would trade if all experts go bankrupt and cease to intermediate capital. This is going to be lower than the current price, because this means that only inefficient households can purchase capital.

1.3.3 Markets

As well as the markets for consumption and for trading units of capital, there is a restricted set of financial markets. In particular, experts and households can only trade risky debt, and not state contingent claims (such as equity). Banks borrow from households at interest rate, r_t. In the case of default their net worth is reduced to zero, and households seize their capital less a proportional default cost.

1.3.4 Households

Households are risk neutral and allowed negative consumption, so their required expected return on any asset is always simply their subjective discount rate, ρ_h. Since experts may default in a crisis, they borrow at a rate above the risk free rate. In the appendix I show that the interest rate charged to an expert is:

$$r_t = \begin{cases} \rho_h & : \phi_t < \frac{1}{1-\frac{\underline{q}}{q_t}} \\ \rho_h + \rho_{e,t}\left(1-(1-\chi)\frac{\underline{q}}{q_t}\frac{\phi_t}{\phi_t-1}\right) & : \phi_t \geq \frac{1}{1-\frac{\underline{q}}{q_t}} \end{cases} \tag{1.4}$$

Where χ is the fraction of expert assets exogenously destroyed during default. $\phi_t \equiv \frac{q_t k_t^b}{n_t}$ is the leverage of a typical expert. The term in brackets is necessarily positive, leading to an interest rate spread, and it is possible to show that the interest rate is increasing in bank leverage. Note that the interest rate charged to any one expert depends optimally on that expert's own leverage, and not aggregate leverage.

Using the household capital evolution equation and conjectured price process we can calculate the household's return on holding capital. Ito's lemma with jumps gives us:[8]

$$d\underline{r}_t^k = \left(\frac{\underline{a}-\underline{\iota}_t}{q_t} + \Phi(\underline{\iota}_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q\right)dt + (\sigma + \sigma_t^q)\,dZ_t + \frac{\underline{q}-q_t}{q_t}df_t \tag{1.5}$$

Note that the realised return on capital could involve a non infinitesimal loss in the case of a jump. The expected return will remain infinitesimal because this only happens with probability of order dt: $E_t df_t = \rho_{e,t}dt$. The expected return for a household is:

$$\mathrm{E}d\underline{r}_t^k = \left(\frac{\underline{a}-\underline{\iota}_t}{q_t} + \Phi(\underline{\iota}_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q + \frac{\underline{q}-q_t}{q_t}\rho_{e,t}\right)dt \tag{1.6}$$

[8] The return is composed of a dividend $\underline{a}$ plus a capital gains term. Capital gains are computed as $\frac{d(q_t\underline{k}_t)}{q_t\underline{k}_t}$ which includes the capital gain from price changes and from changes in the capital stock itself, due to either the shock or investment and depreciation.

In continuous time the investment decision is a static problem, and we can determine the optimal investment rate as the rate that maximises the above return (or equivalently the expected return). Differentiation with respect to $\underline{\iota}_t$ yields:

$$\Phi'(\underline{\iota}_t) = \frac{1}{q_t} \tag{1.7}$$

Thus the optimal investment rate depends only on the current price of capital, and from now on we can think of $\underline{\iota}_t$ as implicitly being defined by q_t. Households are not allowed to short the capital stock so, given the assumption of risk neutrality, either they hold zero capital, or are indifferent about their capital holdings, or want to hold infinite capital (which is ruled out by market clearing). Thus we can summarise the household's capital optimality conditions as:

$$\frac{\underline{a} - \underline{\iota}_t}{q_t} + \Phi(\underline{\iota}_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q + \frac{\underline{q} - q_t}{q_t}\rho_{e,t} \leq \rho_h \tag{1.8}$$

Which holds with equality if the household holds capital. Define $\psi_t \equiv k_t/K_t$ as the share of the total capital stock owned by experts (where k_t is the integral over the identical holdings of the unit mass of experts). Then the above inequality is binding in equilibrium if and only if $\psi_t < 1$.

Finally, it is convenient at this point to ask what the price of capital would have to be if the household was to hold all of the capital stock forever (i.e. if the household was the only agent in this economy). In this case we can use the household's capital FOC, (1.8), to give the capital price. Guessing that in this case the price of capital is constant ($\mu_t^q = \sigma_t^q = 0$) gives:

$$\frac{\underline{a} - \underline{\iota}_t}{q_h} + \Phi(\iota(q_h)) - \underline{\delta} = \rho_h \Rightarrow q_h = \frac{\underline{a} - \iota(q_h)}{\rho_h + \underline{\delta} + \Phi(\iota(q_h))} \tag{1.9}$$

We see that our guess that $\mu_t^q = \sigma_t^q = 0$ is confirmed since there are no time varying elements in the above equation.

1.3.5 Experts

Experts are also risk neutral, and have subjective discount rate ρ_b. Experts, unlike households, must have non-negative consumption. This means that they can become financially constrained, because if they are lacking in net worth they will only be able to expand capital holdings by issuing risk free debt, which comes at the cost of magnifying risk. Denote an expert's net worth by n_t, where this is the market-to-market book value of her assets minus liabilities. I define the current maximised value of her utility by $\theta_t n_t$. Thus θ_t is an expert's value per unit of net worth, which I call experts' "marginal value", and which will be a function of the aggregate state in equilibrium. Conjecture that marginal value follows:

$$d\theta_t = \mu_t^\theta \theta_t dt + \sigma_t^\theta \theta_t dZ_t + df_t \left(\underline{\theta}_t - \theta_t\right) \tag{1.10}$$

Where $\underline{\theta}_t$ is the value of an expert's marginal value following a crash, which is to be determined. An expert earns the following return on capital:

$$dr_t^k = \left(\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q\right) dt + \left(\sigma + \sigma_t^q\right) dZ_t + \frac{\underline{q} - q_t}{q_t} df_t \tag{1.11}$$

This return is maximised by the same choice for investment as households: $\Phi'(\iota_t) = 1/q_t$. They pay an interest rate which depends on how much leverage they take on, as in (1.4). I derive the solution to the expert's problem in the appendix. A very important issue is what happens to the individual expert during a crisis. Remember that the crisis is an aggregate event, and the expert does not think that she has any power to affect whether it happens or not. This is because a crisis is a fall in asset prices, which individual agents take as given. But an expert does have the power to control whether or not she personally goes bankrupt during a crisis. To see this, note that if the price of capital instantaneously falls from q_t to $\underline{q}$, an expert's net worth falls to:

$$\underline{n}_t = \max\left\{n_t - q_t k_t + \underline{q} k_t, 0\right\} = \max\left\{n_t\left(1 - \phi_t\left(1 - \frac{\underline{q}}{q_t}\right)\right), 0\right\} \tag{1.12}$$

The expert will go bankrupt during a crisis if the term inside the max operator is less than zero, leaving $\underline{n}_t = 0$, and will survive the crisis if $\underline{n}_t \geq 0$. Note that while the occurrence of a crisis is a random event, it is completely deterministic whether an expert survives the crisis. If prices fall in a crisis then $1 - \underline{q}/q_t > 0$, meaning that $\underline{n}_t$ decreases in leverage, and an expert can be certain to survive the crisis by choosing low enough leverage. In the extreme, an expert could choose to buy no capital ($\phi_t = 0$) and be certain to always survive a crisis, since then $\underline{n}_t = n_t > 0$.

This introduces a kink into the expert's value function, because there is a threshold leverage choice $\hat{\phi}_t = q_t/(\underline{q} - q_t)$ above which the expert goes bankrupt in a crisis, and below which she does not. Thus a key element of equilibrium will be which region the expert optimally chooses: a crisis cannot be an equilibrium if the existence of a crisis causes all of the experts to deleverage to avoid it. Understanding the conditions under which experts do and do not expose themselves to this crisis risk is clearly an important question. However, the focus of this paper is on understanding the general equilibrium effects of such crisis risk, and as such I make assumptions to ensure that if all other experts are choosing high enough leverage to make a crisis possible then you are also happy to take on that high level of leverage.

In particular, the appendix details the endogenous growth structure behind the reduced-form linear production technologies used in the main text. These use an aggregate capital externality based on Romer (1986), whereby the productivity of any individual expert depends on the total capital being intermediated by the entire *expert sector*. In this case, if all other experts take on enough leverage to expose themselves to a crisis, you have a strong incentive to as well. This is because during a crisis the rest of the experts will have to shed all of their capital, reducing your own productivity during a crisis via the capital externality. In the appendix I prove that this leads to experts optimally maintain high enough leverage to expose themselves to a crisis, if all other experts are doing so.[9]

I relegate the derivation of the expert's problem to the appendix, and present the key results here. The expert's leverage first order condition gives:

$$\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - \rho_h = -\sigma_t^\theta\left(\sigma + \sigma_t^q\right) + \rho_{e,t}\left(1 - (1-\chi)\frac{\underline{q}}{q_t}\right) \tag{1.13}$$

This is identical to BrS' original FOC, with the addition of the final term. The FOC can be interpreted as a simple marginal benefit / marginal cost comparison. On the left hand side is the marginal benefit of increasing leverage, which is the expected excess return on capital over the risk free rate. On the right hand side is the marginal cost of increasing leverage. The first term is common with the original BrS model, and is the increase in risk the expert faces by increasing leverage. The financial friction makes the expert effectively risk averse, and since she has to finance risky capital using less than fully-contingent debt her risk is magnified as she leverages up. The final term is the marginal cost of increasing leverage related to crises. By leveraging up the expert splits her collateral more thinly across her creditors, leading to a higher required interest rate as there is less security per unit lent. The first order condition for consumption gives:

$$\theta_t \geq 1 \tag{1.14}$$

with equality if $dc_t > 0$. The marginal utility gained from consuming today is always one due to the risk neutrality assumption. As in BrS, the expert thus consumes nothing if the marginal value of retaining earnings exceeds one. If the value of retain-

[9] In the appendix I provide a thorough discussion of the conditions under which an expert would allow herself to go bankrupt during a crisis. My technological assumption is a stand in for the idea that disruption in financial markets during a crisis makes it hard to take advantage of the high expected returns that come with temporarily low asset prices. Alternatively, the expectation of bailouts could explain why institutions expose themselves to such risk. The question then becomes why individuals do not reduce their exposure to the crisis. In the words of Citygroup chief executive Chuck Prince, "When the music stops, in terms of liquidity, things will be complicated. But as long as the music is playing, you've got to get up and dance. We're still dancing." (FT, 2007)

ing earnings ever falls to one the expert is indifferent about consuming and holding capital, placing a lower bound of one on θ_t. At the optimum, and when $dc_t = 0$, evaluating the value function gives:

$$\mu_t^\theta = \rho_b - \rho_h \tag{1.15}$$

This gives us our solution for μ_t^θ, and solving the expert's problem now just requires finding a solution for σ_t^θ at every point in the aggregate state space. One important thing to note is that neither of the optimality conditions pins down an exact value for optimal leverage. The expert is in fact indifferent about her choice of leverage as long as (1.13) holds. This is a product of the risk neutrality assumption, and surprisingly even holds when we add crisis risk.

1.3.6 Market clearing

In equilibrium prices adjust to clear the consumption, capital and bond markets. Consumption market clearing requires that expert and household consumption and investment flows equal the flow of production. The risk neutrality of households ensures that this market clears. Capital market clearing requires that expert and household capital demand sums to the supply of installed capital:

$$k_t + \underline{k}_t = K_t \tag{1.16}$$

Where K_t is the current stock of capital, which is an aggregate state variable. In equilibrium, the price of capital adjusts so that the sum of expert and household capital demand equals the existing total installed capital stock. In practice, given the linearity of both sets of agents' policy functions, this means adjusting the price to ensure that they are either indifferent about their capital holdings, or don't want to hold any at all.

Experts will always hold capital in equilibrium, since they are more productive than households. Households may not find it profitable to do so depending on the current return. There are thus two regions of the state space, corresponding to whether or not the household holds capital. If the household doesn't hold capital then we determine the price of capital using the expert's leverage first order condition, (1.13), and need to check that (1.8) holds with inequality. That is, in this region:

$$\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - \rho_h = -\sigma_t^\theta\left(\sigma+\sigma_t^q\right) + \rho_{e,t}\left(1-(1-\chi)\frac{\underline{q}}{q_t}\right)$$

And:

$$\frac{\underline{a}-\underline{\iota}_t}{q_t} + \Phi(\underline{\iota}_t) - \underline{\delta} + \mu_t^q + \sigma\sigma_t^q + \frac{\underline{q}-q_t}{q_t}\rho_{e,t} < \rho_h$$

Since the household is not holding capital, we only need to adjust the price of capital to make sure the expert is indifferent about her leverage choice. In the region where the household does hold capital, both agents' FOCs hold with equality, so we can set (1.8) to bind and subtract it from (1.13) to give:

$$\frac{a-\underline{a}}{q_t} + \underline{\delta} - \delta - \rho_{e,t}\chi\frac{\underline{q}}{q_t} + \sigma_t^\theta\left(\sigma+\sigma_t^q\right) = 0 \tag{1.17}$$

This is the counterpart to BrS' equation (17) when we allow for anticipated crises. This equation is important because it tells us what has to be true to make experts and households simultaneously happy to hold capital. Intuitively, we have to somehow make both experts and households indifferent about their capital holdings, even though experts are more productive and would hence tend to be happier to intermediate. Notice that in the absence of the risk or jump terms this would be impossible since $a > \underline{a}$ and $\underline{\delta} > \delta$, and no q_t could make the above equation hold.

In the original BrS model without crises the capital market is actually cleared by adjusting the level of endogenous risk. To make both sets of agents happy to hold capital the experts' productivity advantage is offset by the fact that they dislike risk, while the household does not. This is where the actual value of expert leverage is determined. Remember it is not pinned down by individual optimisation, since they are individually indifferent about their leverage levels. The level of leverage is determined in general equilibrium, as the level which creates enough risk to satisfy the above equation. Remember that increasing leverage increases risk, by increasing the experts' exposure to risky assets while funding them with risk free debt.

In the model with anticipated crises, there is also the jump term to consider. If we are in the crisis-prone region and $\chi > 0$ then the jump term tends to reduce leverage, all else equal, because it pushes up expert borrowing rates and increases the marginal cost of leverage. In order to reduce the marginal cost to restore indifference we need to reduce the amount of endogenous risk, which is achieved by lower leverage in equilibrium as discussed above.

The other state variable is total bank net worth, N_t. The capital quality shock is i.i.d and therefore there is no need to include its value as a state. The sunspot also doesn't introduce a new state variable conditional on the current state, since it is simply a flow probability of moving between equilibria. Thus we can completely describe the equilibrium of the economy on the state space (N_t, K_t).

1.3.7 State space representation

I noted that the state variables are (N_t, K_t). BrS show that the equilibrium of the economy scales linearly in K_t if we use bank net worth as a proportion of total net worth as a state instead of N_t:

$$\eta_t \equiv \frac{N_t}{q_t K_t} \tag{1.18}$$

In other words, we use (η_t, K_t) as a state. Most variables, such as q_t and ϕ_t, will only depend on η_t in equilibrium. Others, such as total consumption, output and capital demand, will depend on η_t, but also scale linearly in K_t. Conjecture the following law of motion for η_t:

$$d\eta_t = \mu_t^\eta \eta_t dt + \sigma_t^\eta \eta_t dZ_t - \eta_t df_t \tag{1.19}$$

The drift and volatility terms are to be determined along with the drifts and volatilities for q_t and θ_t. The jump term says that in a crisis the value of η_t jumps down to zero. This is because in a crisis N_t jumps to zero, and hence so does $N_t/(q_t K_t)$. The variables q_t, θ_t and ψ_t are functions only of η_t in equilibrium. Given ψ_t and η_t it is easy to calculate leverage as $\phi_t = \psi_t/\eta_t$. Using Ito's lemma we can solve for all of the unknown drifts and volatilities as functions of the parameters and the unknown functions above. I relegate the derivations to the appendix.

1.3.8 Crisis equilibrium

At this point we can construct the crisis equilibrium. Suppose that at some time t experts have leverage ϕ_t, net worth N_t, and face a price of capital of q_t. The multiplicity I investigate is whether a fall in price from q_t to $\underline{q}$ is enough to bankrupt the experts. The fall in price causes net worth to fall from N_t to:

$$\underline{N}_t = N_t - q_t k_t + \underline{q} k_t = N_t \left(1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right)\right) \tag{1.20}$$

If $\underline{N}_t < 0$ the experts go bankrupt. Given knowledge of $\underline{q}$, a crisis is thus possible at the current state if:

$$1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right) < 0 \tag{1.21}$$

The question is now whether the bankruptcy of experts can *justify* why the price fell to $\underline{q}$. This is true if the equilibrium price post bankruptcy is $\underline{q}$. Given that the price is a function of the state of the economy, $\underline{q}$ is an equilibrium object: it is the price that capital trades at once experts go bankrupt. Thus calculating this price requires specifying what happens in the economy after a crisis.

At the moment a crisis hits expert net worth drops to zero, as they do not have enough to pay off their creditors. With zero net worth experts have no money to purchase capital, and they cannot leverage off zero net worth. Hence households have to hold the whole capital stock in equilibrium ($\psi_t = 0$) immediately following a crisis.

However, this is not enough information to price capital, since even the household's first order conditions are forward looking, and so the crisis price of capital will depend on the expectation of future prices.

However, note that experts' policy functions are linear in net worth. This means that once their net worth drops to zero, it must remain there indefinitely. Intuitively, with no net worth they cannot invest, and hence cannot generate any new net worth. This means that, in the absence of any intervention, a financial crisis would be permanent in this model. This makes it easy to price capital: the price when households hold the whole capital stock forever can be simply solved from (1.9). For simplicity, this is the assumption I maintain in the baseline model.

To facilitate non-permanent crisis, I consider the following extension. In order to spur the recovery of the experts, I can introduce an exogenous equity injection to restore them to positive net worth. This can be thought of as originating from either households or from a government sector. Specifically, I assume that following a crisis, the expert sector will receive an equity injection with probability $\rho_r dt$ in any interval dt. This injection is sufficient to restore η to some $\hat{\eta}$, and once it is given no further equity injections are given until another crisis occurs. Once the equity injection is given, capital is priced via the normal equilibrium and hence the price jumps up to $q(\hat{\eta})$. Until this happens only households are holding capital, and hence we can price capital using the household's capital first order condition. The price will be constant at $\underline{q}$ until the experts are recapitalised. Taking this into account, the return a household earns on capital during a crisis is:

$$d\underline{r}_t^k = \left(\frac{\underline{a} - \underline{\iota}_t}{\underline{q}} + \Phi(\underline{\iota}_t) - \underline{\delta}\right) dt + \sigma dZ_t + \frac{q(\hat{\eta}) - \underline{q}}{\underline{q}} dg_t \tag{1.22}$$

The drift and diffusion terms give the return if we remain in a crisis. Note that during a crisis the price is constant at $\underline{q}$, so there are no price appreciation or volatility terms. The last term gives the capital gain the household makes in the even that experts are recapitalised, which is the jump event dg_t. In equilibrium the expected return equals the household's subjective discount rate:

$$\rho_h = \frac{\underline{a} - \underline{\iota}_t}{\underline{q}} + \Phi(\underline{\iota}_t) - \underline{\delta} + \frac{q(\hat{\eta}) - \underline{q}}{\underline{q}} \rho_r \tag{1.23}$$

This gives one equation to determine $\underline{q}$ given $q(\hat{\eta})$. Notice that as $\rho_r \to 0$ the price of capital approaches the price of capital if households are expected to hold all of the capital stock forever. Once $\hat{\eta}$ and ρ_r are chosen, and knowing the equilibrium price function in normal times, $q(\eta)$, we thus have enough information to calculate the price of capital during a crisis. I solve a version of the model with $\rho_r > 0$ in the appendix, and restrict myself to permanent crises ($\rho_r = 0$) for the baseline calibration.

The fundamental cause of crises in this model is the *misallocation of capital*. A crisis causes experts to go bankrupt, which pushes all of the capital stock into the hands of inefficient households. Since these households are inefficient and produce less from the capital stock the price of capital falls to reflect this. This model is thus fundamentally a model of multiplicity due to endogenous misallocation of capital, and a crisis manifests itself as a drop in measured total factor productivity (TFP). An important question is whether modelling the real effects of crises as a drop in measured TFP is empirically correct, a question I tackle in depth in my second chapter.

1.3.9 Selecting equilibria

In a region of the state space where a crisis is possible, between t and $t + dt$ the economy carries on as usual with probability $1 - \rho_{e,t}dt$, and experiences a crisis with probability $\rho_{e,t}dt$. The intensity $\rho_{e,t}$ is a variable that the modeller gets to choose, and intuitively controls the probability that agents coordinate on the belief that a self fulfilling fall in asset prices will happen. A simple assumption would be to set ρ_e to a constant value, which would mean the probability of experiencing a crisis is constant (whenever a crisis is possible). This leads to the following form for $\rho_{e,t}$:

$$\rho_{e,t} = \begin{cases} \rho_e & : \underline{N}_t < 0 \\ 0 & : \underline{N}_t \geq 0 \end{cases} \tag{1.24}$$

So far I have only discussed switching from the "normal" equilibrium to the crisis equilibrium, but is it possible to switch back? The answer is no, as once we enter a crisis expert capital holdings fall to zero, and it is hence not possible for changes in capital prices to push expert net worth around, which was the mechanism for our jumps. Instead, we return to the normal equilibrium once experts' net worth is restored via an equity injection.

1.3.10 Equilibrium

This section is the equivalent to BrS' Proposition II.4 in my model. We solve for the unknown functions q, θ on state space $\eta \in [0, \eta^*]$. All other variables are implicitly defined by (q, θ, η). At any point in the state space where $\underline{N}(\eta) < 0$ the economy may experience a crisis, and does so at the endogenous rate $\rho_{e,t}$ which is solved for along with the other equilibrium variables. η^* is an upper bound on η because at this point returns are so low that experts consume their net worth. There are five boundary conditions:

$$q(0) = q_h \qquad \theta(\eta^*) = 1 \qquad q'(\eta^*) = 0 \qquad \theta'(\eta^*) = 0 \qquad \lim_{\eta \to 0} \theta(\eta) = \infty \tag{1.25}$$

Given the current state, η, and $(q(\eta), q'(\eta), \theta(\eta), \theta'(\eta))$ we can calculate $(q''(\eta), \theta''(\eta))$ using the following procedure.

1. Find ψ such that

$$\frac{a - \underline{a}}{q_t} + \underline{\delta} - \delta - \rho_{e,t} \chi \frac{\underline{q}}{q_t} + \sigma_t^\theta \left(\sigma + \sigma_t^q\right) = 0 \tag{1.26}$$

where:

$$\sigma_t^\eta \eta = \frac{(\psi - \eta)\sigma}{1 - (\psi - \eta)q'(\eta)/q(\eta)} \qquad \sigma_t^\theta = \frac{\theta'(\eta_t)\sigma_t^\eta \eta_t}{\theta_t} \qquad \sigma_t^q = \frac{q'(\eta_t)\sigma_t^\eta \eta_t}{q_t} \tag{1.27}$$

and

$$\rho_{e,t} = \begin{cases} \rho_e & : 1 - \frac{\psi}{\eta}\left(1 - \frac{\underline{q}}{q_t}\right) < 0 \\ 0 & : 1 - \frac{\psi}{\eta}\left(1 - \frac{\underline{q}}{q_t}\right) \geq 0 \end{cases} \tag{1.28}$$

If $\psi > 1$ set $\psi = 1$ and recalculate (1.27).

2. Compute

$$\mu_t^\eta = -\sigma_t^\eta(\sigma + \sigma_t^q + \sigma_t^\theta) + \rho_{e,t}(1 - \chi)\frac{\underline{q}}{q_t} + \frac{a - \iota_t}{q_t} + (1 - \psi_t)(\underline{\delta} - \delta) \tag{1.29}$$

$$\mu_t^q = \rho_h - \frac{a - \iota_t}{q_t} - \Phi(\iota_t) + \delta - \sigma\sigma_t^q - \sigma_t^\theta\left(\sigma + \sigma_t^q\right) + \rho_{e,t}\left(1 - (1 - \chi)\frac{\underline{q}}{q_t}\right) \tag{1.30}$$

$$\mu_t^\theta = \rho_b - \rho_h \tag{1.31}$$

$$q''(\eta) = \frac{2\left[\mu_t^q q(\eta) - q'(\eta_t)\mu_t^\eta \eta\right]}{\left(\sigma_t^\eta\right)^2 \eta^2} \qquad \theta''(\eta) = \frac{2\left[\mu_t^\theta \theta(\eta) - \theta'(\eta_t)\mu_t^\eta \eta\right]}{\left(\sigma_t^\eta\right)^2 \eta^2}$$

Finally, $\hat{\eta}$ is a given parameter, and the crisis price, $\underline{q}$, solves:

$$\rho_h = \frac{\underline{a} - \underline{\iota}_t}{\underline{q}} + \Phi(\underline{\iota}_t) - \underline{\delta} + \frac{q(\hat{\eta}) - \underline{q}}{\underline{q}}\rho_r \tag{1.32}$$

1.3.11 Looking through the equations in partial equilibrium

To begin understanding the effect that anticipated crises have on equilibrium, it is helpful to look through the equations of the model, as summarised above. This allows us to trace where crises enter the equilibrium and gain some intuition as to the effects. Notice that the probability of a crisis, $\rho_{e,t}$, enters via three equations: (1.26), (1.29) and (1.30).

Equation 1.26: Leverage

Starting with equation (1.26), the following result is useful:

Proposition 1. *Suppose we are at a given state η in the interior of the crisis region, and where both experts and households intermediate capital. Then for given $(q(\eta), q'(\eta), \theta(\eta), \theta'(\eta))$ an increase in the probability ρ_e of selecting the crisis equilibrium reduces expert leverage if and only if $\chi > 0$.*

Note that the reason we condition on $(q(\eta), q'(\eta), \theta(\eta), \theta'(\eta))$ in the proof is that this parallels how we actually solve for leverage in the numerical solution. The proof is simple and relegated to the appendix, but the intuition is more subtle and I discuss it here. This is not a true partial equilibrium experiment, in the sense that we study a single agent's actions when we take prices as exogenous. To see this note that I have taken expert value, θ, as given too.

Remember that we form equation (1.26) by subtracting the household's capital first-order condition (if it is binding) from the expert's leverage first order condition. Thus (1.26) holds if both agents are holding capital, and is necessary for them both to be indifferent about their capital holdings. $\rho_{e,t}$ only enters equation (1.26) if $\chi > 0$, and does not if $\chi = 0$. This means that in the absence of exogenous default costs ($\chi = 0$) crisis risk has no immediate impact on who prefers to hold capital.

To see this, note that increasing the probability of selecting a crisis affects the optimality conditions for *both* experts and households. For households, it increases the probability that they suffer a large capital loss on the capital in a crisis, reducing their expected return on capital, as in (1.8). For experts the story is slightly different. While an increase in the probability of a crisis does reduce their expected return, they do not care directly about this. This is because they go bankrupt during a crisis, and hence place no weight on this loss in their optimisation. Instead, experts only care about the probability of a crisis because it affects the spread they pay on their debt. When experts consider taking on an extra unit of leverage they understand that they will have to repay one extra unit at the current interest rate $r(\phi)$, and will also face a marginal increase in their interest payments $r'(\phi)$ because their collateral becomes more thinly spread across their debtors. An increase in the probability of a crisis increases both the current interest rate and the marginal increase.

Interestingly, when $\chi = 0$ these effects are of exactly the same magnitude, and crisis risk does not hurt one group more than the other. Thus when we subtract the two first order conditions from each other, no terms involving $\rho_{e,t}$ remain. However, when $\chi > 0$ this is not true, and a term involving $\rho_{e,t}$ does remain. This term causes us to reduce leverage in response to crisis risk. Intuitively, this is because when we add exogenous default costs the interest rate the expert pays increases faster as she increases leverage, encouraging her to deleverage.

In sum, if exogenous default costs are positive, increasing the probability of a crisis harms both the experts and the household, but harms the expert more on the margin. Given this, to restore both agents to indifference we will need to help the expert, which is achieved by reducing equilibrium leverage. This helps the expert by reducing the amount of endogenous risk, as discussed in Section 1.3.6, which helps experts since they are effectively risk averse.

Thus I have established that in partial equilibrium anticipation of crises reduces expert leverage if and only if $\chi > 0$, which pushes capital into the hands of inefficient households. However, we will see in the numerical section that this partial equilibrium result is overturned by general equilibrium forces in some areas of the state space, meaning that leverage could even increase in response to crisis risk.

Equation 1.30: Capital gains

Equation (1.30) allows us to compute the drift in the capital price, μ_t^q once we have computed the volatilities. Since $1 - (1 - \chi)\underline{q}/q_t > 0$, we see that this equation gives us a higher μ_t^q whenever $\rho_{e,t}$ is higher, all else constant.

The intuition for why this is the case is relatively simple. Equation (1.30) is the expert's leverage first order condition, rearranged to solve for μ_t^q. Increasing $\rho_{e,t}$ increases the expert's marginal cost of leverage by increasing the interest rate the expert pays on debt. In order to return the expert to indifference about her leverage choice we must increase the marginal benefit of leverage, which is done by increasing the expected return to capital. Since we are (for the purposes of this discussion) taking all else as given, this is achieved by increasing μ_t^q which increases capital gains.

This reasoning leads us to expect that we should expect that expected returns *conditional on there not being a crisis* should be higher when agents attach a high probability to there being a crisis in the near future. We can see this by rearranging (1.30) and using the definition of dr_t^k:

$$\mathrm{E}_t\left[dr_t^k | df_t = 0\right] = \rho_h - \sigma_t^\theta\left(\sigma + \sigma_t^q\right) + \rho_{e,t}\left(1 - (1-\chi)\frac{\underline{q}}{q_t}\right) \tag{1.33}$$

While this is true for the expected return conditional on there being no crisis, it is less so for the overall expected return. Nonetheless, the conditional expected return is interesting because it tells us how the economy behaves in the scenario that we don't experience a crisis.

Equation 1.29: Expert net worth

The final equation $\rho_{e,t}$ appears in is the drift of the aggregate state, μ_t^η. The aggregate state is expert net worth as a fraction of the value of the capital stock, and tends to grow when expert net worth grows. Like μ_t^q, it is also increasing in $\rho_{e,t}$ all else equal. This is a reflection of the higher expected returns (conditional on no crisis) that the experts earn when crisis risk increases.

This suggests a potentially interesting trade off: higher crisis risk means that the economy is more likely to suffer a large fall in the net worth of the financial system, but conditional on there not being a crisis we should expect expert net worth to increase faster, because the fear of crises pushes up conditional expected returns.

1.4 Unanticipated Crises

I now move on to analysing the equilibrium of the model and presenting results. Since the model is solved numerically, parameter choices are important, and I thus calibrate the model. I calibrate the model to provide a reasonable description of the data outside of crises times, as I detail below. However, given the rarity of financial crises I do not calibrate the model to match exact properties of financial crisis data. Instead, I provide results for a range of calibrations (specifically, a range of crisis frequencies) and present results across this range. In this section I consider the model with what I call “unanticipated crises”. This is the solution to the model where I set $\rho_e = 0$ so that the sunspot places no weight on selecting a crisis. Nonetheless, the model will still occasionally venture into the region where crises are possible. The model will then never select the crisis equilibrium, and agents will correctly anticipate this in their expectations. Solving the model like this first has the advantage of allowing me to investigate what forces drive the model into states of the world where crises are possible without interference from the effects of crisis fear on the equilibrium itself.

1.4.1 Calibration

I choose the following parameters for the baseline calibration. One unit of time is set to one year. I set $\rho_h = 0.05$ to generate an annual risk free rate of 5%. The deprecation rates are both set to $\delta = \underline{\delta} = 0.05$. Expert productivity is set to $a = 0.1$, which can be considered a normalisation.[10] The remaining non-crisis parameters are set according to the following calibration strategy. Firstly, I attempt to match the following three moments of the US post-war data: 1) a mean quarterly growth rate of logged, HP-filtered[11] output of 0.46%, 2) a standard deviation of quarterly, logged, HP-filtered output of 0.0167, and 3) a standard deviation of the quarterly, logged, HP-

[10] This measures the number of output goods produced from one unit of capital, and thus appropriate scalings of the other model parameters can be found such that different values of a correspond to different definitions of the “number” of output goods.

[11] I use the standard smoothing parameter of 1600, corresponding to quarterly data.

filtered investment-to-capital ratio of 0.0692. The sources for these data are detailed in the data appendix of my second chapter. I choose parameters such that the model without crises ($\rho_e = 0$) generates moments close to these data. Finally, I target that the model with $\rho_e = 0$ spends 10% of its time in states where crises become possible. This part of the calibration is somewhat arbitrary, but ensures that the model is able to generate crises, which is necessary for my exercise.

I parameterise the investment adjustment cost function as quadratic:

$$\iota_t = \Phi_t + \frac{\kappa}{2}\left(\Phi_t - \bar{\Phi}\right)^2 \tag{1.34}$$

κ measures the degree of adjustment costs ($\kappa = 0$ corresponds to no adjustment costs), and $\bar{\Phi}$ is the reference investment rate away from which adjustment costs are paid, which I take to be the average investment rate. I set $\bar{\Phi} = 0.07$ to match the targeted growth rate of output, which requires an investment rate of approximately 0.07. There are four remaining non-crisis parameters: households' productivity, $\underline{a}$, experts' discount rate, ρ_b, the volatility of the capital quality shock, σ, and the degree of adjustment costs, κ. These are chosen to target the four data moments above, and I use a numerical minimisation routine to minimise the squared sum of deviations from the target by varying the four parameters.

This leads to the following parameter choices: $\underline{a} = 0.07456$, $\rho_b = 0.05089$, $\sigma = 0.02591$, and $\kappa = 2.1992$. These lead to the model generating the following values of the targeted moments, when calculated as they are in the data: 1) a mean quarterly growth rate of logged, HP-filtered output of 0.38%, 2) a standard deviation of quarterly, logged, HP-filtered output of 0.0174, 3) a standard deviation of the quarterly, logged, HP-filtered investment-to-capital ratio of 0.0756, and 4) the model spending 10% of its time in the region where crises are possible.

The remaining parameters to choose correspond to aspects of financial crises. For the remainder of the paper I consider the limit case of permanent crises, setting $\rho_r = 0$. This simplification allows me to quickly compute the crisis price from (1.9).

None of the results depend qualitatively on crises being permanent: the anticipation effects of rare but permanent crisis are much the same as the effects of more frequent but shorter crises. To demonstrate this, in the appendix I provide an example of the solution to the model with non-permanent crises. As of yet I have not set the sunspot parameter, ρ_e, or the exogenous destruction cost, χ, as they are not needed for the solution to the model with $\rho_e = 0$. I discuss their values in section 1.5.

Finally, it is worth noting that this calibration leads to Brunnermeier & Sannikov's (2014a) bimodal distribution effectively vanishing. Their paper argued that the economy could endogenously spend a lot of time in very low capitalisation states following bad enough sequences of technology shocks, resulting in a regime that looks a lot like a financial crisis. However, given my calibration the probability of entering this regime is extremely small (see the stationary densities in Figure 1.6), and the crises driven by multiple equilibria are thus effectively the only forms of crises in the model.

1.4.2 Unanticipated crises: proximate causes & crisis-prone region

Figure 1.1 presents the numerical solution to selected variables from the model in the baseline calibration. As expert capitalisation (η) falls, experts intermediate less of the capital stock (ψ falls) which causes the capital price to fall as more of it is held by unproductive households. As the price of capital falls, returns increase for both experts and households. This encourages experts to increase their leverage (ϕ increases).

The bottom right panel of Figure 1.1 plots $\underline{n}_t/n_t$, and crises are thus possible whenever this line is negative (of course, given that $\rho_e = 0$ agents think that crises will never be selected in equilibrium). Crises become possible whenever expert capitalisation, η, is low enough. To understand this result, recall that $\underline{n}_t/n_t$ being negative

Figure 1.1: Selected variables, model with unanticipated crises

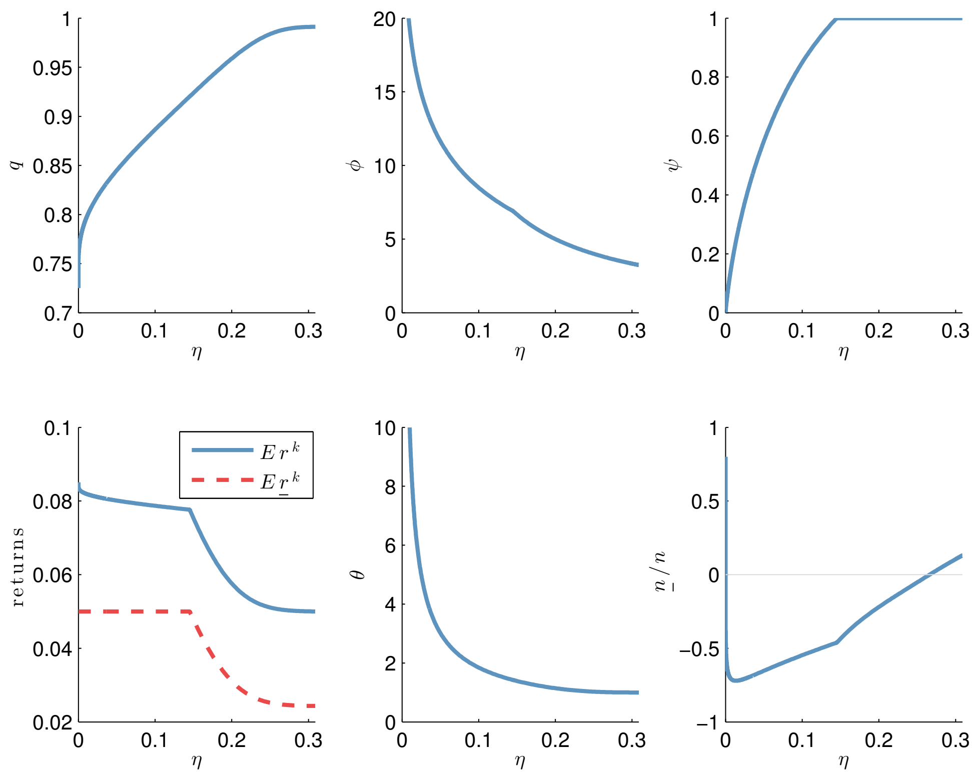

Solution to the model with $\rho_e = 0$. η gives expert net worth as a fraction of the total value of the capital stock. In all panels, η^* is the rightmost limit of the x-axis.

is equivalent to

$$1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right) < 0 \tag{1.35}$$

Since $1 - \underline{q}/q_t < 0$, higher leverage pushes us towards crises being possible, because banks have increased their exposure to the fall in q_t but financed this using fixed debt. For low values of η, where banks are relatively undercapitalised, leverage is high in equilibrium and this gives the model a tendency to predict crises as η falls.

On the other hand, crises are also only possible if the fall in asset prices is large enough, i.e. $\underline{q}/q_t$ is small. Graphically a crisis is a jump from the current value of η down to $\eta = 0$, and so the price falls from its current value down to the value at the intersection with the vertical axis. For low values of η, asset prices are already low, meaning that they don't have as far to fall in a crisis. This means that for a low enough η crises stop being possible. However, for this calibration this region is very small, to the extent that it is not visible in the plot.

The economy naturally gravitates towards higher values of η in the absence of shocks, and spends most of its time around the higher values of η near η^*. Hence in this calibration the economy is immune from crises near the steady state, and only becomes susceptible to crises if the experts become undercapitalised. This happens if a sequence of negative capital quality shocks erode expert net worth.

What does this tell us about the model where $\rho_e > 0$ and we actually do select crises in equilibrium? A crisis will happen in the crisis region, which the above logic tells us we will reach following a bad sequence of negative capital quality shocks. Once we are there, a crisis occurs if the "bad" sunspot is drawn, in which case the experts go bankrupt following a self-fulfilling fall in asset prices. In a crisis, η jumps from its current value down to zero, since aggregate expert net worth falls to zero.

1.4.3 Parameter sensitivity: fundamental causes of crises

Having established that crises become possible when leverage is high and asset prices have far enough to fall, and having identified the region of the state space where crises are possible, I now turn to their fundamental causes. What I mean by this is that I conduct parameter sensitivity to understand what kind of economies are more or less prone to financial crises. I conduct sensitivity for three key parameters, which all have important and economically interesting effects on the size of the crisis region.

The first panel of Figure 1.2 plots $1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right)$ across the state space for three different values of household productivity, $\underline{a}$. The other parameter values are all held at their baseline values. The thin, blue line in all three panels is the baseline calibration, and the dashed red and thick green lines are for a lower and higher value of the parameter respectively. For $\underline{a}$ we see that, holding all else constant, increasing

Figure 1.2: Parameter sensitivity: crisis net worth across changes in three parameters

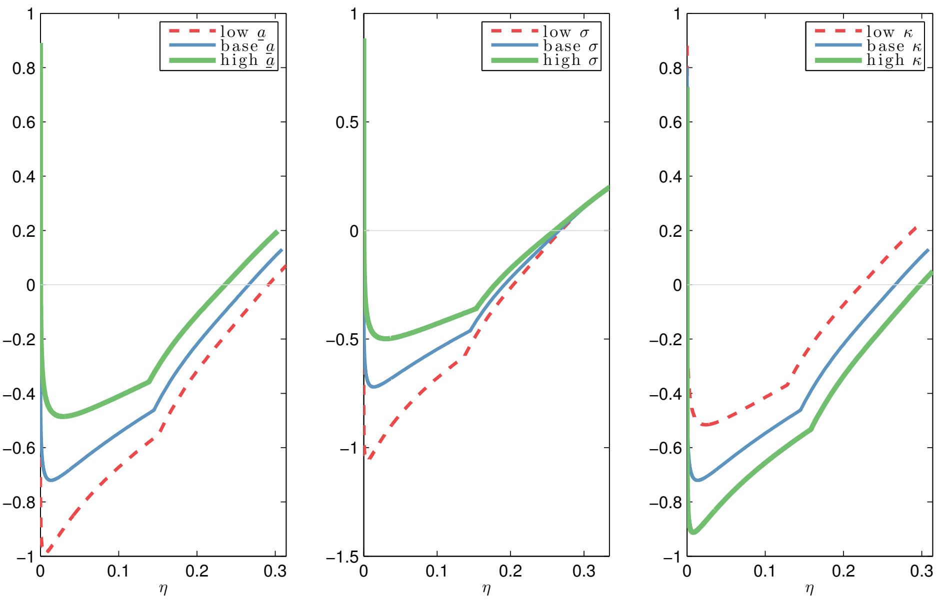

Each panel plots $\underline{n}/n$ across the state space, η. A negative value of any line means a crisis is possible at that η. Each panel calculates this variable for three values of a given parameter, holding all other parameters at their baseline values. For $\underline{a}$. the high and low values refer to 5% deviations from the baseline value, 15% for σ, and 30% for κ.

$\underline{a}$ reduces the size of the crisis region, and decreasing a increases it. This corresponds directly to reducing and increasing the probability of being in the crisis region, as can be seen in Table 1.1 where I give the fraction of time spent in the crisis region under the stationary distribution for each deviation.

I showed in the previous section that crises occur because either 1) leverage is high or 2) the price of capital is high relative to the crisis price $\underline{q}$. Figure 1.11 in the appendix shows that it is the latter effect that is operating here. Increasing $\underline{a}$ increases the productivity of households and hence increases the fire sale price of capital. This makes crises less severe, reducing the range of states for which leverage is high enough to enable a crisis. Figure 1.11 shows that leverage is actually slightly higher after increasing $\underline{a}$, and hence moves in the wrong direction to explain the reduced size of the crisis region. Hence the productivity differential between a and $\underline{a}$ is important in explaining the size of the crisis region because it controls how far prices have to fall in a crisis.

Table 1.1: Fraction of time spent in crisis region

	$\underline{a}$	σ	κ
Low	0.26	0.41	0.03
Baseline	0.10	0.10	0.10
High	0.04	0.04	0.28

Fraction of time spent in crisis region under stationary density for deviations of three parameters from baseline values when all other parameters are held at baseline. For $\underline{a}$. the high and low values refer to 5% deviations from the baseline value, 15% for σ, and 30% for κ.

The second panel demonstrates the effect of fundamental uncertainty (σ) on the crisis region. Increasing fundamental uncertainty reduces the size of the crisis region, and makes entering the crisis region less likely under the stationary distribution. Figure 1.11 shows that, in this case, the effects of leverage and prices both work in the same direction: higher fundamental uncertainty causes experts to deleverage and reduces asset prices, both of which make the system safer from crises.

Finally, the third panel demonstrates the effect of the adjustment cost parameter κ on the crisis region. The larger this parameter is the more costly it is to adjust your capital stock, and hence the less sensitive is investment to the capital price. Or conversely, with high values of κ the capital price is more sensitive to investment. This means that equilibrium price function is steeper when κ is high, as overall investment is increasing as experts get richer. Given our earlier discussion, a steep price function makes crises more likely as it increases the gap between q and $\underline{q}$, even if $\underline{q}$ is fixed. The sensitivity of the crisis region to κ was discussed in Brunnermeier & Sannikov (2014b), who interpret κ as technological illiquidity.

The effects of $\underline{a}$ and σ on the size of the crisis region are also interesting because they mirror earlier results in the original BrS paper. In the original model, reducing $\underline{a}$ increases the level of endogenous risk the model generates. For example, the volatility of the capital price is higher as you reduce $\underline{a}$. My result thus complements the original, by showing that not only is the volatility higher as you reduce $\underline{a}$, so is the risk of a crisis. Similarly for σ, BrS show that as you reduce the level of exogenous risk, the level of endogenous risk increases. I show that as you reduce σ the risk of a crisis increases, complementing the original result.

1.5 Anticipated crises

Having established the mechanics behind crises, I now move on to the general equilibrium effects of anticipating crises. At this point I need to give values for the final parameters of the model, which control various aspects of the crisis. I solve the model setting $\chi = 0$, so there are no exogenous costs of default. This choice is motivated by the fact that since in this model all experts default at the same time, any exogenous default costs would involve a large part of the economy's capital stock being exogenously destroyed. Instead, by restricting $\chi = 0$ I focus on the case where the *value* of the economy's capital stock falls, driven by price effects.

The final parameter to set is the probability of coordinating on a crisis, ρ_e. This parameter controls how likely crises are in the model. Recall that in the baseline calibration crises are permanent, so ρ_e can also be thought of as controlling the expected time until the economy permanently enters the crisis state. Choosing this parameter is tricky for two reasons: Firstly, financial crises are rare, so an appropriate measure of their forward-looking frequency is, by nature, a tough empirical exercise. Secondly, the model features permanent crises, and so I need to choose a relatively low probability of having a crisis to compensate for how severe the crisis state is. I settle for a value of $\rho_e = 0.1$, which implies that the expected time until the economy enters the crisis state is 357 years.[12] This is obviously extremely high, but compensates for the severity of the permanent crisis. In section 1.5.5 I discuss the effects of varying ρ_e, and in the appendix I solve a version of the model with non-permanent crises.

1.5.1 Price & crisis region

Figure 1.3 gives the solution to key model variables for the model with anticipated crises ($\rho_e > 0$) and without ($\rho_e = 0$, which thus repeats the results of the previous section). The top left panel shows us that the price of capital is globally lower when agents anticipate that a crisis is possible. This is intuitive, since the possibility that the price will fall in the future will be reflected in a lower price today. Given that asset prices control the level of investment, this will have important consequences for growth as we shall see later. The lower middle panel plots $1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right)$. Crises are possible if this quantity is less than zero. We thus see that anticipation of crises reduces the size of the crisis region relative to the solution with $\rho_e = 0$, where agents don't anticipate crises, because the red, dashed line is negative for less of the state space than the solid blue line. This can also be seen by comparing the fraction of time spent in the crisis region, which is 0.10 for the model with $\rho_e = 0$, and falls to 0.036 for the model with $\rho_e = 0.1$.[13] Thus the model pushes back against crises once you allow agents to anticipate them. The shrinkage of the crisis region reflects either a rise in $\underline{q}/q_t$ driven by the fall in q_t, or a decrease in leverage. However, in the bottom left panel we see that leverage actually increases compared to the model without crises and hence the effect is driven by the fall in asset prices.

[12] This is calculated starting from an initial value of $\eta_0 = \eta^*$. Details of the computation and simulation procedure are provided in the appendix.

[13] The fraction for the model with $\rho_e = 0.1$ is calculated using the stationary density excluding crisis realisations.

Figure 1.3: Model solution with and without anticipated crises. Key variables

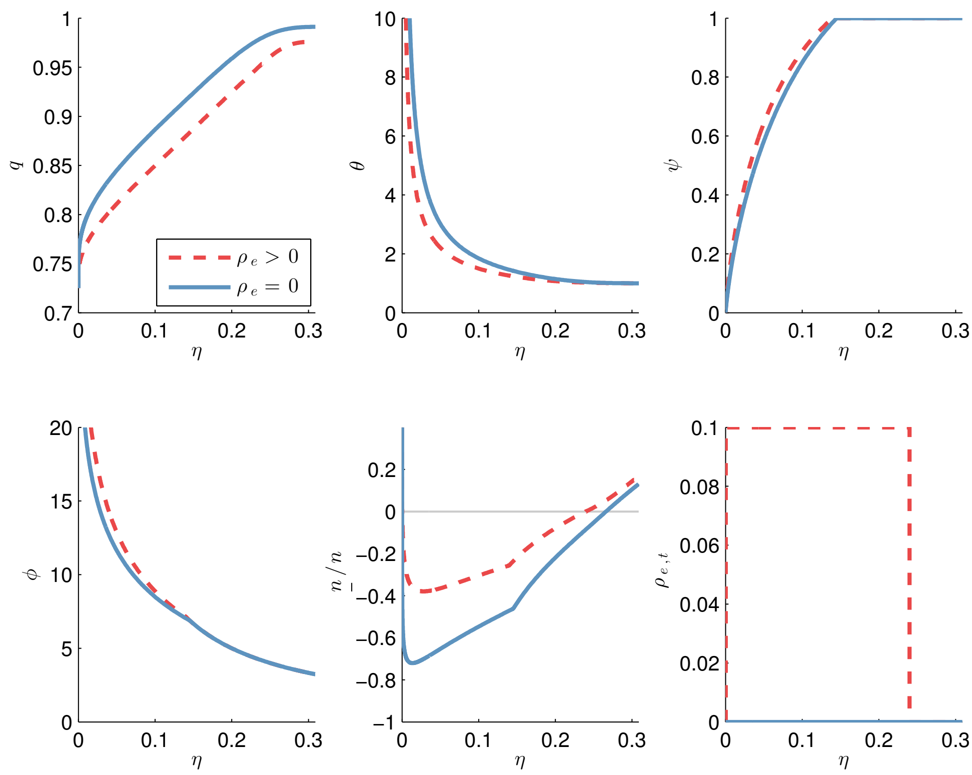

The dashed red line gives the solution to the model with the baseline positive value of ρ_e, meaning the economy will eventually experience a crisis. The solid blue line gives the solution where $\rho_e = 0$ and agents never coordinate on a crisis.

1.5.2 Leverage

Before discussing the main results, it is instructive to understand the effect that the fear of crises has on leverage. The model solution reveals that leverage is higher across the state space in the model where agents anticipate crises. This solution is for the case without exogenous default costs ($\chi = 0$) so the as the earlier discussion of the equations in Section 1.3.11 suggests, we should not expect financial crises to give experts any immediate reason to deleverage. In this section I discuss the general equilibrium forces which lead to higher equilibrium leverage.

The key is that the fear of crises tends to increase expected returns *conditional on there not being a crisis* by reducing asset prices. Even though experts are exposed to the downside risk of crises, they do not take this into consideration since they know they will go bankrupt in these states anyway. Thus from their point of view the main change between the two models is the higher expected returns in the model with anticipated crises, and they leverage up to take advantage of this.

Alternatively, we can see this by looking again at equation (1.26), which is the equation which determines leverage:

$$\frac{a-\underline{a}}{q_t}+\underline{\delta}-\delta-\rho_{e,t}\chi\frac{\underline{q}}{q_t}+\sigma_t^\theta\left(\sigma+\sigma_t^q\right)=0$$

Lower asset prices increase the return differential between experts and households by increasing $\frac{(a-\underline{a})}{q_t}$. This increases the advantage that experts have over households, and requires us to create a disadvantage for experts to return them both to indifference. This is done by increasing leverage, which increases the amount of endogenous risk.

Thus we see that experts lean in to crises, rather than reducing leverage to try and avoid their exposure to them. This is due, of course, to the assumptions that make avoiding crises unprofitable for experts. If we add exogenous default costs ($\chi > 0$) then it is possible that experts might deleverage in some regions of the state space, because exogenous default costs (which are only paid if capital is in the hands of experts) provide us an incentive to put capital in the hands of households instead. I present the solution for $\chi = 0.25$ in the appendix for comparison with the baseline model.

1.5.3 Investment & growth

My first main result is the effect that financial crises have on growth. The left panel of Figure 1.4 plots the policy function for the investment rate, Φ_t, in the models with and without anticipated crises. This picture mirrors the effect that crisis fear had on asset prices, which is unsurprising given the tight link between asset prices and investment in the model. Note that due to the linearity of the production functions in capital, the investment rate becomes independent of the capital stock, and the model features long-run, endogenous growth. The investment rate, along with the depreciation rate, then pins down the growth rate of the capital stock according to the expert and household accumulation equations, (1.1) and (1.2). Aggregating these two equations when $\delta = \underline{\delta}$, as is true in my calibration, yields the evolution of the total capital stock:

$$dK_t = (\Phi(\iota_t) - \delta) K_t dt + \sigma K_t dZ_t \tag{1.36}$$

In the model where agents anticipate crises, the investment rate is everywhere lower because asset prices are always lower. Given the lower asset price, there is less incentive to invest since the (static) profits from investing are lower. Alternatively, we could could think about this as a rough discounted sum. A crisis scenario involves handing the capital stock to the household to intermediate. Given the household's low productivity this means that the discounted sum of profits from investing are lower, leading to lower investment. The growth effects of expected crises are illustrated in the second panel of Figure 1.4. Here I simulate the models, plotting a sample path for output ($Y_t = (a\psi_t + \underline{a}(1 - \psi_t))K_t$). The widening gap between output, driven by the under-accumulation of capital in the crisis economy is clear. For the current calibration, the average quarterly growth rate of output falls from 0.38% to 0.19%. Growth rate effects are likely to have much larger effects on welfare than level effects, and the halving of growth rates adds up over a long enough horizon. Given that other studies have focused on level effects of crises, this suggests an important alternative motivation for policy to address financial crises. Additionally, it could present an explanation for the slow recovery in the aftermath of the recent financial crisis, a time where fear of a repeat was surely escalated, as I discuss further in section 1.5.5.

Figure 1.4: Investment policy function & output sample path

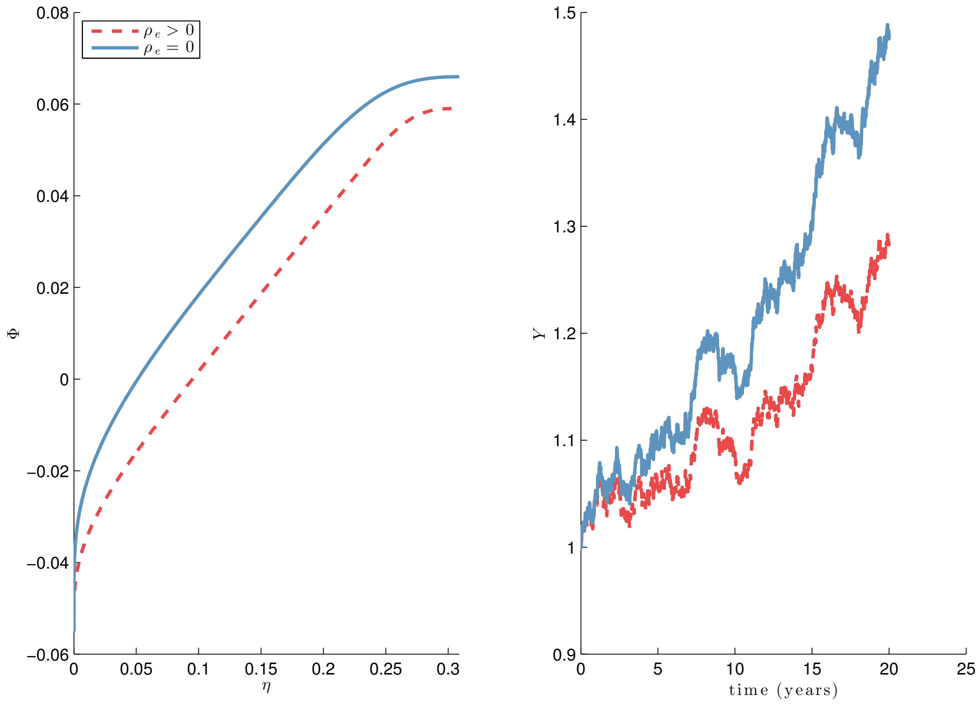

Left panel plots the investment policy function in the models with ($\rho_e > 0$) and without ($\rho_e = 0$) crises. Right panel plots a simulated time path for output for both economies, normalised to one in the first year for both. The path for the economy with crises is constructed so that no crises occur during the sample path.

1.5.4 The financial crisis accelerator

In this section I discuss the financial crisis accelerator. This is the result that the fear of crisis increases the endogenous volatility of the economy. Specifically, in some regions the fear of crises makes the economy more responsive to the exogenous capital quality shock. This can be seen in the first two panels of Figure 1.5 where the volatility terms for both q_t and η_t are plotted. These volatilities give the impact responses to the capital quality shock dZ_t.

The results are very state dependent: for both q_t and η_t, the increase in volatility is larger in the middle of the state space. What drives this result? As explained by

Figure 1.5: The financial crisis accelerator

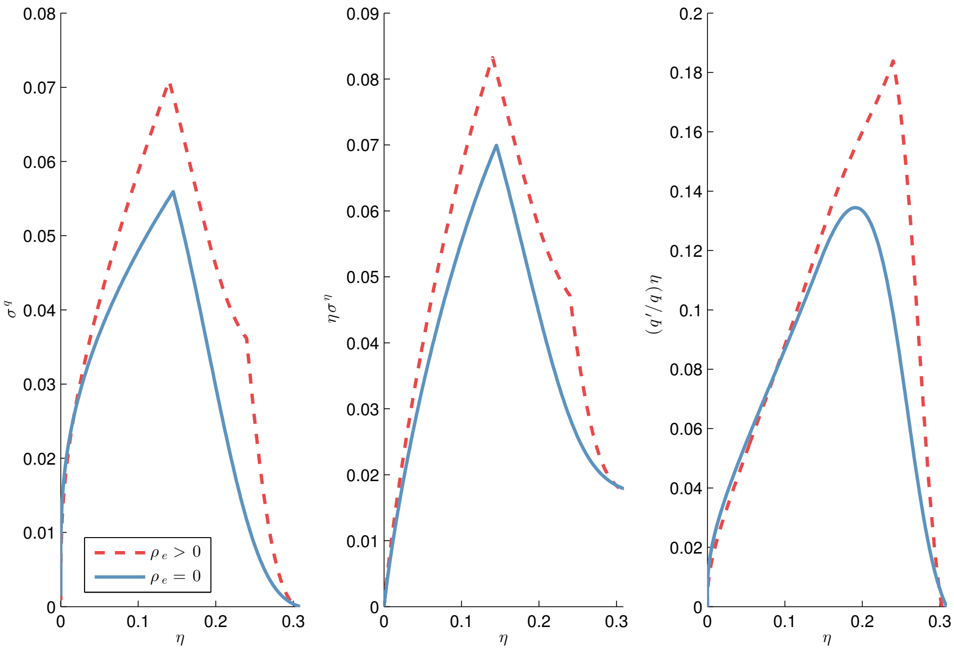

The first two panels plot the loading on the Brownian motion of the capital price and state variable, η. The final panel plots the dependence of σ^q on $\eta\sigma^\eta$, as described by (1.38).

BrS, the financial accelerator in continuous time can be understood as the interaction between the volatility terms of η_t and q_t. For example, a negative shock reduces bank net worth, reducing η_t. But reducing η_t reduces the price q_t as more capital is intermediated by inefficient households. The reduction in q_t further reduces η_t, and the cycle continues. This can be seen by the interdependency between the equations for σ_t^q and σ_t^η (which are derived in the appendix):

$$\sigma_t^\eta \eta_t = \eta_t(\phi_t - 1)(\sigma + \sigma_t^q) \tag{1.37}$$

$$\sigma_t^q = \frac{q'(\eta_t)\sigma_t^\eta \eta_t}{q_t} \tag{1.38}$$

Each depends on the other, and solving the two together gives the solutions for the volatilities in (1.27). The slope of the price function, $q'(\eta_t)$ is important because this tells us how much prices are going to fall in response to a marginal fall in η_t. A steep price function thus gives a severe financial accelerator because prices fall a lot in response to a fall in net worth, making the secondary effect on net worth larger. The other determinant of the size of the multiplier is the current leverage of the experts, ϕ_t. High leverage means that experts have a large exposure to q_t, and makes their net worth more sensitive to changes in asset prices, worsening the financial accelerator.

The changes in the volatility terms between the two solutions can thus be understood by appealing to the changes in leverage and the slope of the price function caused by the fear of crises. As previously mentioned, the model with anticipated crises has higher leverage than the model without, which thus contributes to the increase in volatility. The changes in the slope of the price function are plotted in the last panel of Figure 1.5. In the central region where the volatilities are increased most relative to the model without crises, we see that the slope of the price function is increased.

This region overlaps closely with the region where crises are possible. As we move deeper in to the crisis region we expect to remain there for longer, placing more and more downwards pressure on asset prices because they might suddenly fall if we experience a crisis. This makes q_t very sensitive to our position in the state (η_t) around this region. This is the essence of the financial crisis multiplier: shocks that push the economy closer to (or deeper into) the crisis region will push down asset prices a lot as agents anticipate a possible crash, making the standard financial accelerator more powerful.

It is only in and near to the crisis region that the financial crisis accelerator emerges. This also creates an intuitive asymmetry in the model: starting from the steady state, the financial accelerator is worse in response to negative shocks than it is to positive shocks. This is because a series of negative shocks bring us closer to the crisis region, prompting asset prices to fall faster, harshly eroding net worth and so on. In response to positive shocks we move further away from the crisis region. The probability of crisis in the near future was already close to zero, and remains so, leading to smaller changes in asset prices and a smaller financial accelerator.

Table 1.2 gives the standard deviations of output and the investment rate across several values of ρ_e. This reveals that the financial crisis accelerator effect is rather large for the investment rate (and hence also for asset prices) and quantitatively less important for output itself. In particular, going from the baseline value of ρ_e down to zero more than halves the volatility of investment, but only reduces the volatility of output by around 3%.

Finally, it is worth noting that these results again echo results in the original BrS paper. In their paper they show that the financial accelerator can be made quantitatively more powerful by understanding that their model is prone to occasional prolonged periods of financial distress. The model features a bimodal stationary distribution, where a sufficiently bad series of exogenous shocks can lead to the economy getting trapped with low net worth (low η_t) for a long time. This possibility is what allows asset prices to fall a lot in response to negative shocks, as agents anticipate that this outcome becomes more likely. My result thus complements theirs, because I show that their financial accelerator can also be rationalised by appealing to crises of a self-fulfilling nature.

1.5.5 Stability & growth

In this section I demonstrate how varying the probability of coordinating on a crisis affects equilibrium. Clearly, reducing this probability will, by construction, reduce the likelihood of crises in the model. It also has intuitive effects on the growth rate and volatility of the economy. Table 1.2 gives various moments of the model across a range of values for ρ_e.

The first column gives the solution to the model where agents never coordinate on crises ($\rho_e = 0$), and successive columns increase ρ_e, up to the baseline value in the final column. The first row gives the expected time until the economy enters the crisis state, which is infinity by construction when $\rho_e = 0$, and falling as the probability of coordinating on a crisis is increased. The second row gives the effect on the average quarterly growth rate of output, which is decreasing as the probability of coordinating on a crisis is increased. The final two rows give the effects on the volatilities of output and the investment rate, which are increasing as the probability of coordinating on a crisis is increased, demonstrating the financial crisis accelerator.

Table 1.2: Effects of varying the probability of coordinating on a crisis

	$\rho_e = 0$	$\rho_e = 0.005$	$\rho_e = 0.01$	$\rho_e = 0.02$	$\rho_e = 0.05$	$\rho_e = 0.1$
TTC	∞	1301	904	629	428	357
g_y	0.38%	0.35%	0.32%	0.27%	0.21%	0.19%
σ_y	0.0174	0.0174	0.0175	0.0176	0.0178	0.0180
$\sigma_{I/K}$	0.0756	0.0804	0.1071	0.1079	0.1492	0.1711

TTC is the expected time time until the economy experiences a crisis. g_y refers to the average growth rate of quarterly GDP, σ_Y its standard deviation, and $\sigma_{I/K}$ the standard deviation of the quarterly investment ratio, I_t/K_t.

These results also suggest an interesting behavioural interpretation of the model. If agents' fears of a future crisis exogenously increase after experiencing a crisis (modelled by increasing ρ_e) then the economy will grow slowly, but will be less susceptible to crises in the immediate aftermath. If these fears then decrease if the economy doesn't experience a crisis, growth will gradually recover, and the economy will start becoming more volatile and more susceptible to crises again as time goes on.

1.6 Policy

I now turn to the policy implications of my model. I focus first on prudential (ex ante) policies which aim to reduce crisis risk by limiting expert leverage. I then discuss bailouts, which aim to reduce the ex ante perception of crisis risk by signalling the government's commitment to maintaining asset prices, and a potential market-based solution.

1.6.1 Prudential policy: leverage constraints

Minimally active leverage constraint

Since crises in my model are only possible for high enough expert leverage, policies which limit leverage ex ante are natural candidates for ruling out, or reducing the probability of, financial crises. Remember that a crisis is only possible at time t if equilibrium leverage (ϕ_t) and asset prices (q_t) satisfy

$$1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right) < 0 \tag{1.39}$$

Thus if the government imposes the following regulatory leverage constraint it can completely rule out the possibility of financial crises:

$$\phi_t \leq \bar{\phi}_t^{ma} \equiv \frac{1}{\left(1 - \frac{\underline{q}}{q_t}\right)} \tag{1.40}$$

I call this constraint the "minimally active" leverage constraint. It is active in the sense that it requires active monitoring and adjustment by the regulator: the leverage constraint depends on today's capital price and the capital price during a crisis. It is minimal in the sense that this is the loosest leverage constraint which completely rules out crises. Quantitatively, this constraint says that if asset prices are known to drop by a fraction $1/x$ during a crisis, leverage cannot be higher than x. So if asset prices are thought to drop by a quarter, leverage would be restricted to be no higher than 4. Note that this constraint is state dependent, via q_t.

How does the minimally active leverage constraint affect equilibrium? Trivially, it rules out financial crises. The question that remains is how does it affect the other features of equilibrium? This will be crucial in determining whether or not the policy is welfare improving. I will focus on two key aspects of equilibrium: volatility and growth. As discussed in the introduction, the prevailing view of the effects of prudential leverage constraints is that they would reduce volatility, but at the expense of reducing growth. I show that this is not true in my model.

The details of the model solution with a leverage constraint are relegated to the appendix, and I present only the results here. Given the typical discourse surrounding prudential policies, the results are surprising: the average growth rate of quarterly output rises to 0.38% after the introduction of the policy. This rise, from the original 0.19% of the model with crisis, brings the growth rate of the economy all of the way back up to the growth rate of the economy with no crisis risk. Hence ruling out crisis comes with the benefit of higher growth, not the cost of lower growth.

We can understand why by inspecting Figure 1.6. The top left panel plots the price of capital across the three models. For high levels of expert capitalisation, η, the minimally active leverage constraint increases the price of capital almost all the way from the its original price (dashed red line) to the price in the model without crisis risk (thin blue line). For lower values of η the benefits of the policy on asset prices are smaller. However, the economy spends very little time in this region under the stationary distribution. The intuition for the increase in prices is simple: by ruling out crises the policy removes the possibility of of prices jumping down in the future to $\underline{q}$, which increases prices today. Combined with the shifting right of the stationary distribution past even the distribution of the model without crises this leads to the increase in average prices.

The top right panel plots leverage across the three models, showing that the leverage constraint, by construction, only has large effects on equilibrium leverage in the central region where crises were originally possible. General equilibrium effects lead to small changes outside of this region. Of course, this raises the question of how experts are able to fund increased investment while having their leverage reduced. Consider a region of the state space where experts hold all of the capital stock. Then

Figure 1.6: Effects of the minimally active leverage constraint

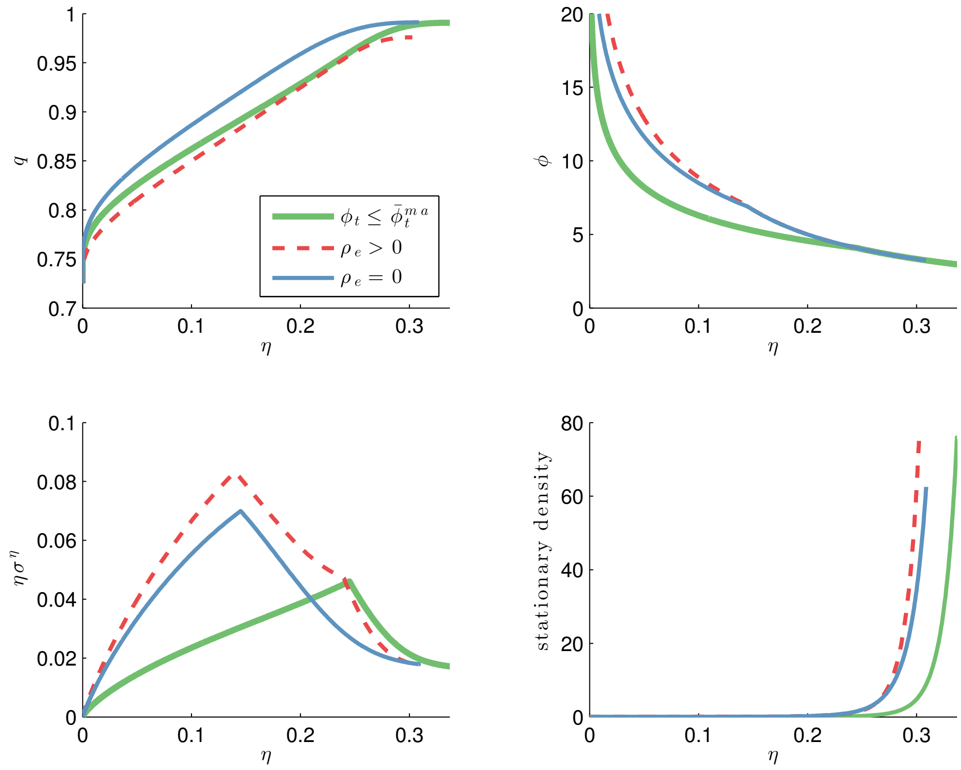

Model solution in models with (dashed red) and without (thin blue) crises, as well as model with minimally active leverage constraint (thick green). Stationary density in the model with crises is calculated ignoring crisis realisatons.

the definition of expert leverage implies that $q_t K_t = \phi_t N_t$. For a given capital stock, this shows us that higher asset prices can only be supported following a reduction in leverage if expert net worth increases *more* than the fall in leverage. The investment first order condition implies that investment is fully tied down by the price of capital, and hence that investment and growth can also only increase if equity rises sufficiently.

Thus an increase in equity, compensating for lower debt, is key to generating increased growth following the implementation of a regulatory borrowing constraint. Is this a reasonable thing to expect? In the model, experts cannot raise equity, and the increase in equity is thus funded by experts paying out net worth (as consumption) less often. This implies that the policy actually *raises* the value to experts of retaining earnings. This is an intuitive idea: if experts know a crisis is coming they have incentive to consume now in order to consume their net worth before it is lost in a crisis. Hence leverage policy could encourage equity by making the financial sector safer, reducing the incentive to withdraw equity as dividends instead of investing it. This can be seen in the bottom right panel, which shows that expert capitalisation is higher under the policy than without it, since the stationary density is shifted to the right.

Finally, the bottom left panel plots the volatility of the aggregate state, η across the three models. In the central region, the leverage constraint reduces volatility relative to both the models with and without crises. From the discussion in Section 1.5.4, this decrease is a direct consequence of the decrease in leverage, which reduces the financial accelerator in that region. However, this reduction is not across the whole state space, and outside of this region there is actually a small increase in volatility. This is due to the increased slope of the price function in certain regions, which leads to a slight exacerbation of the financial accelerator.

Putting all of these effects together makes a strong case for the minimally active leverage constraint in this model: it rules out financial crises, increases growth, and reduces volatility in most regions of the state space. With this in mind, I now turn to looking at the effects on welfare of the policy. The total welfare of experts at a given time is $W_t^e \equiv \eta_t \theta_t q_t K_t$ and that of households is $W_t^h \equiv (1-\eta_t) q_t K_t$.[14] Since the model features two classes of agents there is no single welfare criterion that we can use, so I first examine the impact of the policy on the welfare of each group individually. As a total welfare criterion I select total welfare $W_t \equiv W_t^e + W_t^h$.

Proposition 2. *Any policy which increases the price of capital, q_t, on impact increases household welfare if households are holding capital ($\psi_t < 1$) and leaves house-*

[14] To see this for experts note that an individual expert's maximised value is $\theta_t n_t$, and that the total net worth of the expert sector is $\eta_t q_t K_t$. Since households are risk neutral their welfare is just their net worth, which totals $(1 - \eta_t) q_t K_t$.

hold welfare constant otherwise. Holding the other constant, policies which increase q_t or θ_t on impact increase expert welfare, and expert welfare increases iff

$$\theta_t'\left(1-\phi_t\left(1-q_t'/q_t\right)\right)>\theta_t$$

where primed variables are post policy. Following the unanticipated implementation of a policy, q_t and η_t immediately jump to the values which solve:

$$q_t' = q_{pol}(\eta_t') \tag{1.41}$$

$$\eta_t' = \left(1-\phi_t\left(1-\frac{q_t'}{q_t}\right)\right)\frac{q_t}{q_t'}\eta_t \tag{1.42}$$

The proofs are simple algebra and are omitted. The last part of the proof points out that the implementation of any policy will have an immediate impact on prices, and hence on expert capitalisation. Thus when we evaluate the welfare impact of a policy, we need to take into account two things: Firstly, equilibrium welfare as a function of the state η_t in the new policy regime. Secondly, today's state η_t and how that translates into our state η_t' after the policy is implemented. Welfare varies across the state space, and so whether or not implementing the leverage policy improves welfare looking forward could in principle depend on the state of the economy today. To address this, in Figure 1.7 I plot W_t, W_t^e and W_t^h across the state space in the model with crises and the model with the minimally active leverage constraint, normalising K_t to one. Thus it is important to note that these are welfare changes *on impact* given the current state of the world, and take into account agents expectations of the future evolution of the economy under the equilibrium distribution of the state.

The figure reveals that, regardless of the state today, the policy both increases total welfare and weakly increases the welfare of both sets of agents. It is hence Pareto improving across the whole state space, as well as being a policy that both sets of agents would support. This is because crises hurt both agents, and they are both happy to see them removed. Crises mean that experts will eventually go bankrupt,

Figure 1.7: Welfare: minimally active leverage constraint

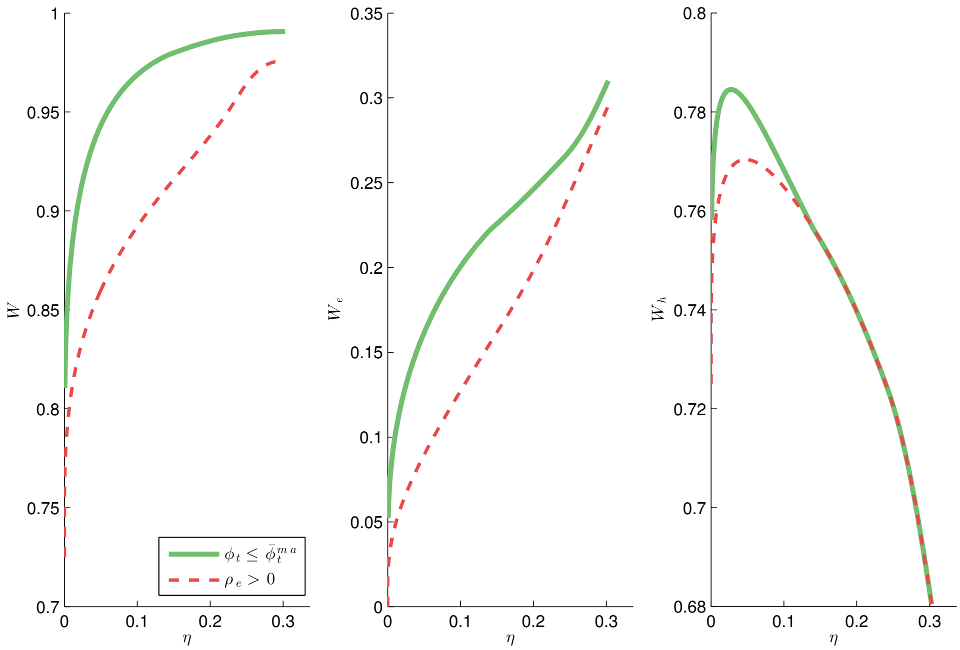

The left panel plots $W(\eta)$ across the state space in the model with crisis risk, and the associated welfare on impact if we are in state η today and the policy is implemented. Hence the thick green line plots $W_{pol}(\eta'(\eta))$ where η' is calculated for every η as in Proposition 2. If the solid green line is above the dashed red line the policy improves welfare on impact. The centre and right panels plot the same for W_e and W_h respectively.

losing their ability to generate net worth. They also mean that households, who will eventually have to intermediate capital, will suffer as the economy operates at lower productivity. Notice that towards the right of the state space, where households don't hold capital, they are indifferent about the implementation of the policy on impact. This is because in this region they do not hold any capital, and hence do not realise any increase in their net worth on impact. Since their welfare is simply their net worth, they also do not realise any increase in their welfare.[15]

[15] Of course, their welfare could change over time as the economy evolves, and will indeed evolve differently with and without policy. However, this is taken into account when computing their current welfare, since it is a forward looking measure.

This result, that the minimally active leverage constraint increases welfare for both agents, is a direct consequence of thinking about the costs arising from financial crises. To see this, note that the same policy actually reduces welfare if it is implemented in an economy without financial crises ($\rho_e = 0$). Of course, this suggests an intuitive condition for the minimally active leverage constraint to increase welfare: crises must be *sufficiently likely*. In particular, there is a threshold likelihood of financial crises above which the policy increases welfare, and below which it does not.

Implementation issues

One issue with the minimally active leverage constraint is that while it improves welfare, it requires a lot of information for the government to implement it correctly. It requires knowledge of current asset prices *and* how far asset prices would fall in a crisis. To address these issues, in this section I investigate how deviations from this policy affect its effects. In particular, I consider a government who attempts to implement the minimally active leverage constraint, but instead accidentally implements:

$$\phi_t \leq \bar{\phi}_t^x \equiv \frac{x}{\left(1 - \frac{\underline{q}}{q_t}\right)} \tag{1.43}$$

For some $x > 0$. This policy is essentially a slightly tighter ($x < 1$) or looser ($x > 1$) version of the minimally active leverage constraint. The policies which are tighter deliver no extra benefit in reducing the probability of there being a crisis, since this has already been driven to zero. I plot the effect on policy of selected variables across a range of values of x in Figure 1.8. Since I am now comparing across many model solutions, I restrict myself to focusing on the effect of policy for a specific starting value of the state today, which I take to be the ergodic mean under the model with crisis and without policy.

The top left panel shows that the improvement in total welfare is maximised (within this class of policies) by choosing the minimally active leverage constraint:

Figure 1.8: Effects of badly implemented policy, selected variables

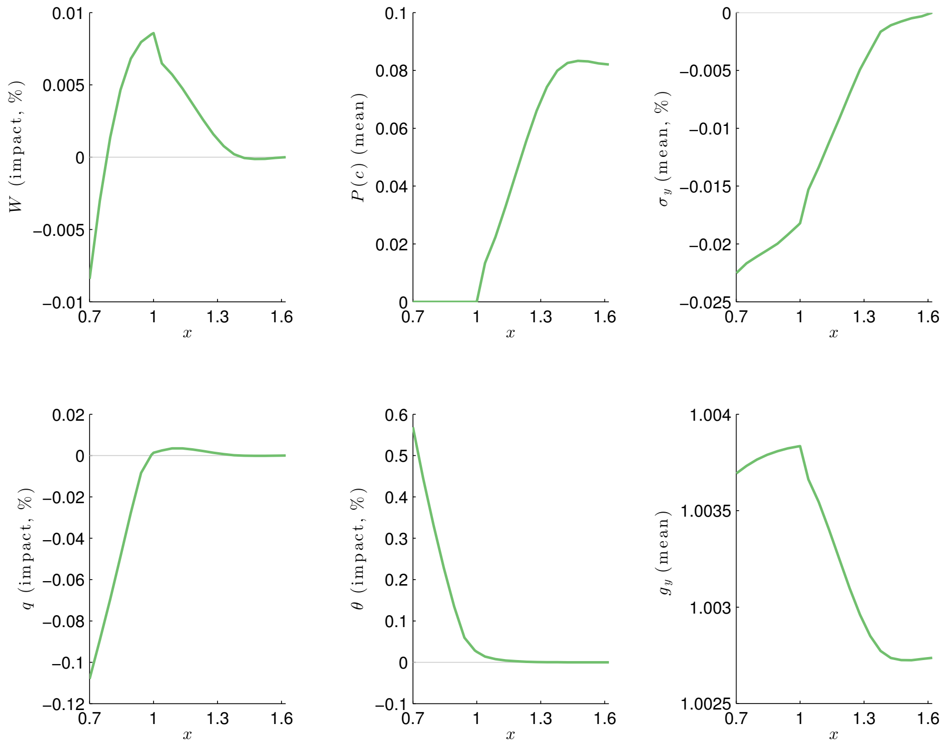

All variables computed from the same value of η, which I take as the ergodic mean of the model without policy. Variables denoted "impact, %" are impact changes, computed in fractional deviations from the value without policy. Variables denoted "mean" are main values from simulations of the model with the corresponding policy (ignoring crisis realisations). $P(c)$ denotes the fraction of time spent in the crisis region.

$x = 1$. Note that at the value of η at which I am making these comparisons, households do not hold any capital, and hence see no welfare gains. Thus the entire welfare gain is driven by an increase in the welfare of experts. Policies which are slightly too tight (x less than but close to 1) still improve welfare, but by less. This reflects the additional distortions that are introduced by leverage constraints which are too tight. However, for these parameter values the policymaker has room for error before policy becomes actually harmful, and policies can be up to 20% too tight and still deliver welfare benefits. In the other direction, policies which are looser than the minimally active constraint still deliver welfare benefits, but they are again smaller.[16] If the policy is so loose that it does not bind at all then the policy trivially delivers no welfare benefits.

The bottom left and centre panels decompose the welfare gain into its q_t and θ_t components. This reveals that the welfare cost of overly tight policies derives from an instantaneous lower asset prices (which are a reflection of the lower present value of output). This is partially offset by increases in the value of net worth to experts, which reflects the gain in instantaneous profits they can make as arbitrage is restricted.

The remaining three panels show how the various policies affect the moments of the economy. The top centre panel shows that tighter policies reduce the fraction of time the economy spends in the crisis region, which eventually falls to zero when $x \leq 1$. The top right panel shows that tighter policies reduce the volatility of output. The bottom right panel shows that, starting from a high $x > 1$, tighter constraints increase average output growth. But average growth is maximised for $x = 1$, and further tightening of the constraint starts to erode the gains to growth.

Overall, these results highlight an interesting *inverse-U* shape in the welfare gains from prudential leverage constraints. Policies which are too loose trivially deliver little or no welfare gains. Intermediate policies deliver welfare gains by ruling out crises, and this benefit is maximised once the probability of having a crisis is reduced to zero. Beyond this point, extra tightness is welfare reducing since it distorts the intermediation of capital without delivering any extra crisis-reduction benefits. The model thus emphasises a role for leverage policy, but also stresses caution in its use.

[16] Note that there is a discontinuous jump up in welfare as x approaches one from above. This is due to the nature of the exercise. Any value of x which is an ε above one has a positive probability of experiencing a crisis, which converges to a number other than zero as $\varepsilon \to 0$. $x = 1$ then delivers a probability of crises of exactly zero.

1.6.2 Ex-post policy: bailouts

Another commonly discussed and contentious policy instrument is ex-post bailouts. The idea behind bailouts is that the fundamental problem during crises is a lack of net worth in the financial system, and bailouts aim to fix this by directly injecting net worth (either for free or at a discounted rate). Bailouts are criticised mainly on ex ante incentive grounds, with the argument being that they incentive risk taking by reducing the punishment banks face when everything goes wrong. In my baseline model this trade off does not exist, and bailouts can be effective at completely ruling out crises without imposing any incentive distortions. The result is summarised in the following proposition:

Proposition 3. *Suppose that in the event of a crisis each expert is recapitalised to their original level of net worth, n_t. Then crises are not possible, and the model equilibrium is identical to the solution without crises (i.e. with $\rho_{e,t} = 0$).*

Proof. The recapitalisation policy rules out crises because it rules out jumps in net worth, and hence η_t: if it were to jump the recapitalisation policy simply jumps us right back to where we started. Since jumps don't happen in equilibrium, experts never receive any bailouts, and the model equations are identical to those with $\rho_{e,t} = 0$. □

This result is very similar in spirit to the original Diamond & Dybvig (1983) result that (in their baseline model) deposit insurance can improve allocations. In my model, bailouts promise to restore asset prices and net worth to their original level in the event of a crisis, completely removing the possibility of a crisis even happening because now agents have no reason to coordinate on the bad equilibrium. Since crises now never happen in equilibrium, experts can never receive any bailouts, which means that their incentives cannot be distorted by the possibility of receiving bailouts.

Of course, as in Diamond & Dybvig (1983), the result that this policy does not induce any distortions is special and due to the simplicity of the setup. Specifically, since bailouts are never needed in equilibrium (just as deposit insurance is never used in their model) there is no incentive to change behaviour to try and receive this bailout. Thus we would expect less successful or well targeted bailout policies, which actually led to bailouts being given in equilibrium (as they are in the real world) to induce ex ante distortions. For this reason we should remain sceptical of bailouts, and future work explicitly assessing the pros and cons is necessary. Indeed, my aggregate-capital externality could be viewed as a stand-in for other frictions, such as bailouts, which encourage banks to allow themselves to get in to trouble during a crisis.

1.6.3 Market based solutions

One possibility which the literature has started to address is that instead of government policies placing limits on the behaviour of the financial sector, the government could encourage the formation of markets to deal with the specific externalities involved.

The typical financial accelerator paper, including this one, assumes that lending is not contingent on aggregate state variables. This is what gives the accelerator power, since following aggregate shocks the value of assets can change dramatically, while the value of debt is fixed. If debt was allowed to be state contingent then the value of debt can also adjust to offset the change in the value of assets, protecting net worth and blunting the accelerator. Dmitriev & Hoddenbagh (2013) show that under the optimal (state contingent) contract, the financial accelerator disappears in the standard Bernanke, Gertler & Gilchrist (1999) model. In my third chapter I show that this is also true in the context of a model with a Gertler & Kiyotaki (2010) style borrowing constraint. Carlstrom, Fuerst, Ortiz & Paustian (2014) take an agnostic view on the degree of indexation of debt, and perform a structural estimation to pin down the value in the context of their model, again finding that higher indexation reduces the financial accelerator. Finally, Kilenthong & Townsend (2014) argue for market based solutions to price externalities in a general theoretical framework.

The takeaway from this literature is that if we are to believe in the power of financial frictions, we need to be confident that markets are sufficiently less than complete. While this may ultimately be an empirical question, I provide some additional theoretical insights here in the context of my model. The model solved above features only defaultable debt, which gives a less than fully state contingent set of assets to trade on. I show in this section that if we add a second "insurance" asset, which pays off during a crisis, then the same frictions which make anticipated crises possible also mean that no individual expert would be willing to take out insurance against the possibility of a crisis. This result this stands in contrast to the results above, and highlights a limit on the power of market based solutions.

In particular, consider the following insurance asset. If the asset is held from t to $t + dt$ and there is no financial crisis, the holder pays a premium $r_t^I dt$. If there is a crisis then the asset pays out one unit of the consumption good. This is a classic insurance contract over the event of a financial crisis happening. The household provides this insurance contract to the experts, who may choose any non-negative amount of insurance.

Proposition 4. *In the baseline model, an expert will never choose to hold positive amounts of the insurance asset, and the equilibrium with insurance is identical to the equilibrium with only defaultable debt.*

The proof is relegated to the appendix, and I discuss the intuition here. This insurance contract allows an expert to transfer wealth between future states of the world: do I want money tomorrow if things work out, or if there is a crisis? However, if anticipated crises are to exist in equilibrium, we require frictions which make the value of wealth to an expert during a crisis, $\underline{\theta}$, low. In fact, $\underline{\theta} = 1$ in the baseline model, its lowest possible value, because the aggregate capital externality reduces an expert's ability to produce if all other experts are bankrupt. This ensures that experts are willing to take on high enough leverage to allow themselves to go bankrupt during a crisis. What Proposition 4 establishes is that under the frictions which allow crises to happen in the baseline model, the addition of an additional insurance market is unable to provide any extra protection. The intuition is simple: a crisis is not a good time to have net worth, so experts have no incentive to use insurance to transfer wealth to that state of the world.

Of course, the above result relies as crucially on the assumption of the aggregate-capital externality as does the very existence of crises in my model. The point is that the conditions that make anticipated crises possible in my model are the very conditions which make the above market based solution infeasible.

1.7 Conclusion

In conclusion, I study the ex ante effects of the fear of future financial crises. I show theoretically that this "crisis fear" has both negative growth and business cycle effects. Financial crises push capital away from experts and towards less productive households, worsening the allocation of capital. Thus the possibility of future crises lowers the expected return on capital. This lowers asset prices, investment and growth today, even if experts are currently well enough capitalised to survive a crisis. The model features endogenous growth, leading to permanent effects of crises on growth. The externality that generates endogenous growth is also crucial for generating crises, by reducing the productivity of surviving experts in crises and hence encouraging them to overleverage and allow crises to happen in equilibrium. The possibility of future crises also creates a state-dependent "financial crisis accelerator" in which shocks which push the economy closer to crisis lead to more severe financial accelerator effects than those that push the economy away from crisis.

The model has implications for policy, and shows that explicitly taking into account agents' understanding that there could be future crises can overturn the received wisdom about the tradeoffs of prudential policy. In particular, in my model, restrictions on expert leverage can remove the possibility of financial crises and simultaneously increase growth. This is in contrast to the standard view that leverage constraints should reduce growth by restricting the ability of the financial sector to intermediate funds. While this effect still operates in my model, leverage constraints also encourage growth by making the system safer and promoting the retention of net worth by financial institutions. This strengthens the case for prudential policy, and future quantitative work should address the importance of this effect relative to the traditional growth-harming effects of prudential policy in richer model structures.

Appendices

1.A Endogenous growth

We can consider the linear production functions as the reduced form of a simple endogenous growth model. In particular, consider the experts' production function $y_t = ak_t$. Suppose instead that experts produce according to $y_t = z\hat{K}_t^{1-\alpha}k_t^{\alpha}l_t^{1-\alpha}$, where l_t is their labour choice, and $\hat{K}_t$ is aggregate expert capital, which an individual takes as given. There is thus a capital externality: experts don't take into account that their capital choice affects the productivity of other experts. Experts hire labour at wage w_t. Assume that households inelastically supply one unit of labour to experts, and one unit to households.

An expert chooses labour to maximise static profit: $\pi_t = z\hat{K}_t^{1-\alpha}k_t^{\alpha}l_t^{1-\alpha} - w_t l_t$. This yields the first order condition $w_t = (1-\alpha)z\hat{K}_t^{1-\alpha}k_t^{\alpha}l_t^{-\alpha}$. After optimising labour, an expert's profit function becomes linear in individual capital:

$$\pi_t = \left((1-\alpha)^{\frac{1-\alpha}{\alpha}} - (1-\alpha)^{\frac{1}{\alpha}}\right) w_t^{\frac{\alpha-1}{\alpha}} z^{\frac{1}{\alpha}} \hat{K}_t^{\frac{1-\alpha}{\alpha}} k_t \tag{1.44}$$

Imposing market clearing ($l_t = 1$) and $\hat{K}_t = k_t$ this profit becomes $\pi_t = \alpha z k_t$. Given the linearity of both the individual and equilibrium profit function in k_t, we see that this model is isomorphic to the baseline BrS model with $a = \alpha z$, and where output (y_t) is replaced with profit (π_t). Doing the same with the household production function yields the same result, with $\underline{a} = \alpha\underline{z}$. Thus we are able to reinterpret BrS' model as an endogenous growth model based on Romer (1986) under certain parameter restrictions.[17]

[17] The main restriction is setting the exponent on $\hat{K}_t$ equal to the labour share. As discussed in Ennis & Keister (2003) this restriction yields linear production, which means that aggregate capital does not have to be considered a state variable and removes transitional dynamics from the capital stock. Additionally, my assumption that labour is supplied inelastically to each class of agents is a simplification. The assumption removes interactions between the two groups through the wage, and can be removed at the expense of making the reduced form productivities effectively dependent on the aggregate state.

1.B Derivations

1.B.1 Household derivations

Risky debt

I derive the required interest rate in discrete time and take the limit. One unit is lent today, and next period $1 + r_t dt$ is repaid unless there is default. If there is default the expert's assets are seized and split amongst the lenders. The expert will have assets worth $(1-\chi)q_{t+dt}k_{t+dt}$ where χ is destroyed. The expert borrowed $d_{t+dt} = q_t k_{t+dt} - n_t$, so the assets which can be seized per unit lent is

$$(1-\chi)\frac{q_{t+dt}k_{t+dt}}{q_t k_{t+dt} - n_t} = (1-\chi)\frac{\frac{q_{t+dt}}{q_t}\phi_t}{\phi_t - 1}$$

Where $\phi_t \equiv q_t k_{t+dt}/n_t$. Suppose we are in a region where crises are possible. Then the expert defaults if the bad sunspot is drawn. A good sunspot is drawn with probability $P_g = e^{-\rho_{e,t}dt} \simeq 1 - \rho_{e,t}dt$, and a bad with probability $P_b = 1 - e^{-\rho_{e,t}dt} \simeq \rho_{e,t}dt$. The household discounts the future between t and $t+dt$ with factor $\beta = e^{-\rho_h dt} \simeq 1 - \rho_h dt$. The expected return on risky debt must equal:

$$1 = \beta P_g(1 + r_t dt) + \beta P_b \mathrm{E}_t \left[(1-\chi)\frac{\frac{\underline{q}_{t+dt}}{q_t}\phi_t}{\phi_t - 1} \middle| f_{t+dt} = 1 \right] \tag{1.45}$$

Where the expectation term is the expectation conditional on there being a crisis at $t + dt$. This expectation is for the different values of q_{t+dt} we might have depending on the value of the other shocks to the economy, and I denote the price by $\underline{q}_{t+dt}$ to make it clear that this is the crisis price. Taking the limit as $dt \to 0$, using the approximations above and noting that $dt^2 = 0$:

$$1 = 1 - \rho_h dt - \rho_{e,t}dt + r_t dt + \rho_{e,t}dt(1-\chi)\frac{\underline{q}_t}{q_t}\frac{\phi_t}{\phi_t - 1} \tag{1.46}$$

Where I have also used that $\underline{q}_{t+dt} = \underline{q}_t + d\underline{q}_t$ and $d\underline{q}_t$ is of order dt. Rearranging and dividing by dt:

$$r_t = \rho_h + \rho_{e,t}\left(1 - (1-\chi)\frac{\underline{q}_t}{q_t}\frac{\phi_t}{\phi_t - 1}\right) \tag{1.47}$$

Note in the special case of full destruction, $\chi = 1$, we have simply $r_t = \rho_h + \rho_{e,t}$. Also runs are only possible if:

$$\bar{N} = N_t - q_t k_{t+dt} + \underline{q}_t k_{t+dt} < 0 \Rightarrow \phi_t > \frac{1}{1 - \frac{\underline{q}_t}{q_t}} \tag{1.48}$$

and in this region $r_t > \rho_h$, i.e. the expert pays a premium for default risk. As leverage increases the interest rate increases, reaching a maximum of $\rho_h + \rho_{e,t}\left(1 - (1-\chi)\frac{\underline{q}_t}{q_t}\right)$.

1.B.2 Expert derivations

Conjecture that marginal value follows:

$$d\theta_t = \mu_t^\theta \theta_t dt + \sigma_t^\theta \theta_t dZ_t + df_t\left(\underline{\theta}_t - \theta_t\right) \tag{1.49}$$

Where $\underline{\theta}_t$ is an expert's marginal value following a crash, which is to be determined. Experts' value can be expressed as:

$$\rho_b \theta_t n_t = \max_{dC_t \geq 0, \phi_t \geq 0} \left\{dC_t + \mathrm{E}_t d(\theta_t n_t)\right\} \tag{1.50}$$

Using Ito's lemma the last term becomes:

$$\mathrm{E}_t d(\theta_t n_t) = \mathrm{E}_t\left[d\theta_t n_t + \theta_t dn_t + dn_t d\theta_t + df_t\left(\underline{\theta}_t \underline{n}_t - \theta_t n_t\right)\right] \tag{1.51}$$

Giving:

$$(\rho_b + \rho_{e,t})\theta_t n_t = \max_{dC_t \geq 0, \phi_t \geq 0} \left\{dC_t + \mathrm{E}_t\left[d\theta_t n_t + \theta_t dn_t + dn_t d\theta_t\right] + \rho_{e,t}\underline{\theta}_t \underline{n}_t\right\} \tag{1.52}$$

Where it is understood that the jump terms are excluded from $d\theta_t$ and dn_t in the above equation. dn_t follows:

$$dn_t = \left(dr_t^k \phi_t + (1-\phi_t)\, r_t - dc_t\right) n_t \tag{1.53}$$

Remember that there is implicitly a jump here contained in dr_t^k. Also define $dc_t = dC_t/n_t$. Excluding the jump term from dr_t^k and using (1.4) and (1.11) we can write dn_t as

$$\frac{dn_t}{n_t} = \left(\left(\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q\right) dt + \left(\sigma + \sigma_t^q\right) dZ_t\right) \phi_t + ...$$

$$... + \left(\rho_h + \rho_{e,t}\right)\left(1-\phi_t\right) dt + \rho_{e,t}(1-\chi)\frac{\underline{q}_t}{q_t}\phi_t dt - dc_t \tag{1.54}$$

Note here that the interest rate r_t is calculated assuming that the expert has taken on enough leverage to go bankrupt during a crisis. If the expert takes on low enough leverage she can survive a crisis, in which case she only pays interest $r_t = \rho_h$ and the equation above is the same just setting $\rho_{e,t} = 0$. Let's calculate expectations of the moments conditional on $df_t = 0$:

$$\mathrm{E}_t dn_t = \left[\left(\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q\right)\phi_t + \right.$$

$$\left. \left(\rho_h + \rho_{e,t}\right)\left(1-\phi_t\right) + \rho_{e,t}(1-\chi)\frac{\underline{q}_t}{q_t}\phi_t - \frac{dc_t}{dt}\right] n_t dt \tag{1.55}$$

$$\mathrm{E}_t d\theta_t = \mu_t^\theta \theta_t dt \tag{1.56}$$

$$\mathrm{E}_t dn_t^2 = \left(\sigma + \sigma_t^q\right)^2 \phi_t^2 n_t^2 dt \tag{1.57}$$

$$\mathrm{E}_t dn d\theta = \sigma_t^\theta \left(\sigma + \sigma_t^q\right) \phi_t n_t \theta_t dt \tag{1.58}$$

$$\mathrm{E}_t d\theta_t^2 = \left(\sigma_t^\theta\right)^2 \theta_t^2 dt \tag{1.59}$$

Plugging these in, and dividing by n_t:

$$
\begin{aligned}
(\rho_b + \rho_{e,t})\theta_t = \max_{dc_t \geq 0, \phi_t \geq 0} dc_t + ... \\
\mu_t^\theta \theta_t dt + ... \\
\left[\left(\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q \right) \phi_t + (\rho_h + \rho_{e,t})(1 - \phi_t) + \rho_{e,t}(1 - \chi)\frac{\underline{q}_t}{q_t}\phi_t - \frac{dc_t}{dt} \right] \theta_t dt + ... \\
\sigma_t^\theta (\sigma + \sigma_t^q) \phi_t \theta_t dt + ... \\
\rho_{e,t}\underline{\theta}_t \frac{\underline{n}_t}{n_t}
\end{aligned}
\tag{1.60}
$$

Let's assume for now that $\frac{\underline{n}_t}{n_t} = 0$, i.e. that the expert allows herself to go bankrupt during a crisis. In general, it actually depends on ϕ_t. I discuss this in the next section, where we consider the conditions under which experts will allow themselves to go bankrupt during a crisis. The leverage first order condition gives:

$$
\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - (\rho_h + \rho_{e,t}) + \rho_{e,t}(1 - \chi)\frac{\underline{q}_t}{q_t} = -\sigma_t^\theta (\sigma + \sigma_t^q) \tag{1.61}
$$

And the first order condition for consumption gives:

$$
\theta_t \geq 1 \tag{1.62}
$$

with equality if $dc_t > 0$. At the optimum, and when $dc_t = 0$, evaluating the value function gives:

$$
\rho_b \theta_t = \mu_t^\theta \theta_t dt + \rho_h \theta_t dt + \rho_{e,t}\underline{\theta}_t \frac{\underline{n}_t}{n_t} \tag{1.63}
$$

And since $\underline{n}_t = 0$ this becomes:

$$
\mu_t^\theta = \rho_b - \rho_h \tag{1.64}
$$

1.B.3 Will an expert optimally allow herself the risk of going bankrupt?

In this section I discuss under what conditions an expert will allow herself to take on enough leverage such that she would go bankrupt during a crisis. Indeed, this is necessary for anticipated crises to be possible in equilibrium. To do this we need to delve a bit deeper into the crisis value term $\underline{\theta}_t \frac{\underline{n}_t}{n_t}$. There is a kink here since the expert has limited liability. If the expert goes bankrupt this value must drop to zero because the expert's total net worth is wiped out, but if the expert chooses low enough leverage she will survive, and this value will be positive:

$$\underline{\theta}_t \frac{\underline{n}_t}{n_t} = \begin{cases} \underline{\theta}_t \left(1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right)\right) & : 1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right) \geq 0 \\ \underline{\theta}_t \cdot 0 = 0 & : 1 - \phi_t \left(1 - \frac{\underline{q}}{q_t}\right) < 0 \end{cases} \tag{1.65}$$

If she takes on low enough leverage she also only pays the risk free rate on her borrowing, so her net worth evolves according to:

$$\frac{dn_t}{n_t} = \left(\left(\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma \sigma_t^q \right) dt + \left(\sigma + \sigma_t^q\right) dZ_t \right) \phi_t + ...$$

$$... + \rho_h \left(1 - \phi_t\right) dt - dc_t \tag{1.66}$$

There is a clear cost-benefit decision here: you can take on high leverage, in which case you make large profit as long as there is no crisis, but large losses during a crisis, or you can take on low leverage and make low profit in normal times and less losses in a crisis. A key variable is how much a unit of net worth is worth to an expert during a crisis: $\underline{\theta}_t$. If a unit of net worth is worth a lot during a crisis (as we might expect, since returns are high) then this will push experts towards caution. If the expert is happy with a level of leverage in the region where she doesn't go bankrupt during a crisis then we can show that the following FOC holds:

$$\frac{a-\iota_t}{q_t}+\Phi(\iota_t)-\delta+\mu_t^q+\sigma\sigma_t^q-\rho_h-\rho_{e,t}\left(1-\frac{\underline{q}_t}{q_t}\right)\underline{\theta}_t+\sigma_t^\theta\left(\sigma+\sigma_t^q\right)=0 \qquad (1.67)$$

Compare this to the leverage FOC in the region where she does go bankrupt:

$$\frac{a-\iota_t}{q_t}+\Phi(\iota_t)-\delta+\mu_t^q+\sigma\sigma_t^q-\rho_h-\rho_{e,t}\left(1-\frac{\underline{q}_t}{q_t}\right)-\rho_{e,t}\chi\frac{\underline{q}_t}{q_t}+\sigma_t^\theta\left(\sigma+\sigma_t^q\right)=0 \quad (1.68)$$

This shows us the differences between the costs of increasing leverage on the margin in the two regions. If the expert has low enough leverage to survive a crisis (first equation), increasing leverage hurts because it increases the losses in a crisis, which are valued at the marginal value of net worth during a crisis. If the expert has already chosen high enough leverage to go bankrupt, then increasing leverage by an extra unit hurts in a different way: it increases borrowing costs.

Proposition 5. *If* $(\underline{\theta}_t-1)\left(1-\frac{\underline{q}_t}{q_t}\right)\leq\chi\frac{\underline{q}}{q_t}$ *then experts find it optimal to choose high enough leverage to go bankrupt during a crisis.*

Proof. Suppose that a crisis is possible in equilibrium. Then in equilibrium we know that (1.68) holds because this is the required optimality condition for the other experts to be willing to accept a crisis. We need to verify that an individual expert is willing to choose leverage that leads her to go bankrupt during a crisis. To verify this it is sufficient to check that lowering leverage does not increase today's value. Small changes in leverage lead to no change in value, because remember that experts are locally indifferent about their leverage choices in equilibrium. But what about a change large enough to avoid bankruptcy? This is not profitable as long as (1.67) holds either with equality, or instead has the terms on the left hand side greater than zero. This means that in this region increasing leverage weakly increases value, and given the linearity this means the expert increases leverage all the way in to the region where she does go bankrupt. Mathematically, we require:

$$\frac{a-\iota_t}{q_t}+\Phi(\iota_t)-\delta+\mu_t^q+\sigma\sigma_t^q-\rho_h-\rho_{e,t}\left(1-\frac{\underline{q}_t}{q_t}\right)\underline{\theta}_t+\sigma_t^\theta\left(\sigma+\sigma_t^q\right)\geq 0 \tag{1.69}$$

Substituting in (1.68) and cancelling terms leaves:

$$(\underline{\theta}_t-1)\left(1-\frac{\underline{q}_t}{q_t}\right)\leq\chi\frac{\underline{q}}{q_t} \tag{1.70}$$

□

The intuition behind the result is quite simple. Recall that $\underline{\theta}_t$ is the marginal value of net worth during a crisis. (1.70) says that this value cannot be too high, otherwise experts would want to deleverage (all the way to zero leverage, in fact) to take advantage of this. This benefit of reducing leverage must be weighed against the cost, which is the lost expected revenue from lending. In equilibrium this can be derived from the optimality of the other experts, giving the above expression. In the baseline model, with $\chi=0$, this condition requires that $\underline{\theta}_t=1$. This is a very strong requirement, which follows from the way the model is constructed. In particular, recall that since experts don't face leverage constraints in equilibrium they must be indifferent about their leverage choices. Given their risk neutrality, this means that they derive no utility from any of their lending, which is why $\underline{\theta}_t$ has to be so low in order to convince them to lend even in the face of a potential crisis. This would not be true in a richer model, in which experts derived more explicit benefits from lending, where $\underline{\theta}_t$ would be allowed to be higher.

The capital externality which creates endogenous growth actually also ensures that $\underline{\theta}_t=1$ in the baseline model, making crises possible in equilibrium. This is summarised in the following proposition:

Proposition 6. *In the basline model with endogenous growth (as described in Appendix 1.A) and with permanent financial crises ($\rho_r=0$), $\underline{\theta}_t=1$.*

Proof. From an individual expert's point of view all of the other experts go bankrupt in a crisis, and therefore are unable to hold any capital. Since an individual expert has zero mass this means that the total capital held by experts, $\hat{K}_t$ is zero. Due to the production externality, this means that the productivity of a surviving expert is zero in a crisis: $y_t = z\hat{K}_t^{1-\alpha} k_t^{\alpha} l_t^{1-\alpha} = 0$. Since the expert is unable to produce, now or for the rest of time, she might as well consume her net worth, leading to $\underline{\theta}_t = 1$. $\square$

The idea behind this admittedly stylised assumption is that the disruption in financial markets during a crisis would make it hard for a surviving bank to function efficiently. It also matches the empirical fact that the value of being a surviving bank (as measured by the market value of bank equity) appears to be very low: In the US bank equity values fell by an average of 80% during the crisis and have remained persistently low. It is also worth noting that this result does not hold exactly for non-permanent crises, because then experts' net worth can have higher value, even if they cannot produce today. This is because they might want to hold on to their wealth in order to benefit from positive returns when the other experts are bailed out, and $\hat{K}_t$ becomes positive again. It is easy to show that, for any $\hat{\eta}$, $\underline{\theta}_t$ is falling in the expected length of the crisis, so there is always a long enough crisis (small enough ρ_r) to ensure that (1.70) holds.

To see the importance of the aggregate-capital externality in allowing crises, it is instructive to think about the case where the reduced form production functions ($y_t = ak_t$ and $\underline{y}_t = \underline{a}\underline{k}_t$) are the true production functions, and there is thus no feedback from aggregate expert capital to individual expert productivity:

Proposition 7. *Consider a model without the aggregate-capital externality, meaning that $y_t = ak_t$ and $\underline{y}_t = \underline{a}\underline{k}_t$ are the true individual expert and household production functions respectively. Then $\underline{\theta}_t = \infty$.*

Proof. To see this, note that while the economy is in the crisis capital is priced by the household, at $\underline{q}$, which is constant (until the exogenous recapitalisation restores the economy to positive η). Since $\underline{q}$ and $\underline{\theta}$ are constant there is also effectively no risk to the surviving expert from investing (the expert only cares about covariance risk of net worth with θ_t). There also cannot be another financial crisis while we are in a crisis, by construction, so the expert can borrow risk free at $r_t = \rho_h$. Since the household is investing, we know that $\mathrm{E}_t d\underline{r}_t^k = \rho_h$, and using the definitions of $d\underline{r}_t^k$ and dr_t^k we can show that

$$\mathrm{E}_t dr_t^k = \rho_h + \left(\frac{a - \underline{a}}{\underline{q}} + \underline{\delta} - \delta \right) dt > r_t = \rho_h$$

I.e. the expected profit from increasing leverage is positive ($\mathrm{E}_t dr_t^k - r_t > 0$). Without any risk there is no force that creates a cost of leverage to experts, and a suviving expert would thus choose infinite leverage, making infinite instantaneous profit, leading to $\underline{\theta}_t = \infty$. □

This proposition highlights the fundamental issue making it hard to generate expected financial crises in this model: without any other frictions, it is great to be the only surviving expert in a financial crisis. With asset prices so low you can make huge amounts of profit. The aggregate-capital externality powering endogenous growth is a way to shut this down, by making it bad to be the only surviving expert. Other more realistic assumptions could replace this, but the general idea is that disruptions in financial markets during a crisis should reduce the value of being a surviving expert. For example, imposing a borrowing constraint during crises would reduce the ability of experts to take advantage of the temporarily high returns. Another possibility is that the expectation of bailouts is what leads experts to allow themselves to get in trouble during a crisis. Indeed, this is the focus of several theoretical papers, for example Farhi & Tirole (2012), Acharya & Yorulmazer (2007), and Mailath & Mester (1994). Indeed, this concern is empirically validated, as shown by Duchin & Sosyura (2014) who use the TARP program to show that individual banks increase the riskiness of their portfolios in response to signals that they might receive government aid in the future. Future work could incorporate this mechanism as a potential rationalisation of crises in my model.

1.B.4 Equilibrium derivations

Derivation of μ_t^q, σ_t^q, μ_t^θ, and σ_t^θ

Using Ito's lemma on $q_t = q(\eta_t)$:

$$dq_t = q'(\eta_t)d\eta_t + \frac{1}{2}q''(\eta_t)d\eta_t^2 + \left(\underline{q} - q_t\right) df_t \tag{1.71}$$

Using the conjectured law of motion for η_t, (1.19), gives:

$$dq_t = \left(q'(\eta_t)\mu_t^\eta\eta_t + \frac{1}{2}q''(\eta_t)\left(\sigma_t^\eta\right)^2\eta_t^2\right)dt + q'(\eta_t)\sigma_t^\eta\eta_t dZ_t + \left(\underline{q} - q_t\right) df_t \tag{1.72}$$

Equation coefficients from the above equation with the conjectured law of motion for q_t, (1.3), gives:

$$\mu_t^q = \frac{q'(\eta_t)\mu_t^\eta\eta_t + \frac{1}{2}q''(\eta_t)\left(\sigma_t^\eta\right)^2\eta_t^2}{q_t} \qquad \sigma_t^q = \frac{q'(\eta_t)\sigma_t^\eta\eta_t}{q_t}$$

We can do the same exercise for θ_t to calculate μ_t^θ and σ_t^θ as:

$$\mu_t^\theta = \frac{\theta'(\eta_t)\mu_t^\eta\eta_t + \frac{1}{2}\theta''(\eta_t)\left(\sigma_t^\eta\right)^2\eta_t^2}{\theta_t} \qquad \sigma_t^\theta = \frac{\theta'(\eta_t)\sigma_t^\eta\eta_t}{\theta_t}$$

Derivation of μ_t^η and σ_t^η

Now I need to use Ito's lemma multiple times to work out the evolution of η_t using the definition $\eta_t \equiv N_t/(q_tK_t)$. I first need the individual evolutions of N_t, q_t and K_t. q_t is already given by (1.3). K_t is total capital ($K_t \equiv K_t^b + K_t^h$) which it is easy to show evolves via:

$$\frac{dK_t}{K_t} = (\Phi(\iota_t) - \delta\psi_t - \underline{\delta}(1-\psi_t))\,dt + \sigma dZ_t \tag{1.73}$$

This comes from aggregating (1.1) and (1.2). The evolution of N_t, total bank capital, is just the aggregate version of (1.53):

$$\frac{dN_t}{N_t} = dr_t^k\phi_t + (\rho_h + \rho_{e,t})\,(1-\phi_t)\,dt + \rho_{e,t}(1-\chi)\frac{\underline{q}_t}{q_t}\phi_t dt - dc_t \tag{1.74}$$

We need to use Ito's Lemma including jumps to deal with the jump df_t in the net worth and price evolution. Remember that if $df_t = 1$, N_t jumps to zero, and q_t jumps to $\underline{q}$. Ito's lemma gives:

$$d\eta_t = \frac{dN_t}{q_tK_t} - \frac{d(q_tK_t)N_t}{(q_tK_t)^2} - \frac{dN_t d(q_tK_t)}{(q_tK_t)^2} + \frac{N_t d(q_tK_t)^2}{(q_tK_t)^3} + \left(\frac{0}{\underline{q}K_t} - \frac{N_t}{q_tK_t}\right) df_t \tag{1.75}$$

Where it is understood that the jumps have been removed from dN_t and dq_t. Rearranging gives:

$$\frac{d\eta_t}{\eta_t} = \frac{dN_t}{N_t} - \frac{d(q_tK_t)}{q_tK_t} - \frac{dN_t}{N_t}\frac{d(q_tK_t)}{q_tK_t} + \frac{d(q_tK_t)^2}{(q_tK_t)^2} - df_t \tag{1.76}$$

Using Ito's Lemma on $d(q_tK_t)$ and $d(q_tK_t)^2$gives:

$$\frac{d(q_tK_t)}{q_tK_t} = \frac{dq_t}{q_t} + \frac{dK_t}{K_t} + \frac{dq_t}{q_t}\frac{dK_t}{K_t} \tag{1.77}$$

$$\frac{d(q_tK_t)^2}{(q_tK_t)^2} = \frac{dq_t^2}{q_t^2} + \frac{dK_t^2}{K_t^2} + 2\frac{dq_t}{q_t}\frac{dK_t}{K_t} \tag{1.78}$$

Plugging this and the assumed dq_t equation into the $d(q_tK_t)$ terms gives:

$$\frac{d(q_tK_t)}{q_tK_t} = (\Phi(\iota_t) - \delta\psi_t - \underline{\delta}(1-\psi_t) + \mu_t^q + \sigma\sigma_t^q)\,dt + (\sigma + \sigma_t^q)\,dZ_t \tag{1.79}$$

$$\frac{d(q_tK_t)^2}{(q_tK_t)^2} = (\sigma_t^q)^2dt + \sigma^2dt + 2\sigma\sigma_t^qdt = (\sigma + \sigma_t^q)^2dt \tag{1.80}$$

Now using definition of dr_t^k to simplify a bit:

$$\frac{d(q_tK_t)}{q_tK_t} = dr_t^k - \frac{a-\iota_t}{q_t}dt - (1-\psi_t)(\underline{\delta} - \delta)dt \tag{1.81}$$

Calculating the cross term:

$$\frac{dN_t}{N_t}\frac{d(q_tK_t)}{q_tK_t} = \phi_t(\sigma + \sigma_t^q)^2dt \tag{1.82}$$

Putting all this together:

$$\begin{aligned}\frac{d\eta_t}{\eta_t} &= (\phi_t - 1)(dr_t^k - \rho_h dt - \rho_{e,t}dt) + \rho_{e,t}(1-\chi)\frac{\underline{q}_t}{q_t}\phi_t dt + \frac{a-\iota_t}{q_t}dt + (1-\psi_t)(\underline{\delta} - \delta)dt \\ &\quad - (\phi_t - 1)(\sigma + \sigma_t^q)^2dt - dc_t - df_t \end{aligned}\tag{1.83}$$

Which is exactly BrS' equation for the evolution of η_t, plus the extra $\rho_{e,t}$ terms. Equating terms with the guessed form for $d\eta_t$ in (1.19):

$$\sigma_t^\eta = (\phi_t - 1)(\sigma + \sigma_t^q) \tag{1.84}$$

$$\mu_t^\eta = (\phi_t - 1)(Edr_t^k - \rho_h - \rho_{e,t} - (\sigma + \sigma_t^q)^2) + \rho_{e,t}(1-\chi)\frac{\underline{q}_t}{q_t}\phi_t + \frac{a-\iota_t}{q_t} + (1-\psi_t)(\underline{\delta} - \delta) \tag{1.85}$$

Now if banks are holding positive leverage we can use their leverage FOC to simplify μ_t^η to:

$$\mu_t^\eta = -\sigma_t^\eta(\sigma + \sigma_t^q + \sigma_t^\theta) + \rho_{e,t}(1-\chi)\frac{\underline{q}}{q_t} + \frac{a-\iota_t}{q_t} + (1-\psi_t)(\underline{\delta} - \delta) \tag{1.86}$$

1.B.5 Boundary conditions

The boundary conditions are identical to those in the baseline model of BrS' paper, and the interested reader is referred to their Proof of Proposition II.4.

1.C Proofs

Proof of Proposition 1. Firstly note that $\frac{\partial(\sigma^{\eta}\eta)}{\partial\phi} > 0$, and thus $\frac{\partial\sigma^q}{\partial\phi} > 0$ and $\frac{\partial\sigma^\theta}{\partial\phi} < 0$ since $q' > 0$ and $\theta' < 0$ respectively. Now implicitly differentiate (1.26) with respect to ρ_e:

$$\phi'(\rho_e)\left[(\sigma^\theta)'(\sigma+\sigma^q)+(\sigma^q)'\sigma^\theta-\left(1-\frac{\underline{q}}{q}\right)\rho_e\chi\frac{\underline{q}}{q}\right]=\chi\frac{\underline{q}}{q}\left[1-\phi\left(1-\frac{\underline{q}}{q}\right)\right] \quad (1.87)$$

The term in square brackets on the right hand side is negative whenever crises are possible. The term in square brackets on the left hand side is negative because σ^θ and $(\sigma^\theta)'$ are negative, and $1-\underline{q}/q$ is positive whenever $\underline{q} < q$. □

Proof of Proposition 4. Since the household provides the insurance and is risk neutral, the premium for the insurance contract must satisfy $r_t^I = \rho_{e,t}$ so that the household breaks even.[18] Note that if the expert has an optimal plan that involves her going bankrupt during a crisis, then limited liability implies that she will never hold any of the insurance asset – it costs her in normal times and provides no benefits during a crisis. Thus the only way she might hold any is if it gives benefit in a plan where she will remain solvent during a crisis. In this region, and including the insurance asset, the expert's optimisation problem is now:

$$(\rho_b+\rho_{e,t})\theta_t = \max_{dc_t\geq 0,\phi_t\geq 0,\phi_t^I\geq 0} dc_t + ...$$

$$\mu_t^\theta\theta_t dt + ...$$

$$\left[\left(\frac{a-\iota_t}{q_t}+\Phi(\iota_t)-\delta+\mu_t^q+\sigma\sigma_t^q\right)\phi_t+\rho_h\left(1-\phi_t\right)-r_t^I\phi_t^I-\frac{dc_t}{dt}\right]\theta_t dt + ...$$

$$\sigma_t^\theta\left(\sigma+\sigma_t^q\right)\phi_t\theta_t dt + ...$$

[18] To see this, note that from t to $t+dt$ the household earns the expected insurance premium $(1-\rho_{e,t}dt)r_t^I dt = r_t^I dt$, and has expected payout $\rho_{e,t}dt$. Setting expected profit to zero gives $r_t^I dt - \rho_{e,t}dt = 0 \Rightarrow r_t^I = \rho_{e,t}$.

$$\rho_{e,t}\underline{\theta}_t\left(1-\phi_t\left(1-\frac{q}{q_t}\right)+\phi_t^I\right) \tag{1.88}$$

Where k_t^I is the units of the insurance asset held, and $\phi_t^I = k_t^I/n_t$. Imposing that $r_t^I = \rho_{e,t}$, we see that the expert will optimally choose to hold zero insurance ($\phi_t^I = 0$) as long as $\underline{\theta}_t \leq \theta_t$, which is always satisfied if $\underline{\theta}_t = 1$ which is true in the baseline model (see Proposition 6). In this case insurance also doesn't affect the expert's choice of leverage or consumption relative to the case without insurance. □

1.D Model solution with regulatory leverage constraint

This section outlines the solution of the model when experts face an exogenous borrowing constraint of the form $\phi_t \leq \bar{\phi}_t$. In the model this is interpreted as a regulatory leverage constraint, but the solution also applies to leverage constraints derived from limited commitment problems. We still conjecture that marginal value follows:

$$d\theta_t = \mu_t^\theta \theta_t dt + \sigma_t^\theta \theta_t dZ_t + df_t\left(\underline{\theta}_t - \theta_t\right) \tag{1.89}$$

Experts' value can be expressed as:

$$\rho_b \theta_t n_t = \max_{dC_t \geq 0, 0 \leq \phi_t \leq \bar{\phi}_t} \left\{dC_t + \mathrm{E}_t d(\theta_t n_t)\right\} \tag{1.90}$$

Using previous arguments we can show that:

$$(\rho_b + \rho_{e,t})\theta_t = \max_{dc_t \geq 0, 0 \leq \phi_t \leq \bar{\phi}_t} dc_t + ...$$

$$\mu_t^\theta \theta_t dt + ...$$

$$\left[\left(\frac{a-\iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q\right)\phi_t + (\rho_h + \rho_{e,t})(1-\phi_t) + \rho_{e,t}(1-\chi)\frac{q_t}{q_t}\phi_t - \frac{dc_t}{dt}\right]\theta_t dt + ...$$

$$\sigma_t^\theta \left(\sigma + \sigma_t^q\right) \phi_t \theta_t dt$$

Where I have imposed that $\frac{n_t}{n_t} = 0$, i.e. that the expert allows herself to go bankrupt during a crisis. The leverage first order condition gives:

$$\frac{a - \iota_t}{q_t} + \Phi(\iota_t) - \delta + \mu_t^q + \sigma\sigma_t^q - (\rho_h + \rho_{e,t}) + \rho_{e,t}(1-\chi)\frac{\underline{q}_t}{q_t} + \sigma_t^\theta \left(\sigma + \sigma_t^q\right) = \lambda_t \geq 0 \quad (1.91)$$

Where λ_t is the lagrange multiplier on the leverage constraint. As before, the first order condition for consumption gives:

$$\theta_t \geq 1 \quad (1.92)$$

with equality if $dc_t > 0$. At the optimum, and when $dc_t = 0$, evaluating the value function now gives:

$$\mu_t^\theta = \rho_b - \rho_h - \lambda_t \bar{\phi}_t \quad (1.93)$$

The rest of the model is the same, except we must be careful in the derivation of μ_t^η to use the new leverage first order condition.

1.E Numerical solution and simulation

The algorithm to solve the model is based on BrS' original algorithm, which is detailed after their statement of Proposition II.4. The only difference is that I have additional terms in some of my equations relating to the crisis price $\underline{q}$. In the baseline parameterisation with permanent crises this can be solved for at the beginning of the code, and passed as a parameter to the rest of the algorithm. Following BrS, the algorithm searches over values of $q'(0)$ to find the value which satisfies the required boundary conditions. Another difference from BrS is that I use a Newton-based algorithm to update my guesses for $q'(0)$ until I reach convergence, whereas they use a bisection algorithm.

The model is simulated by discretising time. For example, I simulate η_t using its transition:

$$d\eta_t = \mu_t^\eta \eta_t dt + \sigma_t^\eta \eta_t dZ_t - \eta_t df_t$$

I choose a small value for dt, and draw values of dZ_t from a normal distribution with mean zero and standard deviation $\sqrt{dt}$. Given η_t, the value at the next interval of time is found by the approximation $\eta_{t+dt} \simeq \eta_t + d\eta_t$. The jump process is approximated by a random variable which takes value one in every period with probability $1 - e^{-\rho_{e,t} dt}$ and zero otherwise. The model moments (average growth and standard deviations) are calculated from simulations of 5000 years, with $dt = 1/120$ (the results are unchanged by picking smaller values of dt). The expected time to experience a crisis is calculated by repeatedly simulating the economy, starting from η^*, and calculating how long it takes for the economy to experience a crisis, and averaging this over the trials.

1.F Model solution with $\chi > 0$

If I choose $\chi > 0$ then we can no longer use a constant value for the probability of coordinating on a crisis, as I did when $\chi = 0$. This is because this leads to discontinuous changes in leverage at the point where crises become possible. To deal with this, in this case I instead parameterise $\rho_{e,t}$ as:

$$\rho_{e,t} = \begin{cases} \rho_e \left|\frac{\underline{N}_t}{N_t}\right| & : \underline{N}_t < 0 \\ 0 & : \underline{N}_t \geq 0 \end{cases} \tag{1.94}$$

This parameterisation has appealing economic features, as well as being mathematically useful. Economically, it says that agents are more likely to coordinate on a crisis equilibrium the less well capitalised the banking sector is. Mathematically, this removes a discontinuity from the model around the point where a crisis just becomes possible. Conditional on this functional form, I have one free variable to choose which is the slope term ρ_e. The larger this parameter is the more likely we are to coordinate on the crisis equilibrium. Figure 1.9 plots the solution to the model in this case. The results are qualitatively similar to the results of the baseline model, except for leverage: leverage is now lower in some regions due to the experts' desire to deleverage to avoid paying the exogenous default costs.

Figure 1.9: Model solution with $\chi = 0.25$. Key variables

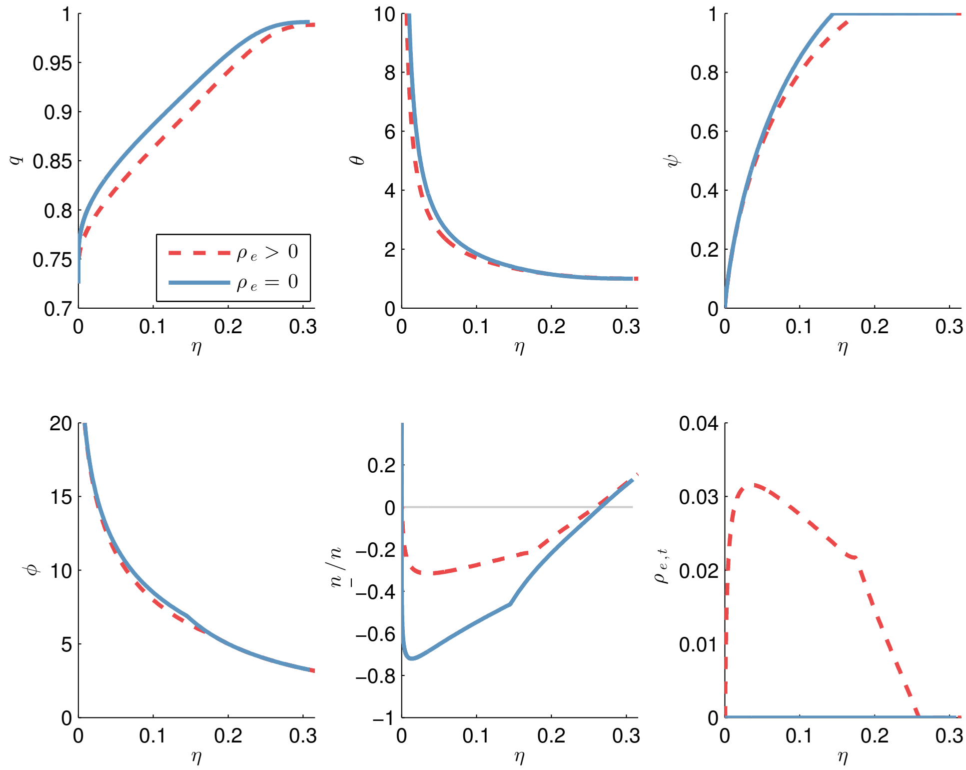

Model solution with $\chi = 0.25$. *Solid blue line plots the solution without crises (*$\rho_e = 0$*) and dashed red line plots the solution with crises (*$\rho_e = 0.1$*).*

1.G Model solution with non-permanent crises

In this section I solve the model with non permanent crisis. I keep the same parameters as the baseline model, including setting $\rho_e = 0.1$. I set the level of recapitalisation of the experts to $\hat{\eta} = 0.00005\eta^*$, and the flow intensity of recapitalisation to $\rho_r = 1.2427$. With these numbers, the chance of being recapitalised within one year is 71%, two years is 92%, and essentially 100% within around five years. Recapitalisation to that value of $\hat{\eta}$ implies that it takes roughly 15 years for the economy to naturally recover from $\hat{\eta}$ back to η^*, giving a total time from crisis to complete recovery of something around 20 years.

This generates a value of the crisis price of $\underline{q} = 0.783$, which is higher than the crisis price when crises are permanent (0.725). Since the crisis price is now slightly higher, crises are less likely, and the economies with and without crises spend 5% and 2.1% of their time in the crisis region under the stationary density (ignoring crisis realisations). Figure 1.10 plots key variables from the model solution. By comparison with Figure 1.3 it can be seen that the model solutions are qualitatively very similar.

Figure 1.10: Model solution non permanent crises

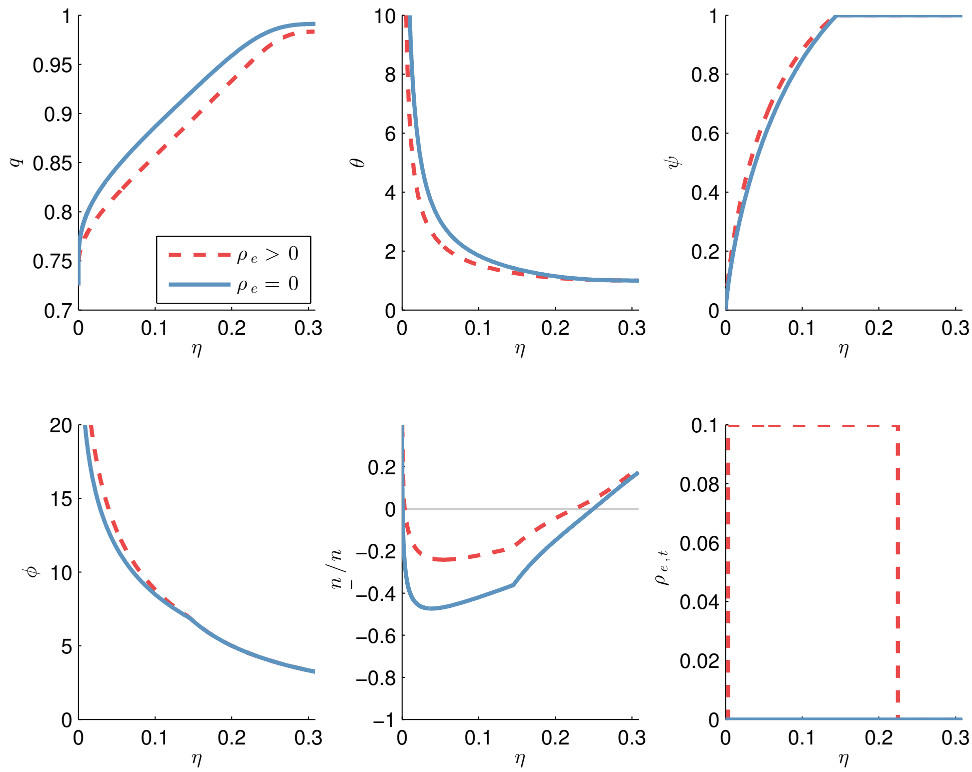

The dashed red line gives the solution to the model with the baseline positive value of ρ_e, meaning the economy occasionally experiences crises. The solid blue line gives the solution where $\rho_e = 0$ and agents never coordinate on a crisis.

1.H Miscellaneous Graphs and figures

Figure 1.11: Parameter sensitivity: crisis net worth, q and ϕ across changes in three parameters

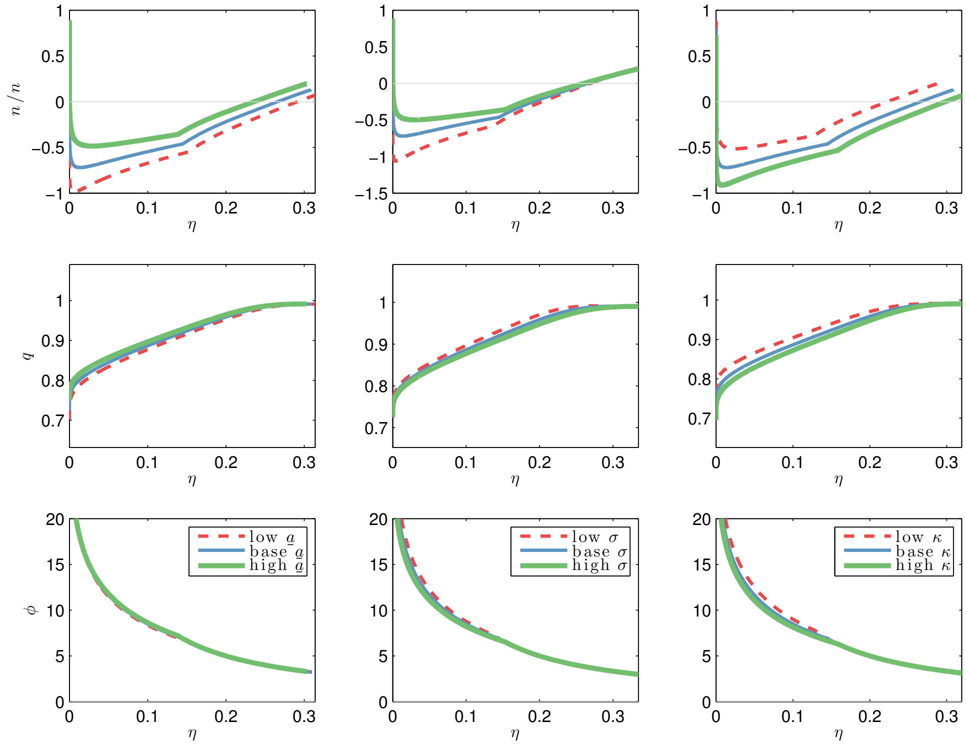

The top row plots the crisis region across parameter changes for three parameters. The middle and lower rows decompose this into changes in asset prices and leverage.

Chapter 2

Real Wages and the Manifestation of Financial Crises

2.1 Introduction

The recent financial crisis caused a highly synchronised recession across much of the developed world. However, beneath the surface there are differences in how countries experienced this decline. In this paper I document a new stylised fact: countries which experienced larger declines in Total Factor Productivity (TFP) during the crisis experienced less severe falls in hours worked. In other words, some countries experienced the recession mostly as a collapse in employment, and others mostly as a collapse in measured productivity. How can we rationalise this heterogeneous behaviour in response to a common global shock?

I present differential wage adjustment in response to the crisis as a potential explanation for this fact. I show that countries with larger falls in real wages during the crisis tend to be those with TFP, and not labour market, problems. Motivated by this second fact, I offer a parsimonious explanation of the negative TFP-hours correlation using a model of firm heterogeneity and differential wage adjustment in response to a financial shock.

The intuition is relatively simple. For a given level of wages, a financial crisis reduces the ability of firms to borrow to fund investment, reducing the capital stock, marginal product of labour, and hence demand for labour. If wages adjust relatively little, this will lead to large falls in hours worked in equilibrium. On the other hand, if wages adjust downwards by a lot this offsets the fall in demand for labour, leading to smaller falls in hours. However, lower wages also shield firms from having to shut down or downsize in response to the financial crisis, which leads to a worsening allocation of resources and hence lower measured TFP.

My model thus features endogenous TFP movements in response to a financial shock. This comes from a composition effect, based on the model of Buera & Moll (forthcoming). In particular, the model features firms who are heterogeneous in their productivities, and who decide whether or not to produce based on their profitability. Firms with the lowest productivity levels choose not to produce, leading to measured TFP being endogenously determined by the set of firms who are producing. In the model, a fall in the real wage increases profitability for all firms, leading unproductive firms to start producing and reducing measured TFP. The extent to which this happens depends on the extent of wage adjustment during the crisis, with larger falls in wages leading to larger falls in measured TFP. Differential wage adjustment thus generates a negative TFP-hours correlation in my model following a financial crisis, consistent with the data.

I present cross-country correlations consistent with this story, as well as a more detailed case-study look at the US and UK. For the cross-section analysis, I use data from the OECD over the crisis period. I first present four simple correlations over the whole period: a) TFP and hours are negatively correlated, b) TFP and real wages are positively correlated, c) hours and real wages are negatively correlated, d) real wages and the (nominal) price level are negatively correlated. Of these, (a) is the main correlation to be explained, and (b) and (c) support my focus on differential wage adjustment as a potential explanation. (d) suggests an explanation for why different countries experienced differential wage adjustment, which I had up to now taken as exogenous. In an environment with downwards nominal wage rigidity, countries which run higher inflation will have their real wages reduced, which I find evidence for in the cross section. I also construct a short country-year panel using the cross-country data, which I use to control for country and time fixed effects. This shows that the correlations are not driven by preexisting country-specific factors, and gives me enough power to show that the correlations are statistically significant.

After establishing these correlations in the cross section, I focus in particular on two countries: the United States and the United Kingdom. Both of these countries have come under scrutiny in recent years because of the nature of their experiences of the Great Recession, and they lie at opposite ends of the spectrum in terms of their labour market and TFP experiences making them useful examples. The US has seen an unusually large fall in employment and hours. Between 2008 and 2011 total hours fell by roughly 10%. On the other hand, over that period TFP remained robust, and even increased slightly relative to trend.

In the UK the labour market performed relatively better, with hours only falling by just over 5% over the same period, but the TFP performance has been dismal. TFP has fallen by over 5% relative to trend, in what has been deemed the UK's "productivity puzzle". At the same time period, the real wage behaviour of the two countries has been very different, with the US seeing wages grow in line with their trend over the period, and the UK seeing wages fall by 6% relative to trend, consistent with my proposed explanation.

For these two countries I construct business cycle wedges following Chari, Kehoe & McGratten's (2007) accounting procedure. This exercise shows that the US' recession can be explained mostly through the labour wedge, and the UK's recession mostly through the efficiency wedge. My model is also consistent with this evidence: I show that following a financial shock, the model generates only a labour wedge and no efficiency wedge if wages are fully rigid, and only an efficiency wedge and no labour wedge if wages are fully flexible. The model is thus also able to jointly rationalise the behaviour of the efficiency and labour wedges in the US and UK during the crisis by again appealing to differential wage adjustment.

Focusing on these two countries also allows me to look closer at various aspects of the data. Firstly, I decompose the labour wedge into distortions on the firm and household side, following the procedure of Karabarbounis (2014), and show that the labour wedge is driven primarily by households being off of their labour supply curves, which is consistent with the mechanism in my model, which relies on sticky wages leading to households' labour supply being rationed. Secondly, I take a closer look at the role of inflation and nominal wages in determining the evolution of real wages, and discuss institutional changes in the two countries that could contribute to their patterns of wage adjustment. Finally, I discuss the robustness of my results to composition effects driven by the firing of low productivity workers during recessions.

The rest of the paper is organised as follows. I review related literature in section 2.2. In section 2.3 I present the international cross-sectional evidence. In section 2.4 I present further results for the US and UK, including the wedges decomposition. Section 2.5 contains the model and results, and in section 2.6 I conclude.

2.2 Related literature

This paper is related to many theoretical and empirical papers that relate financial crises to labour markets and productivity. One main contribution of the paper is to clarify the role of wages in the transmission of financial shocks to productivity. In some papers, such as Buera & Moll (forthcoming), a financial crisis manifests itself as a fall in productivity. In others, such as Petrosky-Nadeau (2013), a crisis manifests itself as a rise in productivity. How can different models give such different predictions for productivity, and what other data can we use to discipline what kind of response of productivity is appropriate? My model and empirics highlight the role of wage adjustment as a driver of misallocation and productivity.

On the theoretical side, the paper is related to papers such as Khan & Thomas (2013) which emphasise endogenous productivity from misallocation across heterogeneous firms in response to financial shocks. They consider a flexible wage economy, whereas I also consider sticky wages, allowing me to compare economies with differing degrees of wage adjustment. Arellano, Bai & Kehoe (2012) also consider a model with heterogeneous firms, and consider both flexible and sticky wages. Their focus is on matching labour market outcomes, and they do not compare endogenous productivity across their flexible and sticky-wage variants. Buera, Fattal-Jaef & Shin (2014) also generate a fall in TFP following a financial shock. In an extension to their model, they add sticky wages and show that unemployment increases further in their model when wages are sticky. TFP in their model appears to fall slightly less when wages are sticky, consistent with my results, although they do not discuss nor attempt to explain this. Other papers discuss the empirics and theory behind the effect of financial shocks on the composition of firms and workers. For example, Siemer (2014) finds that young, small firms suffer disproportionately more during the recent crisis, and builds a model to explain this.

I perform Chari, Kehoe, & McGratten's (2007) business cycle accounting exercise over the Great Recession period for the US and the UK. This exercise has been performed by other authors (Ohanian, 2010, and Chadha & Warren, 2013), and my results are consistent with their findings. Other papers have investigated the UK's productivity puzzle. Pessoa & Van Reenen (2014) argue that the fall in output per worker in the UK is related to the UK's large fall in wages, which encouraged substitution away from capital towards labour, but they do not discuss the potential for wages to affect TFP.

My paper also builds on several components which have theoretical and empirical grounding in other papers. Firstly, the transmission of the financial crisis to the firm sector has been documented in numerous studies. Using German data, Dwenger, Fossen, & Simmler (2015) show that firms who banked with banks hit harder by losses to proprietary trading activities reduced their investment significantly more than other firms. Focusing instead on the heterogeneity in firms' financing decisions, Giroud & Mueller (2015) show that the majority of job losses associated with falling house prices during the Great Recession are concentrated among firms with high leverage, who are feasibly more at risk during financial tightenings.

Secondly, I utilise sticky wages. Olivei & Tenreyro (2007, 2010) find indirect evidence for wage stickiness by exploiting known timing conventions for wage setting. Countries, such as the US and Japan, in which wages are known to be reset around the same time at all firms, experience larger responses to monetary policy shocks immediately following the wage reset than in other quarters. Countries without such a known convention experience no such pattern. Kaur (2014) studies nominal wage setting in informal agricultural markets in India, and finds that nominal wages tend to rise following positive rain shocks, but not fall in response to negative shocks. In a cross section of US counties, Mian & Sufi (2014) show that counties which were harder hit by the collapse of the housing bubble had no larger wage adjustment than other counties, and larger unemployment increases, suggesting a role for sticky wages. Other papers study downwards nominal wage rigidity by looking for a spike at zero nominal wage changes, and an associated missing mass below zero. Daly & Hobijn (2014) document an increase in this spike during the Great Recession in the US. Finally, Druant et al. (2009) provide survey evidence on the wage setting practices of a sample of European firms. They find that 29.7% of firms have a policy of adjusting wages for inflation, and that only half of these use automatic indexation. Additionally, most firms who use indexation index to historical inflation numbers. Overall, this suggests a sizable fraction of firms for whom inflation could have effects on real wages.

Sticky wages, both nominal and real, have also been used in many theoretical models to amplify employment fluctuations. Galí (2011) embeds staggered wage setting into the standard New Keynesian framework. Den Haan, Rendahl, & Riegler (2015) build a model where precautionary money demand rises during recessions, pushing down the price level and increasing real wages if nominal wages are sticky, leading to further increases in unemployment. Schmitt-Grohé & Uribe (2015) show that nominal wage rigidity causes overborrowing in small open economies when combined with a currency peg.

Finally, my paper is related to the literature on directed technical change. This literature emphasises the role of factor prices in determining where investment in factor-augmenting technical change is directed. For example, if real wages are low (relative to the prices of other inputs) this would lead to firms using relatively more labour for production, increasing the demand for innovations which improve the efficiency of labour. Acemoglu (2002) lays the theoretical foundation for the modern literature. Hanlon (2015) provides an empirical test of the theory using the impact of the U.S. Civil War on the British cotton textile industry. My paper is similar in spirit, in that I also emphasise the role of a factor price (the real wage), but I focus on efficiency effects via the allocation of existing resources, instead of via directed technical change.

2.3 International evidence

In this section, I present the international evidence at the core of the paper, including the negative TFP-hours correlation and the supporting evidence leading to my proposed explanation. I first present simple correlations, then partial correlations which allow me to control for country and time fixed effects, and finally perform robustness checks.

2.3.1 Data

The data are from the OECD dataset "Growth in GDP per capita, productivity and ULC". I use data on TFP, hours worked, real wages, population and prices. For some countries I use wage data from different sources to increase the sample length, and these changes as well as other details are contained in the appendix. TFP is not utilisation adjusted, and is calculated allowing different weights in the production function for different sub-types of capital. From the total of 20 countries in the database I drop Switzerland due to missing wage data, and Korea due to irregularities in the wage data. Ireland is initially left in the sample, but is dropped as an outlier.[1] This leaves a total of 17 countries in the baseline comparison.

I study a cross section of countries during the Great Recession. The OECD data is annual, and runs up to 2011, starting at different dates for different countries, the earliest being 1970. I take the Great Recession to be the period from 2008 to 2011, and study correlations between various log-changes over this period.

I detrend TFP using a constant growth rate estimated for each country using all available pre-crisis data. The motivation for this is twofold: Firstly to control for the pre-crisis growth rate of each country. Secondly, I am interested in how a financial crisis affects TFP, and hence want to study deviations from longer term productivity trends. Consistent with standard macroeconomic models where the wage grows at the same rate as productivity in the long run, I also detrend the real wage rate using the same deterministic time trend. Hours worked is expressed in per capita terms, and not detrended. The price level is expressed in levels and is not detrended.

[1] Ireland has been dropped as an outlier because its experience has been extreme relative to the rest of the sample. Ireland experienced the worst fall in hours (over 20%) in the sample, and third worst fall in TFP (over 10%). While the negative TFP-hours correlation holds robustly across the rest of the sample, Ireland is clearly a counterexample of a country which experienced both very bad TFP and very bad labour market performance and is worthy of further independent study. The exercises with Ireland included are presented in the appendix.

2.3.2 Simple correlations

I perform two exercises with this data. The first is a simple cross-country comparison. For each country, I construct the log change in variable x over the Great Recession as $\tilde{x}_i = \log(x_{i,2011}/x_{i,2007})$. I plot selected relationships between my four variables (detrended TFP, z, hours per capita, l, detrended real wages, w, and the price level, P) in Figure 2.1. It is worth noting at this point that I am focusing purely on correlations, and am not making causal statements. I also am not controlling for any other covariates. Additionally, given the small sample size it is hard to show significance of the correlations. I thus hold off reporting coefficients and significance levels until the next exercise in which I use a panel structure to exploit more variation and control for country and time fixed effects.

The top left panel gives the key correlation: countries which experienced larger falls in hours over the recession tended to experience less severe declines in TFP. Interestingly, the UK and Finland appear as slight outliers, experiencing worse TFP growth than other countries for a given fall in hours. This suggests that perhaps part of the UK productivity puzzle may lie in factors very specific to the UK.

The top right and bottom left panels give the correlations between wages and TFP and hours respectively. Countries which experienced higher wage growth experienced higher TFP growth. On the other hand, countries which experienced higher wage growth experienced larger falls in hours. Notice that these three panels all speak against a simple TFP shock interpretation of the data: in that case we would expect a positive correlation between hours, wages and TFP. Indeed, the downward slope of the relationship between hours and wages suggests that movements along the labour demand curve dominate the evolution of labour markets over the period. These correlations motivate my choice of differential wage adjustment as a potential explanation of the negative TFP-hours correlation.

Finally, the bottom right panel plots the correlation between price changes and real wage changes over the period. Countries with higher inflation experienced lower real wage growth. This suggests a role for sticky nominal wages combined with differential inflation outcomes in determining real wages. These four figures trace out all of the elements of my story: countries with higher real wage growth in the crisis experienced worse falls in hours but better TFP growth, with variation in wages partly driven by inflation.

Figure 2.1: Relationships between selected variables across countries.

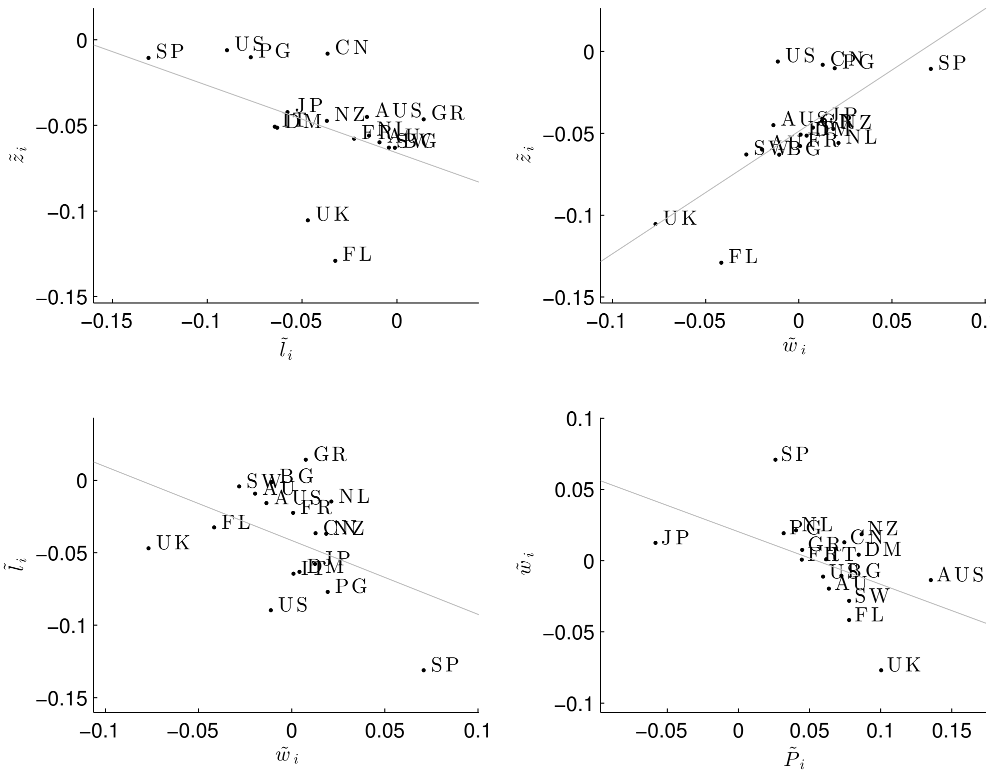

Lines are OLS lines of best fit between the two variables. The country names refer to: AUS = Australia, AU = Austria, BG = Belgium, CN = Canada, DM = Denmark, FL = Finland, FR = France, GR = Germany, IT = Italy, JP = Japan, NL = Netherlands, NZ = New Zealand, PG = Portugal, SP = Spain, SW = Sweden, UK = United Kingdom, US = United States.

2.3.3 Partial correlations

While I am not looking to uncover causal relationships, there are obviously problems with taking such a simplistic cut of the data. For example, the correlations above could be driven not by differential wage behaviour during the recession (as my story claims), but simply by preexisting differences across countries. I would thus like to construct evidence that the above correlations hold within a hypothetical country over the recession. I do this by creating a short panel structure from my data. Instead of just looking across countries, I now look across both countries and time. Specifically, for each variable x I construct the log change from 2007 to year t in country i: $\tilde{x}_{i,t} = \log(x_{i,t}/x_{i,2007})$. I do this for $t = \{2008, 2009, 2010, 2011\}$ giving me four years of data across 17 countries, and a total of 68 data points. Intuitively, this lets me look at the relationship between variables both across and within countries.

Partial correlations capture the relationships between variables after controlling for their relationships with other variables. For example, the relationship between TFP and hours after controlling for the fact that some of their correlation derives from the fact that they both depend on country characteristics. I control for country, f_i, and time, q_t, fixed effects, allowing me to focus on variation in the variables unrelated to country characteristics, and the year of the recession. The first step is to regress each of the variables of interest on the (common) control variables and calculate the residual:

$$\hat{x}_{i,t} = \tilde{x}_{i,t} - \beta'_x X_{it,} \tag{2.1}$$

$\hat{x}_{i,t}$ is the residual for variable $\tilde{x}$. $X_{i,t}$ is the set of controls which contains only the country and time fixed effects. β'_x is the OLS estimator of $\tilde{x}$ on X. The residual is thus the component of $\tilde{x}$ not explained by the control variables. The partial correlation between variables $\tilde{x}^1$ and $\tilde{x}^2$ is then the correlation between their residuals, $\hat{x}^1$ and $\hat{x}^2$, which is the correlation between the components of $\tilde{x}^1$ and $\tilde{x}^2$ not explained by X.

In Figure 2.2 I repeat Figure 2.1, replacing the values of the variables with the residuals used in the partial correlations (each country now has four datapoints in each plot, one per year). The relationship between the variables are still apparent in this exercise, highlighting that the correlations appear to be driven by events *during* the Great Recession, and not driven by preexisting country characteristics.

Figure 2.2: Relationships between selected variables: Panel structure

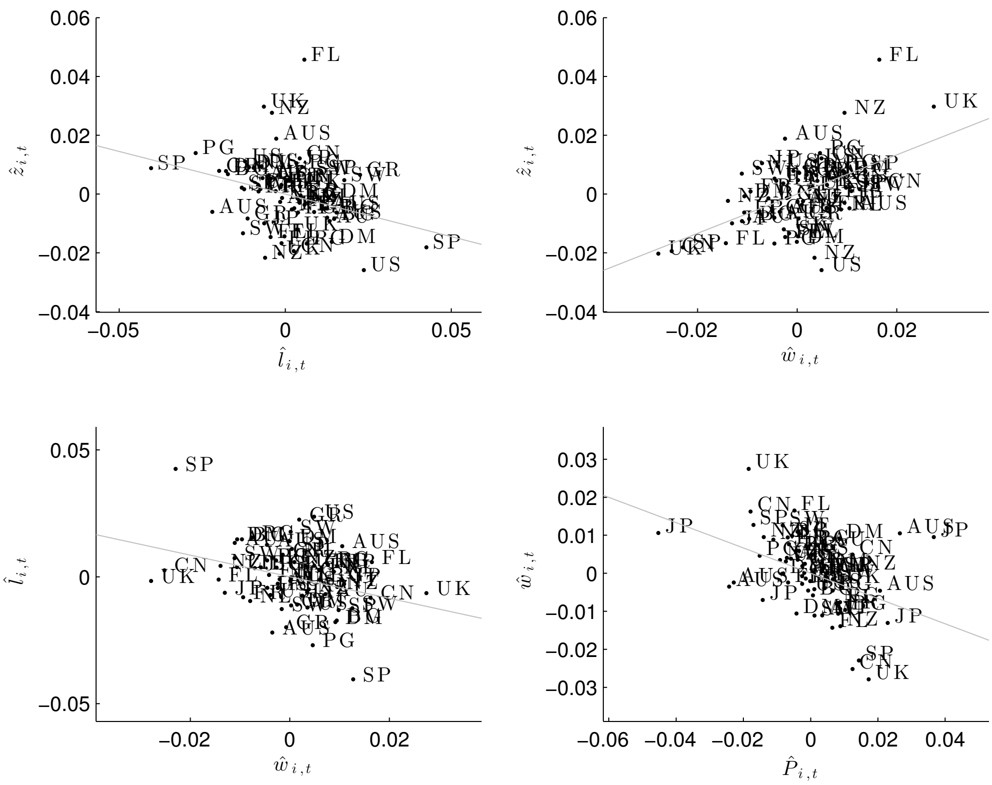

Lines are OLS lines of best fit between the two variables. The country names refer to: AUS = Australia, AU = Austria, BG = Belgium, CN = Canada, DM = Denmark, FL = Finland, FR = France, GR = Germany, IT = Italy, JP = Japan, NL = Netherlands, NZ = New Zealand, PG = Portugal, SP = Spain, SW = Sweden, UK = United Kingdom, US = United States.

The panel structure, by giving me more power, also allows me to more precisely measure the sizes of the partial correlations, and test their precision. Table 2.1 gives the estimated correlations between all four variables. All of the correlations are significant at at least the 5% level.[2] The strong negative correlation between TFP growth and hours growth is apparent in the correlation coefficient of -0.2963. TFP and wages are positively and negatively related to real wages respectively, with partial correlations of 0.5304 and -0.3307. Finally, the correlation between real wages and prices is -0.4265.

Table 2.1: Partial correlations

	TFP	l	w	P
TFP	1	-0.2963**	0.5304***	-0.3616**
l		1	-0.3307**	0.3405**
w			1	-0.4265***
P				1

, **, and * denote significance at the 10%, 5%, and 1% levels respectively.*

2.3.4 Robustness

To check the robustness of my results I perform several checks. Firstly, I drop several countries and groups of countries in turn to check if the results are driven by outliers. For example, in Figure 2.2 Spain, the US, Finland and the UK all appear towards the extremes of the plots. The signs and magnitudes are unaffected by dropping these countries individually or as a group. There might also be concern with the reliability of wage data. I also perform the exercise comparing TFP and hours to prices instead of real wages. If nominal wages are sticky then higher prices should lead to lower real wages, identifying a similar effect. This gives a negative relation between TFP and prices, and positive between hours and prices, which supports the original result.

Table 2.2 gives the results of two additional robustness exercises. One concern

[2] The panel structure could introduce serial correlation into the errors, which could be be a problem for the significance tests. However, given the relatively small sample size, clustering the standard errors is problematic.

with the partial correlations I have presented is that, even after controlling for time and country fixed effects, I am comparing countries who have suffered different-sized financial crises. Perhaps a more ideal comparison would be to control for the size of the financial crisis and see if, conditional on this, the same negative relationship between TFP and hours emerges. To this end, I collect data on the "credit intermediation ratio" from the OECD National Accounts.[3] This is the ratio of loans from the financial sector to the non-financial sector to the total liabilities of the non-financial sector. In other words, it is a measure of the ability of non-financial firms to raise funds from the financial sector. This ratio has strong predictive power for output: in the sample of countries used the correlation between the growth rates of output and the credit ratio is 0.4685, significant above the 0.1% level. Panel A of Table 2.2 gives the four main partial correlations, where I now control for both fixed effects and the credit intermediation ratio. The correlation between TFP and hours is now stronger and more statistically significant, and the other correlations are at least as large and significant as before.

I next perform an exercise designed to see if my proposed mechanism for the correlation between TFP and hours fits the data. Specifically, my proposed mechanism links both TFP and hours to movements in wages: lower wages lead to higher hours and misallocation which reduces TFP. Thus I should see that the partial correlation between TFP and hours is reduced towards zero if I also control for wages, since I claim that their negative relationship derives from wage movements. Panel B of Table 2.2 gives values of this partial correlation where I variously control for wages and prices, in addition to the fixed effects.[4] The partial correlation is halved once

[3] The item can be found under Financial Dashboard, Financial Indicators - Stocks, Private Sector Debt. The data is not available for New Zealand, so I drop it for this exercise. Additionally, there are concerns with the data for Finland. Specifically, a casual plot of output and the credit ratio during the crisis reveals a strong, positive relationship for all countries, except for Finland who experienced a severe recession while the credit ratio increased making it a severe outlier. Finland is thus dropped, but the results are unaffected by including it.

[4] Note that while wages and prices are both endogenous objects, controlling for them in this manner is allowable since I am still only computing partial correlations. This exercise measures the correlation between TFP and hours which is unrelated to their mutual correlation with wages.

Table 2.2: Robustness

2.A: Controlling for financial variables

	corr(z, l)	corr(z, w)	corr(l, w)	corr(w, P)
	-0.4574***	0.5286***	-0.5192***	-0.4286***

2.B: Controlling for wages and prices

	Baseline	w	P	$[w, P]$
corr(z, l)	-0.2963**	-0.1511	-0.1975	-0.1147

, **, and * denote significance at the 10%, 5%, and 1% levels respectively. Panel 2.A gives partial correlations after controlling for country and time fixed effects, and for the "credit intermediation ratio" in country i at time t. New Zealand and Finland are dropped, the former due to missing data and the latter due to concerns with the credit data. Panel 2.B gives the partial correlation between TFP and hours after controlling for fixed effects, and the variables given in the column headers.*

wages are controlled for and becomes statistically insignificant, thus supporting that the relationship between TFP and hours is related to a large extent with wage movements. I also control for prices, since prices should be related to real wages if nominal wages are sticky, and again find that the relationship is smaller and insignificant. Controlling for both prices and wages together leads to the largest reduction in the partial correlation between TFP and hours.

Finally, all results are robust to changing the detrending assumptions for hours, wages and prices. The results are weakened is TFP is not detrended, however detrended TFP appears to be the appropriate measure as discussed above. I want to investigate how the financial crisis caused TFP to change, and hence want to study deviations from existing long-term trends in TFP.

Overall, the international evidence paints an interesting picture. During the crisis TFP and hours are negatively related, and this seems to have a lot to do with real wages. This is true in simple cross sections, and also in a country-year panel structure after controlling for both fixed effects and a credit variable. Finally, real wages are negatively related to prices, consistent with nominal wage stickiness.

2.4 Case study: US vs UK

In this section, I take the US and the UK as two case studies with which to take a closer look at the results of the previous section. I first perform a business cycle accounting exercise, and show that business cycle "wedges" are correlated with wages in a way we would expect given the cross-sectional evidence from the previous section. I then decompose the real wage changes into nominal wage and inflation components, and discuss institutional factors behind wage stickiness in the two countries. In the appendix I perform other exercises: a) a decomposition of the labour wedge, b) a discussion of the robustness of my results to composition effects in the real wage data, and c) a discussion of why wages, and not real unit labour costs, are the appropriate measure for my exercise.

Here I do not use the OECD data from the previous section, and instead use national accounts data from the US and the UK. This has the advantage of giving longer time series, including data later into the Great Recession than the OECD data, which allows me to examine both the initial phase of the crisis and the following "recovery".

2.4.1 Business cycle accounting

In this section, I perform Chari, Kehoe & McGratten's (2007) business cycle accounting procedure on the US and the UK over the Great Recession. This procedure specifies a "prototype economy", which is a standard real business cycle model augmented with hypothetical wage and capital income taxes. These two taxes, along with the level of TFP and government spending, give four "wedges". The idea is then to select four time series from the data, and back out the values of these four wedges such that the equilibrium of the prototype economy exactly matches the data on these time series. I provide a brief description of the procedure here, and provide a more detailed derivation and discuss data sources in the appendix.

Prototype economy

The prototype economy features a standard Cobb-Douglas production function in capital and labour:

$$y_t = e^{\tau_t^e} k_{t-1}^{\alpha} \left((1+g_z)^t l_t\right)^{1-\alpha} \tag{2.2}$$

Here a lowercase refers to per-capita variables, and y_t, k_{t-1}, and l_t refer to per-capita output, capital, and labour respectively. g_z is the (estimated) long run trend growth rate of TFP, and the Solow residual at time t is given by $e^{\tau_t^e}(1+g_z)^t$. The efficiency wedge, τ_t^e, is thus defined as the (log) deviation of TFP from its long term trend, and can be backed out directly from the production function given data on output, capital, and labour.

The labour wedge is a hypothetical percentage tax on labour income, τ_t^l. The prototype economy has period utility function $U(c,l) = \log(c) - v(l)$, which leads to a standard labour optimality condition augmented with the labour wedge:

$$v'(l_t)c_t = (1-\tau_t^l)(1-\alpha)\frac{y_t}{l_t} \tag{2.3}$$

This equation equates the marginal rate of substitution between labour and consumption to the marginal product of labour, adjusted for the labour wedge. The labour wedge can then be measured from the data using this equation and data on output, hours, and consumption. The full wedges procedure also requires computing an investment wedge and government spending wedge.[5] Once the wedges are computed, the final procedure is to solve and simulate the prototype economy subject to the

[5]The full procedure requires estimating a law of motion for the wedges and solving for the policy functions of the prototype economy, which depend on this law of motion.

measured wedges. Trivially, simulating the economy with the realised sequence of all four wedges will lead to paths for output, consumption, capital, and hours which exactly replicate the data. The importance of each individual wedge can be evaluated by simulating the prototype economy over a given period subject only to movements in any one wedge, a procedure which I detail in the appendix along with the parameter and functional form assumptions.

Wedges

As is commonly found in these exercises,[6] the investment and government wedges account for very little movement in output, and hence I focus on the efficiency and labour wedges. Figure 2.3 summarises the results, with the top row showing graphs for the US, and the bottom the UK. The first column summarises the exercise from the beginning of the financial crisis up to the most recently available data. My data is quarterly, so I start in the first quarter of 2008.[7] The top left panel shows the dramatic fall in US output over this period relative to trend: a rapid fall of nearly 7% over the first year, followed by a further gradual decline. We see that the labour wedge is able to account very well for the fall in output, even slightly over-predicting it initially, and continuing to provide a reasonable account over the whole period. The efficiency wedge, on the other hand, is not able to account for the dynamics of output at all, and predicts that output should be slightly above trend over the whole period.

The UK shows the opposite pattern, with the efficiency wedge doing a better job at explaining output than the labour wedge. The bottom left panel documents the dramatic fall in output relative to trend, which is around 9% within the first

[6]See, for example, the original Chari, Kehoe, & McGratten (2007) paper.

[7]The NBER dates the US recession as starting in December 2007, and the results are robust to moving the start point by a few quarters in either direction. Interestingly, the data actually suggests that, relative to the long term trend, the US (but not UK) economy started slowing down in early 2006, and that this slowdown is associated with the efficiency wedge. Since I want to focus on financial crises , and this period precedes the events of 2008, I will not focus on it in this paper.

year, and consistently two percentage points worse than the US over the period. The efficiency wedge explains the initial fall well, and continues to explain the bulk of the fall in output over the period. The labour wedge, while getting the sign of the output movements right, cannot match the magnitude of the fall, and additionally predicts that output should have returned to trend by 2014.

The remaining two columns provide additional information. The middle column plots the data on hours worked. A reflection of the more important labour wedge in the US is the worse performance of hours during the recession: it fell by nearly 11% in the US, whereas it fell by just over 5% in the UK. As is to be expected, the labour wedge is important for explaining the movements in hours in both the US and the UK: it is well known that the benchmark RBC model is unable to generate the required movements in hours worked, which explains why the labour wedge still plays a role.

The final column plots the wedges themselves. For the efficiency wedge, I plot the exponential of the wedge, which gives the deviation of TFP from trend. For the US, we see a peak increase of over 2% in detrended TFP during the crisis, which is eventually reversed, but still leaves detrended TFP less than 1% below trend by the end of the sample period. The UK, on the other hand, has TFP fall by 5% from trend within the first year, and continue falling to around 7% below trend by the end of the sample period. The labour wedge increases in both countries, but by roughly twice as much in the US. Additionally, the US labour wedge remains severely elevated at the end of the sample period, whereas the UK labour wedge actually ends the sample less severe than at the beginning.

To summarise, as in the OECD data the key empirical observation I want to explain is the negative TFP-hours correlation: Why did the US experience a much larger fall in hours than the UK, while experiencing a much less severe response of measured TFP? The accounting exercise above allows me to also express this question in terms of wedges: Why did the financial crisis manifest itself more in the labour

Figure 2.3: Business cycle accounting results

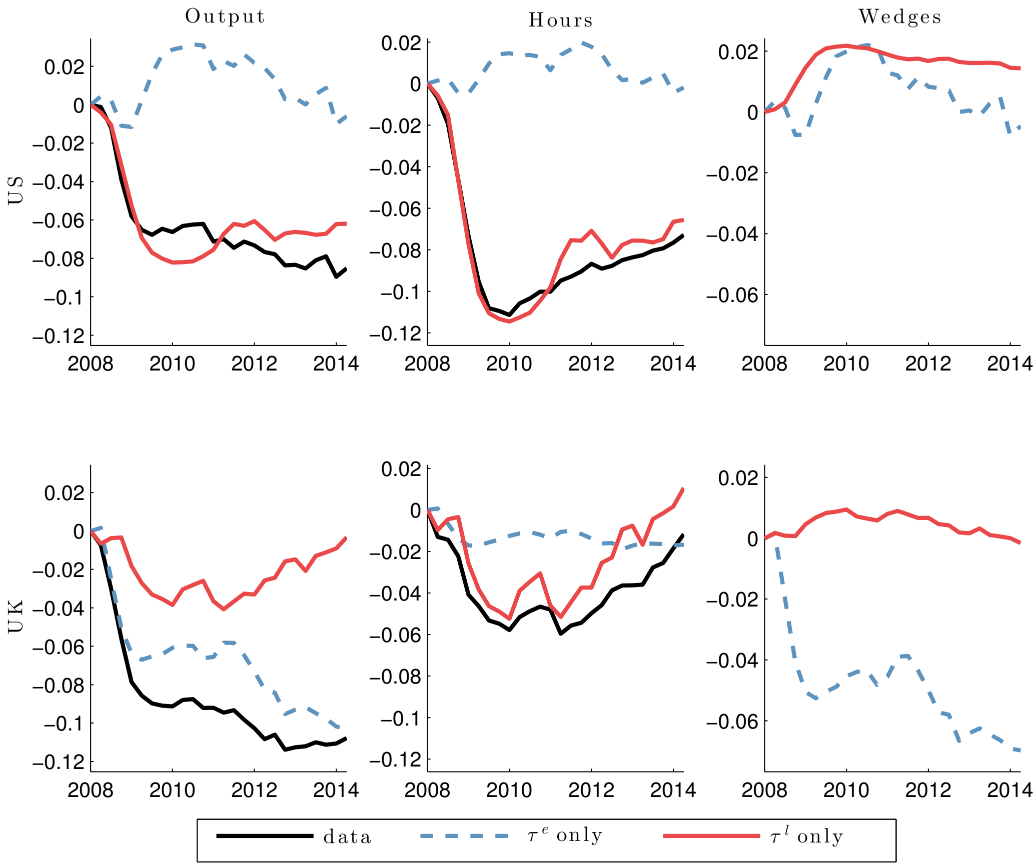

All variables are expressed as a fractional deviation from the value in the initial period. The first two columns plot output and hours, and their simulated paths subject only to one wedge. The final column plots the exponential of the efficiency wedge (giving the deviation of TFP from trend), and the value of the labour wedge. Output is expressed as the deviation from the estimated trend.

wedge for the US, and the efficiency wedge for the UK?

Wages & wedges

Using national accounts data, I construct a measure of average real wages for the two countries, and plot this against the wedge and output data in Figure 2.4. All of the variables, including real wages, are detrended using the average trend growth in TFP.

The first column covers the whole sample period, repeating the first column of

Figure 2.3 with the added real wage data. The data is surprising: despite the US having nearly double the fall in hours of the UK, detrended real wages only fell by 6% over the whole sample period, whereas they fell by 10% in the UK. Additionally, in the first two years of the crisis, during which the US saw its dramatic decline in hours, real wages (detrended and in levels) actually *increased* by over 2% in the US. The UK, on the other hand, has seen a nearly secular decline in detrended real wages over the whole period.[8]

This provides the complementary result to my wage correlations in the OECD data. There I showed that wages were positively correlated with TFP and negatively correlated with hours over the crisis. Here I show that high wages are associated with the recession manifesting itself in the labour wedge, and falling wages are associated with the recession manifesting itself in the efficiency wedge.

Given that the behaviour of wages in the US changes over the sample period – initially rising, and then starting a decline – it is interesting to check if these correlations hold over time within the US too. That is, does the role of the labour and efficiency wedges change in the US when detrended real wages start to fall? The last two columns of Figure 2.4 demonstrate that this is the case, by splitting the sample in half. For the US, we see that in the first three years, when detrended real wages remained elevated, the labour wedge explains the behaviour of output, while the efficiency wedge cannot. For the last three years, where detrended real wages are falling, the labour wedge is actually improving and now it is the efficiency wedge that better explains the fall in output. This time-series relationship between the wages and the role of the two wedges exactly matches the cross-country relationship I described earlier between the US and the UK.

[8] One concern with aggregate wage data is that it contains composition bias if workers on the lowest wages are fired first. In the appendix, I summarise the evidence on composition effects during the crisis for the US and UK. Adjusting for estimates of composition effects leads instead to US wages remaining roughly constant for the first two years of the crisis, and does not affect the UK estimates. Thus there still appears to be a large difference in wage behaviour after controlling for composition.

Figure 2.4: Real wage behaviour and the output decomposition

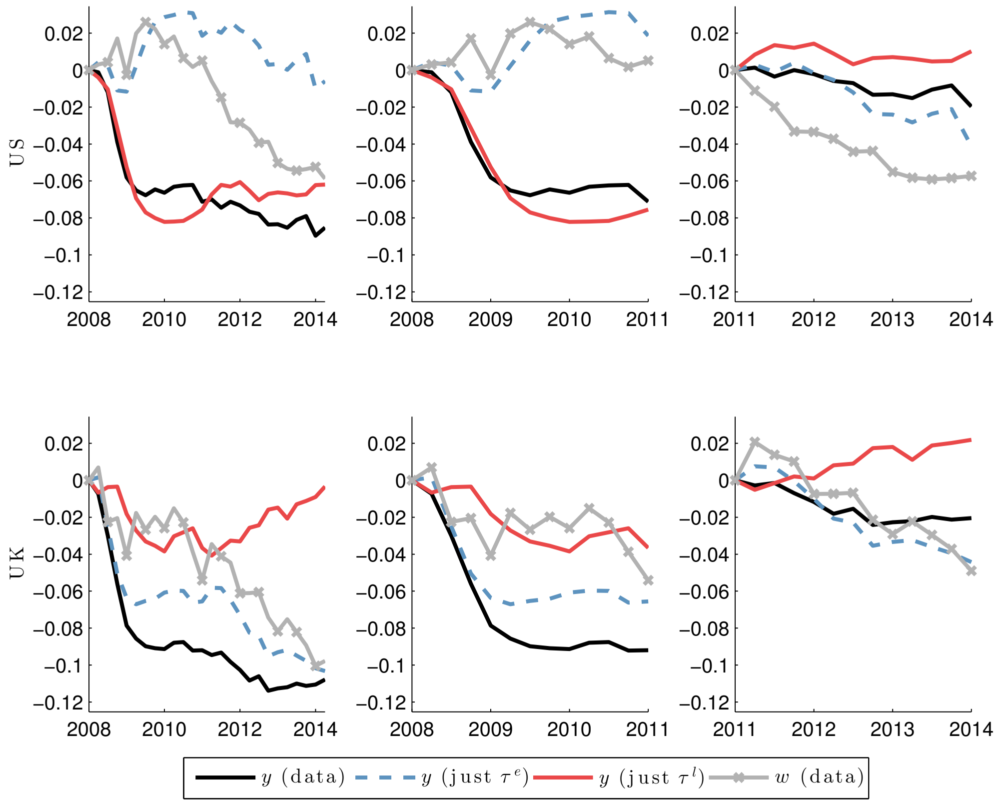

All variables are deviations from the estimated constant trend growth rate, expressed as a fractional deviation from the value in the initial period.

Finally, in the appendix I provide a decomposition of the labour wedge following Karabarbounis (2014). This exercise uses wage data to decompose the labour wedge into two wedges, one reflecting distortions to the household's labour supply first order condition, and the other distortions to the firm's labour demand first order condition. This exercise shows that the measured labour wedge is explained almost entirely by distortions to labour supply, and not demand, during the Great Recession for both countries. This is consistent with unemployment caused by rationing due to wage stickiness, which is precisely the form of unemployment I will have in my model.

2.4.2 Inflation, institutional factors, and real wage determination

In my model I will be using the idea of wage stickiness, and it is thus important to try and understand what drives wages, if it is not market clearing. In this section I decompose real wage growth in both countries into components driven by nominal wage growth and inflation. This allows me to look for patterns hinting at what kind of frictions in wage adjustment are relevant, for example nominal or real. I measure prices using the GDP deflator, with the source listed in the appendix. Using the definition of real wages, $w_t = W_t/P_t$, in any interval $[t_0, t_1]$ I calculate the following decomposition of real wage growth between t_0 and t:

$$\log\left(\frac{w_t}{w_{t_0}}\right)\frac{1}{t_1 - t_0} = \log\left(\frac{W_t}{W_{t_0}}\right)\frac{1}{t_1 - t_0} - \log\left(\frac{P_t}{P_{t_0}}\right)\frac{1}{t_1 - t_0} \tag{2.4}$$

I plot the results of this decomposition over three time periods in Figure 2.5. Note that by plotting the negative of the inflation component the graphs are constructed so that adding up the nominal wage and inflation components at any time recovers the real wage. Whenever nominal wage growth is larger than inflation we have real wage growth, and the real wage line thus lies above zero, and vice versa.

For $t = t_1$, this equation decomposes the average yearly real wage growth in $[t_0, t_1]$ into average nominal wage growth and average inflation. This decomposition is given by the final values on each graph: for example, the top left panel shows that between 2002 and 2008 the US experienced average annual real wage growth of roughly 1.5%, nominal wage growth of 4% and inflation of 2.5%. For $t_0 < t < t_1$ the equation also decomposes real wage growth into the same two components, however since we are dividing by $(t_1 - t_0)$ and not $(t - t_0)$ the numbers should not be interpreted as yearly averages. Yearly averages can be recovered by multiplying by $(t_1 - t_0)/(t - t_0)$, and the graphs are instead meant to provide an illustration of each variable's path during the sample period.

Figure 2.5: Decomposing the real wage

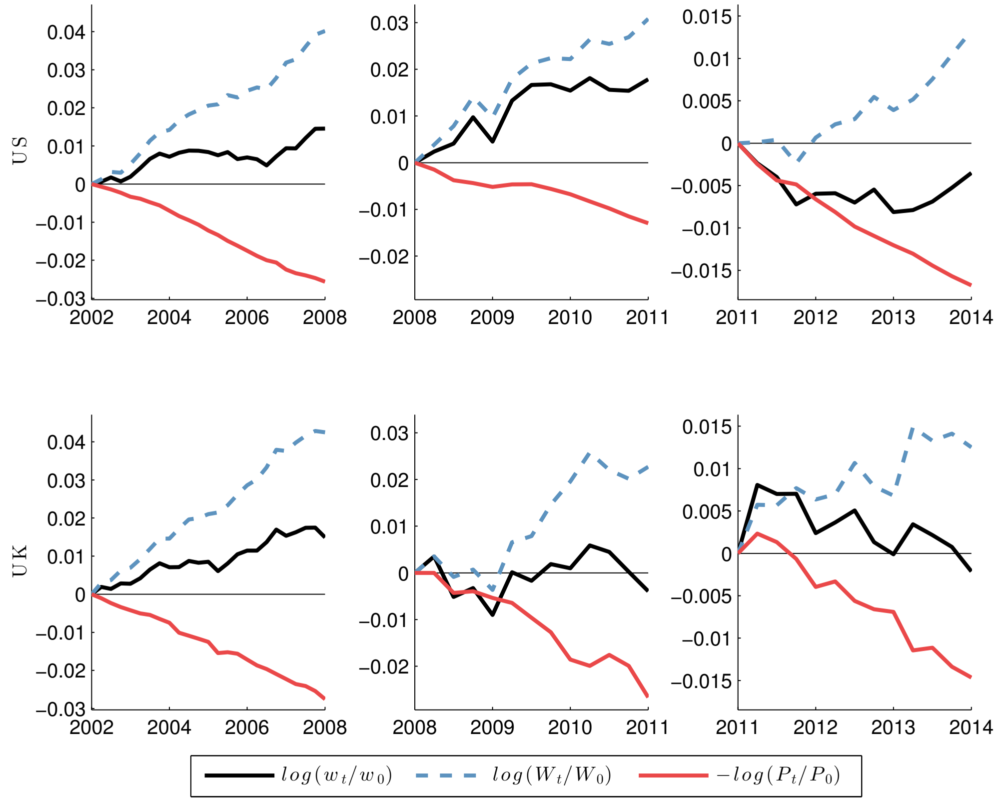

All variables are expressed as a fractional deviation from the value in the initial period. Variables are not detrended. Instead of inflation I plot the negative of inflation, such that adding up the nominal wage and inflation lines will give you the real wage at any point in time.

The first column provides a decomposition of real wage growth pre-crisis. I take the period 2002-2008 since this corresponds to the period between the NBER's dating of the end of the 2001 recession and start of the 2008 recession.[9] The second and third columns decompose real wage growth for the two halves of the Great Recession sample period.

One fact that emerges from these graphs is that the US and the UK experienced

[9]The UK did not experience an equivalent recession in 2001, and the date range is taken to be the same for comparability.

similar paths for nominal wages pre and post-crisis despite their different labour market outcomes. Specifically, nominal wage growth was running at around 4% a year in both countries in the six years preceding the crisis, and dropped to roughly 2%-3% a year over the next three years. It then dropped to around 1.5% a year over the final three years of the sample. Given the differences in inflation and unemployment between the two countries over the crisis, this similarity is somewhat surprising.

Given the similarities in the behaviour of nominal wages, inflation explains a significant portion of the differences in real wages across the two countries over this period. While inflation ran at similar levels pre-crisis, inflation between 2008 and 2011 ran at over 2.6% per year in the UK, and only 1.3% in the US. Coupled with slightly lower nominal wage growth in the UK, this translates to real wage growth roughly 2% higher per year in the US over this period. If we calculate the fraction of the difference in the two countries' real wage growths during this period explained by each component, 63% is explained by the difference in inflation rates. An alternative calculation is to calculate a hypothetical real wage growth gap during this period if nominal wages in each country continued to grow at their 2002-2008 rates. In this case, the hypothetical wage growth gap is 1.1%, whereas the true gap is 2.2%. Doing the same but holding inflation fixed instead generates a hypothetical wage gap of 1%. Thus there remains a significant role for nominal wage differences in explaining the real wage differences, but with inflation playing an equal if not slightly larger role.

Overall, the similarity of the nominal wage paths of the two countries despite their vast differences in unemployment suggests a role for nominal wage stickiness. Decomposing real wage growth into nominal wage and inflation components reveals that both components do play a role in accounting for the differences between the two countries, but that inflation plays a slightly larger role.

It is actually not important for my story whether the (lack of) wage adjustment is driven by nominal wage stickiness and inflation, or other causes. This leaves open the interesting possibility that institutional factors might play a role in explaining the portion of the real wage difference between the US and the UK explained by nominal wage changes. In particular, there are several factors that could potentially explain why wages in the US could have recently become more (downwardly) rigid, while wages in the UK have become less rigid.

In the UK, there has been a trend towards increasing labour market flexibility since the 1980s, which has been relatively untested since the last recession before the Great Recession was the 1990-91 recession. Blundell, Crawford & Jin (2014) summarise these changes, which include "the increasing number of welfare-to-work programmes available to jobseekers, the more stringent job search conditions attached to benefits claimed by the unemployed, those with disabilities and lone parents, and, more recently, the increase in the state pension age for women". More formally, Gregg, Machin, & Fernandez-Salgado (2014) show that unemployment has become more of a moderating force on real wages in the UK, and even identify a structural break in the unemployment-wage relationship around 2003. They argue that part of the reason that wages have become more responsive to unemployment is declining union membership, which has been falling steadily since the 1980s. They show that union wages are less responsive to unemployment than non-union wages, and hence that declining union membership can partly explain why overall wages have become more flexible.

In the US, on the other hand, some papers have argued that the extension of unemployment benefits during the Great Recession could have put upwards pressure on wages. Hagedorn, Karahan, Manovskii, & Mitman (2015) use an identification strategy that exploits a policy discontinuity at state borders to estimate that unemployment in the US could have been as much as 2.5pp higher in 2011 due to the extension of benefits. Using the same methodology, Hagedorn, Manovskii & Mitman (2015) argue that 61% of the increase in employment in 2014 can be attributed to the expiration of the benefit extension. They claim that this is due to a "macro effect" whereby more generous unemployment benefits increase the outside option of workers, increasing their bargaining power and hence pushing up wages. Consistent with this, they find a statistically significant, positive relationship between wages and benefits using their identification strategy.

In a recent paper, Mulligan (2015) analyses the effects of taxes and fiscal policies on the incentives of firms to hire and workers to provide labour in the US and UK. He finds that during the first three years of the crisis, changes in taxes and subsidies in the US reduced employees' reward to work, whereas changes in the UK actually increased employees' reward to work. These policies, by relatively improving the outside option of not working in the US and worsening it in the UK, could explain why real wages behaved so differently in the first three years of the crisis. Interestingly, much of these changes were reversed in both countries during the second half of the crisis, which could explain why I find that the the two countries are more similar in the second half of the crisis.

Of course, the above policy changes only say that wages in the US could have become more rigid, while those in the UK have become less so. They do not say that wages in the US have become *more rigid than* those in the UK. However, combined with the differential inflation experiences of the two countries, they paint an interesting picture of the forces that could have pushed us in that direction.

2.5 Model

In the remainder of the paper I build a model to explain the above correlations. The model is an adaptation of Buera & Moll (forthcoming). They set up a model with heterogeneity across firms which they use to show how a credit crunch can be transmitted into the efficiency wedge by disrupting the allocation of resources across firms. The main difference between my model and theirs is that I consider flexible vs. sticky wages, and thus show how the transmission in their model is dependent on assumptions on wage setting, which I exploit for my main result.[10] In particular, I will consider two extreme assumptions on wage setting, completely flexible and completely fixed wages, and study how the transmission of the credit crunch differs in these two cases. These results provide benchmarks between which different countries fall, depending on their level of wage adjustment during the crisis. Time is discrete and the horizon infinite, and I consider only perfect foresight economies in which there is no aggregate uncertainty. There is idiosyncratic uncertainty at the firm level, but this will not lead to aggregate uncertainty due to the law of large numbers and my assumption of a continuum of firms. The model is set up as a real model, and nominal concerns will only enter via the effect of inflation on the real wage.

2.5.1 Household

Since the focus of the model is on the firm side, the household structure is left stylised. There is a representative household with discount factor β, period utility function $U(c_t, L_t^s) = \log(c_t) - v(L_t^s)$ where c_t is consumption and L_t^s hours worked. Labour disutility satisfies $v'(L^s) > 0$ and $v''(L^s) \geq 0$, as is standard.[11] The household receives labour income and any profit paid out by firms, and can borrow or save using a risk free bond, which is traded with firms. It can also choose consumption, and must give some income as equity injections for firms. The household's budget constraint is thus:

$$c_t + b_t = w_t L_t^s + r_{t-1} b_{t-1} + D_t \tag{2.5}$$

[10] I also model the firms as firms owned by households, whereas they model them as a seperate species. I also allow households to lend to firms, while they do not.

[11] The assumption of log utility is only for consistency with the use of log utility in the empirical wedges exercise. All of the results below go through for an arbitrary consumption utility function satisfying the usual conditions.

Where b_t is savings, r_t the real interest rate, and D_t the net income received from firms. The first order condition for bonds yields the familiar Euler equation:

$$u'(c_t) = \beta u'(c_{t+1})r_t \tag{2.6}$$

The household takes wages as given, and chooses how much labour to supply. Since I assume sticky wages in some cases, and given the simple labour market structure, we might end up with labour market rationing in some states of the world. The household understands this and takes it into account in its optimisation. If the household is not rationed in its labour supply in equilibrium in period t, then its labour supply first order condition must hold:

$$w_t = v'(L_t^s)c_t \tag{2.7}$$

However, if in equilibrium the market doesn't clear the household understands that it will only be able to work $L_t^s \leq L_t^d$ hours, where L_t^d is aggregate labour demand from firms, and the labour supply condition will not hold at time t.

2.5.2 Firms

There is a continuum of mass one of firms, all owned by the household. They each operate a constant returns to scale production function in capital and labour, and have heterogeneous productivities. Firm i has production function $y_t = z_t k_{t-1}^{\alpha} l_t^{1-\alpha}$ where I suppress i subscripts for clarity. z_t is firm level TFP, k_{t-1} capital operational at time t, and l_t labour hours. They can raise funds by issuing risk-free debt, but face stylised borrowing and equity constraints. Specifically, firms cannot raise any equity, and can only borrow up to a multiple λ_t of net worth, n_t. Note that net worth is defined as the marked-to-market book value of firms assets less their liabilities. Additionally, I assume that they cannot short capital, so that $k_t \geq 0$.[12] As is standard in the

[12]This restriction follows naturally from the production function, if we assume it does not admit negative capital. Alternatively, this can be viewed as a restriction on the firms to trade claims which short their future output.

literature, I assume that firms exogenously exit with probability $1-\sigma$ each period, at which point they pay out all their net worth as dividends. Given that firms may always be financially constrained in the future they will never pay out dividends until they exit,[13] so their balance sheet constraint is simply:

$$k_t = n_t + d_t \tag{2.8}$$

Where d_t is borrowing, and capital is funded by net worth and borrowing. To derive the evolution of net worth, it is first helpful to study the firm's labour demand choices. Firms can hire and fire workers without frictions, and take the (common) real wage, w_t, as given. Static profit is maximised, for a given level of z_t and k_t, by choosing labour to satisfy:

$$\pi_t = \max_{l_t} z_t k_{t-1}^{\alpha} l_t^{1-\alpha} - w_t l_t \tag{2.9}$$

Solving this problem gives:

$$\pi_t = \pi(z_t, w_t) k_{t-1} \tag{2.10}$$

Where

$$\pi(z_t, w_t) = \left((1-\alpha)^{\frac{1}{\alpha}-1} - (1-\alpha)^{\frac{1}{\alpha}}\right) z_t^{\frac{1}{\alpha}} w_t^{1-\frac{1}{\alpha}} \tag{2.11}$$

Profit per unit of capital is increasing in productivity ($\pi_1(z_t, w_t) > 0$) and decreasing in wages ($\pi_2(z_t, w_t) < 0$). The key thing to note here is that profit is linear in capital, which will be useful later. Optimised labour and output are also given by:

$$y_t = (1-\alpha)^{\frac{1}{\alpha}-1} z_t^{\frac{1}{\alpha}} w_t^{1-\frac{1}{\alpha}} k_{t-1} \tag{2.12}$$

$$l_t = (1-\alpha)^{\frac{1}{\alpha}} z_t^{\frac{1}{\alpha}} w_t^{-\frac{1}{\alpha}} k_{t-1} \tag{2.13}$$

Following Buera & Moll, I make the following assumptions on the timing of firms' capital choices and productivity realisations. I assume that firms choose capital one

[13] A formal proof of this statement and the conditions under which it holds is given in the appendix.

period in advance of production. Less standardly, I assume that firms know their idiosyncratic productivity one period in advance, and hence know their productivity at $t+1$ when choosing capital to be used in $t+1$.[14] This assumption means that firms know at time t their return on capital between t and $t+1$:

$$r^k_{t+1} = \pi(z_{t+1}, w_{t+1}) + 1 - \delta \tag{2.14}$$

Notice that, due to labour optimisation and the constant returns to scale assumption, an individual firm's return on capital is independent of its capital choice. We can now state the evolution of firm net worth as:

$$n_{t+1} = (\pi(z_{t+1}, w_{t+1}) + 1 - \delta)k_t - r_t d_t \tag{2.15}$$

Which, combined with the balance sheet constraint, yields:

$$n_{t+1} = ((\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t)\phi_t + r_t)\, n_t \tag{2.16}$$

Where I have defined leverage as $\phi_t = k_t/n_t$. The no-shorting and borrowing constraints can be expressed in terms of a constraint on leverage: $0 \leq \phi_t \leq (1+\lambda_t)$.This allows me to define the firm's problem. Let firm value at time t for a given net worth and next-period productivity be $V_t(n_t, z_{t+1})$, where the time subscript for the overall function allows for the possibility of a non-recursive path for the aggregate variables. It is easy to show that value in this framework is linear in net worth, and so I will impose this from the start, letting $V_t(n_t, z_{t+1}) = v_t(z_{t+1})n_t$, where $v_t(z_{t+1})$ can be interpreted as value per unit of net worth. A full derivation can be found in the

[14]This assumption is useful for the analytical results. Relaxing it in versions of the model which are solved numerically does not qualitatively change the results.

appendix. $v_t(z_{t+1})$ can be expressed recursively as:

$$v_t(z_{t+1}) = \max_{0 \leq \phi_t \leq 1+\lambda_t} \left\{ \mathrm{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} (1 - \sigma + \sigma v_{t+1}(z_{t+2})) \right] \times \right.$$

$$\left. ((\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t)\phi_t + r_t) \right\} \tag{2.17}$$

The expression inside the expectation gives the expected value of a unit of net worth next period. Firms exit with probability $(1-\sigma)$, in which case a unit of net worth is paid out as dividends and is worth one. With probability σ the firm remains active, and a unit of net worth is worth $v_{t+1}(z_{t+2})$. The term afterwards is your return on net worth, conditional on a choice of ϕ_t. Importantly, this is known at time t and hence out of the expectation. This allows for a very simple solution to the firm's problem. As long as the value of being a firm is positive, the optimal leverage policy is "bang-bang":

$$\phi_t = \begin{cases} 1 + \lambda_t & : \pi(z_{t+1}, w_{t+1}) + 1 - \delta \geq r_t \\ 0 & : \pi(z_{t+1}, w_{t+1}) + 1 - \delta < r_t \end{cases} \tag{2.18}$$

Intuitively, if a firm is sufficiently productive tomorrow it will want to invest. Due to the linearity of profits, it will want to invest as much as possible, and will hit its borrowing constraint, leading to leverage of $1 + \lambda_t$. If the firm is not productive enough it won't want to invest, and will invest zero in capital and simply carry its net worth over using bonds. This defines a threshold productivity $\underline{z}_{t+1} = \underline{z}(w_{t+1}, r_t)$ above which firms invest and below which they don't, given by:

$$\pi(\underline{z}(w_{t+1}, r_t), w_{t+1}) + 1 - \delta = r_t \tag{2.19}$$

Importantly for my result, this threshold productivity is increasing in the wage ($\underline{z}_1(w_{t+1}, r_t) > 0$), which follows from the definition of the profit function. This is intuitive: a higher wage increases the wage bill, reducing the profit per unit of capital.

This increases the minimum productivity required to make investment profitable.

2.5.3 Aggregating the firm sector

To be able to easily aggregate the firm sector it is necessary to assume that firms' productivities are i.i.d. over time. I hence assume that productivities are distributed according to a CFD $F(z)$. I also restrict myself to looking at cases where the wage is either at or above the market-clearing wage, and do not consider parameters or shocks which would lead to a wage stuck below the market-clearing level. This means that firms are always on their labour demand curves. In other words, I assume that in the model with sticky wages it is always the household who is rationed, and never firms. Thus I do not distinguish between equilibrium hours, L_t, and labour demand, L_t^d. With these assumptions I can express the evolution of total firm net worth, N_t, as:

$$N_{t+1} = (1-\sigma)\Big[(1 - F(\underset{\sim}{z}(w_{t+1}, r_t))) \Big((1+\lambda_t) \Big(a w_{t+1}^{1-\frac{1}{\alpha}} \mathrm{E}_t \Big[z_{t+1}^{\frac{1}{\alpha}} | z_{t+1} \geq \underset{\sim}{z}(w_{t+1}, r_t)\Big] + 1 - \delta\Big) - \lambda_t r_t\Big) + F(\underset{\sim}{z}(w_{t+1}, r_t)) r_t\Big] N_t + w_e \tag{2.20}$$

The term proceeded by $(1-\sigma)$ is the net worth carried over by surviving firms. w_e is the total exogenous equity injection given to new firms. I assume that a mass $(1-\sigma)$ of new firms are created each period to keep the total mass constant, each receiving equity $w_e/(1-\sigma)$. The constant a is defined by:

$$a \equiv \Big((1-\alpha)^{\frac{1}{\alpha}-1} - (1-\alpha)^{\frac{1}{\alpha}}\Big)$$

In the appendix, I show that it is possible to aggregate the firm sector up to a representative firm. Defining Y_t, L_t and K_t as total firm output, labour demand and capital demand respectively, we get an aggregate production function:

$$Y_t = e^{\tau_t^e} K_{t-1}^{\alpha} L_t^{1-\alpha} \tag{2.21}$$

τ_t^e is the efficiency wedge, which is endogenous in the model. I discuss the precise formula for τ_t^e in the next section. The labour demand decisions also aggregate up to a representative firm's first order condition:

$$(1-\alpha)Y_t = w_t L_t \tag{2.22}$$

Finally, total capital demand is given by:

$$K_t = (1+\lambda_t)(1 - F(\underline{z}(w_{t+1}, r_t)))N_t \tag{2.23}$$

This is the leverage of investing firms, $(1+\lambda_t)$, multiplied by their total net worth, which is a fraction $(1 - F(\underline{z}(w_{t+1}, r_t)))$ of total net worth, N_t.

2.5.4 Wedges

The wedges are defined in an analogous way to the empirical definitions in Section 2.4. The efficiency wedge is defined by the production function, (2.21). By combining the individual firm policy functions I show in the appendix that the efficiency wedge is given by:

$$\tau_t^e = \log \mathrm{E}\left[z_t^{\frac{1}{\alpha}} | z_t \geq \underline{z}(w_t, r_{t-1})\right]^{\alpha} \tag{2.24}$$

Quite intuitively, TFP (the exponential of the efficiency wedge) is approximately the average productivity of firms who decided to invest and produce, up to a Jensen's inequality term. Hence the model is able to generate endogenous movements in the efficiency wedge in response to a financial shock, as long as this affects the distribution of firms who are producing. (2.24) defines a relationship $\tau_t^e = \tau^e(\underline{z}_t)$, which it is easy to show satisfies $\tau^{e\prime}(\underline{z}) > 0$. Intuitively, the higher the productivity threshold, the higher the average productivity of producing firms. The threshold for deciding to invest, $\underline{z}(w_t, r_{t-1})$ depends on the wage and real interest rate, and so it is via changes in these prices that the model is able to generate movements in TFP. As in the empirical exercise, the labour wedge is defined by:

$$v'(L_t)c_t = (1-\tau_t^l)(1-\alpha)\frac{Y_t}{L_t} \tag{2.25}$$

Note that I have implicitly assumed that the labour wedge is measured using the same utility function as that of the model. This, of course, means that if we impose flexible wages and labour market clearing the labour wedge will trivially be zero. With sticky wages, however, we have a chance to generate a labour wedge.

2.5.5 Market clearing

Goods market clearing requires that output is either invested or consumed:

$$c_t + K_t = Y_t + (1-\delta)K_{t-1} \tag{2.26}$$

Intuitively, we can think of this equation as determining consumption and hence pinning down the real interest rate. The labour market market may or may not clear, depending on the whether I assume flexible or sticky wages. If I assume flexible wages, then labour market clearing requires that:

$$L_t^s = L_t \tag{2.27}$$

We can then roughly think of the household's labour first order condition pinning down wages. If instead I assume sticky wages, then I simply impose an exogenous sequence of wages, $\{w_t\}_{t=0}^{\infty}$. Since I restrict myself to cases of rationing on the side of unemployment, I thus dispense with the household's labour first order condition in this case, and define unemployment as:

$$U_t = L_t^{s*} - L_t \tag{2.28}$$

Where L_t^{s*} is the hypothetical amount of labour that the household would like to supply if it could, which is backed out from the household labour first order condition, (2.7), using equilibrium consumption and the wage.

2.5.6 Results

To gain intuition, I turn to analytical results from the model's steady state. My goal is to show how a financial shock is translated into wedges and unemployment, depending on whether wages are flexible or sticky. Variables with a time subscript denote steady state values.

It is helpful to define steady state labour supply and demand at this point. The steady state labour demand curve is defined for a given wage as $L^d(w;\lambda)$. The dependence on λ is also made explicit as this is what I will be varying in my comparative statics. Labour demand is found by combining the aggregated firm equations (2.20), (2.21), (2.22), and (2.23), using the steady state interest rate $r = 1/\beta$ from (2.6). The steady state labour supply curve is defined as $L^s(w;\lambda)$. This is given by the household's labour supply equation, (2.7). This equation contains consumption, which found from the resource constraint, (2.26) using the values for output and capital solved for during the derivation of labour demand. The following lemma establishes properties of steady state labour demand and supply:

Lemma 1. *Labour demand is downwards sloping in the wage* ($L_1^d(w;\lambda) < 0$) *and increasing in the borrowing limit* ($L_2^d(w;\lambda) > 0$)*. Labour supply is decreasing in the borrowing limit* ($L_2^s(w;\lambda) < 0$)*.*

The proofs are relagated to the appendix, and I discuss the intuition here. Labour demand is downwards sloping in the wage because high wages encourage substitution towards capital. Labour demand increases as the borrowing constraint is relaxed because this increases firms' net worth, increasing their steady state capital, and hence marginal product of labour and labour demand. On the other hand, labour supply is decreasing as the borrowing limit is relaxed because a looser borrowing constraint leads to higher output and consumption in equilibrium. Higher consumption leads to lower labour supply, at a given wage, via (2.7). This can be interpreted as an income effect: the looser borrowing constraint raises household income, so households demand more leisure and hence supply less labour. It is harder to prove that steady state labour supply is upwards sloping in the wage ($L_1^s(w;\lambda) > 0$), and I impose this as a weak regularity assumption, which should be considered a maintained assumption in the following propositions.[15]

Two other objects are useful. Firstly, define $\underline{z}_{ss}(w) \equiv \underline{z}(w, 1/\beta)$. This is the productivity threshold in steady state, which now depends only on the wage since the real interest rate is fixed at $1/\beta$. It is simple to prove that the productivity threshold is increasing in the wage ($\underline{z}'_{ss}(w) > 0$). Higher wages lead to lower profit per unit of investment, increasing the productivity required to break even. The second object is the relationship between TFP and threshold productivity, $\tau^e = \tau^e(\underline{z})$, defined by (2.24). It is worth noting at this point that neither relationship, $\underline{z}_{ss}(w)$ or $\tau^e(\underline{z})$, depend directly on λ, so tightening the borrowing constraint will only affect productivity indirectly through changes to the real wage. These relationships, along with the labour demand and supply curves, are plotted in Figure 2.6.

I present my results in three propositions. The first two consider the effect of a tightening of the borrowing constraint in steady state in the extreme cases of fully flexible and fully rigid wages. The third then considers intermediate wage adjustment. I provide formal proofs in the appendix, and more informal discussions below. I start

[15]Labour supply could be downwards sloping in the wage if consumption increased enough in response to an increase in the wage. This could be possible, since increasing the wage reduces firm profits, net worth, and hence ability to invest. From the resource constraint, $c = Y - \delta K$, so if δK decreases by more than Y consumption could increase. This does not happen in reasonable calibrations of the model.

with the case of fully flexible wages:

Proposition 8. *If wages are fully flexible, then in steady state a credit crunch (a reduction in λ from λ_0 to λ_1) leads to:*

1. *A fall in the efficiency wedge, τ^e*
2. *No change in the labour wedge, τ^l, which remains at zero*
3. *No change in unemployment, U, which remains at zero*

Proposition 8 is the counterpart to Buera & Moll's result in my model: a credit crunch manifests itself as a fall in TFP. Due to the assumed flexibility of the wage rate, the labour wedge and unemployment remain at zero because the labour market clears. The intuition behind the fall in TFP deserves some explanation, especially as it will highlight the differences between the flexible and sticky wage cases. Figure 2.6 gives a graphical proof of the fall in TFP. The left hand panel plots the inverse of labour supply and demand, with solid lines corresponding to the original λ_0, and the dashed lines the credit crunch ($\lambda_1 < \lambda_0$). The credit crunch reduces labour demand and increases labour supply, leading to an unambiguous fall in the real wage from w_0 to w_1.

The remaining panels trace the transmission from the real wage to TFP. The top right panel plots the relationship between the wage and threshold productivity, $\underline{z}_{ss}(w)$. The fall in wages from w_0 to w_1 increases profits, reducing the productivity threshold from $\underline{z}_0$ to $\underline{z}_1$ and encouraging less productive firms to start investing. This leads to a composition effect which endogenously reduces TFP, as detailed in the bottom right panel which plots $\tau^e(\underline{z})$. The reduction in the productivity threshold for investment leads to a fall in the average quality of producing firms, and hence a fall in measured TFP from τ^e_0 to τ^e_1.

I now turn to the case of fully rigid wages. For simplicity, I assume that the wage is fixed at a level which leads to full employment before the credit crunch in

Figure 2.6: Credit crunch in steady state

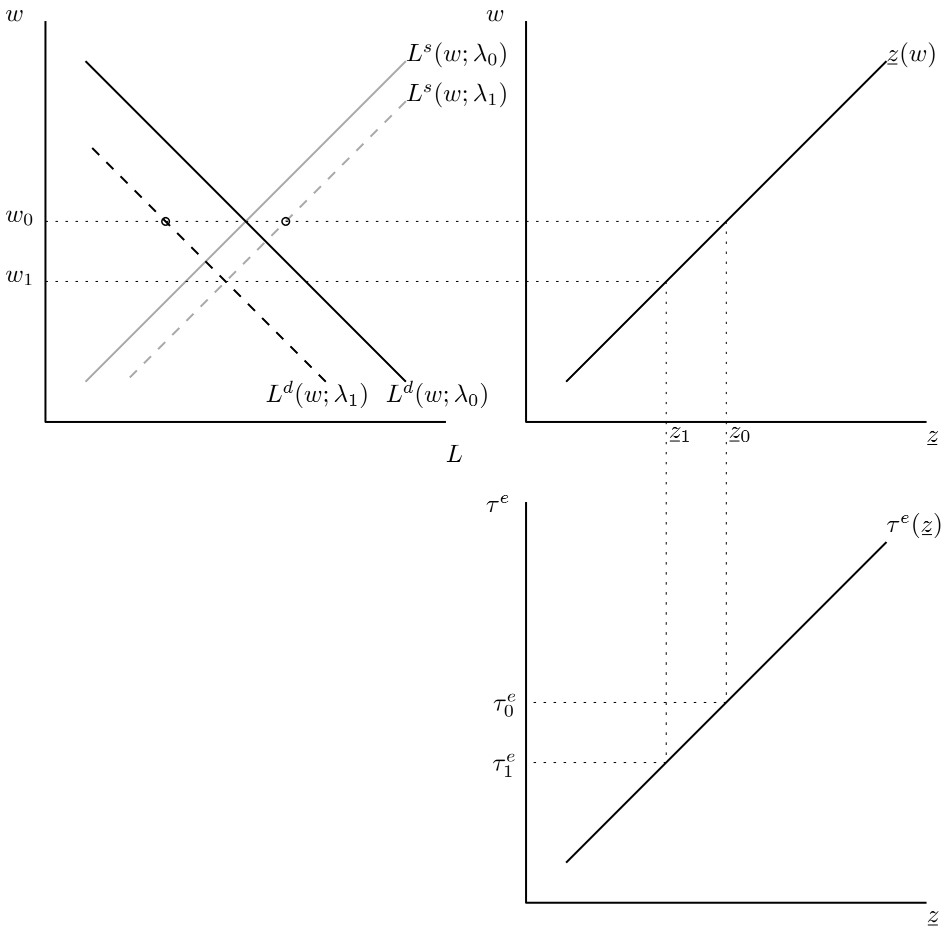

Model steady state for two values of the borrowing constraint, λ. Tightening the borrowing constraint leads to lower wages (top left panel), lower threshold productivity (top right panel), and lower measured TFP (bottom right panel).

the graphical presentation, but this is not important for the results. I can prove the following proposition:

Proposition 9. *If wages are fully rigid at w_0, then in steady state a credit crunch (a reduction in λ from λ_0 to λ_1) leads to:*

1. *No change in the efficiency wedge, τ^e*

2. *An increase in the labour wedge, τ^l*

3. An increase in unemployment, U

With the wage fully rigid at w_0, we get no change in the efficiency wedge, i.e. no fall in TFP. This follows directly from the fixed wage: since the wage does not move, there is no change in profitability (*per unit invested*) and hence no change in the productivity threshold, which remains at $\underline{z}_0$, or TFP, which remains at τ^e_0. Intuitively, the credit crunch was transmitted to TFP *via* the wage fall encouraging less productive firms to start investing, and this channel is shut off if wages are rigid.

With sticky wages the credit crunch now manifests itself as problems with labour markets. First, it increases unemployment. This is demonstrated graphically by noting that the shifts in the labour demand and supply curves induced by the reduction in λ lead to a wage which is now above market clearing level. This leads to unemployment, equal to the horizontal distance between $L^s(w_0; \lambda_1)$ and $L^d(w_0; \lambda_1)$. Intuitively, the reduction in firms' borrowing capacity reduces their investment in capital, which reduces the marginal product of labour, reducing their demand for labour. At the same time, the reduction in consumption caused by the reduction in output leads to an increase in labour supply. At the fixed wage we are rationed, with equilibrium hours falling by the fall in labour demand.

This manifests itself as an increase in the labour wedge, as defined by (2.25). To see this, note that since firms remain on their labour demand curves, $w = (1-\alpha)Y/L$ giving:

$$\tau^l = 1 - \frac{v'(L)c}{w_0}$$

The reductions in c and L (which reduces $v'(L)$) lead to an increase in τ^l. Intuitively, since households are rationed from working and would like to work more, we need an increase in the hypothetical labour tax to justify why hours are so low given the higher marginal product of labour.

Having established the response of the economy individually under the two extreme wage assumptions, I now make a comparative statement, which can be used to interpret the international data I presented.

Proposition 10. *Following a credit crunch* (a *reduction in* λ *from* λ_0 *to* λ_1) *in steady state which leads to rationing unemployment, the larger the fall in real wages:*

- *The larger the fall in TFP,* τ^e
- *The smaller the fall in hours worked,* L

Since I only consider economies with full employment or rationing unemployment, we will always be on the labour demand curve. Since this curve is downwards sloping, any fall in real wages will increase hours relative to the initial fall caused by the reduction in λ, reducing the total fall in hours. Additionally, any fall in wages reduces TFP by reducing $\underline{z}$ and hence τ^e, as discussed above.

These results together give us a useful lens with which to make sense of the data I presented at the beginning of the paper. The first two propositions consider the extreme cases of completely flexible and rigid wages. I show through these examples the importance of wage adjustment in determining how a credit crunch manifests itself in the real economy. If wages are fully flexible then we see a fall in the efficiency wedge (TFP) and no distortions in the labour market, meaning no labour wedge or unemployment. On the other hand, if wages are fully rigid then we see no change in the efficiency wedge, and an increase in the labour wedge and unemployment. These results thus help explain the behaviour of wedges in my two prototype examples, the US and the UK during the Great Recession, and set the stage for the larger cross country comparison.

The third proposition says that, for a given size of financial crisis, an economy whose wages adjust (downwards) more will experience a less severe fall in hours, and a more severe fall in TFP. This proposition is thus consistent with the cross country evidence I presented for the Great Recession. The model predicts the correct correlations between TFP, hours, and real wages, and I am thus able to explain *jointly* the behaviour of productivity and labour markets during the Great Recession by exploiting only differences in wage adjustment.

2.6 Conclusion

In this paper, I document a new stylised fact: countries which experienced larger declines in Total Factor Productivity during the Great Recession experienced less severe falls in hours worked. I also show that countries with larger falls in real wages during the crisis tend to be those with TFP, and not labour market, problems. Motivated by this second fact, I offer a parsimonious explanation of the negative TFP-hours correlation using a model of firm heterogeneity and differential wage adjustment in response to a financial shock. My model is motivated by cross-country data during the recent crisis from the OECD, and a more detailed case-study analysis of the US and UK. My motivating evidence can be summarised as four correlations: During the crisis: a) TFP and hours are negatively correlated, b) TFP and real wages are positively correlated, c) hours and real wages are negatively correlated, and d) real wages and the price level are negatively correlated.

I present a model which can explain the negative TFP-hours correlation by appealing to the other three correlations. For a given level of wages, a financial crisis reduces the ability of firms to borrow to fund investment, reducing the marginal product of labour, and hence labour demand. If wages adjust relatively little, this will lead to large falls in hours worked in equilibrium. If wages adjust downwards by a lot this offsets the fall in labour demand, leading to smaller falls in hours. However, lower wages also shield firms from having to shut down or downsize in response to the financial crisis, which leads to a worsening allocation of resources and hence lower measured TFP. In the model, this comes from a composition effect, based on Buera & Moll (forthcoming). In particular, the model features firms who are heterogeneous in their productivities, and who decide whether or not to produce based on their profitability. Firms with the lowest productivity levels choose not to produce, leading to measured TFP being endogenously determined by the set of firms who are producing. In the model, a fall in the real wage increases profitability for all firms, leading unproductive firms to start producing and reducing measured TFP. The extent to which this happens depends on the extent of wage adjustment during the crisis, with larger falls in wages leading to larger falls in measured TFP. Differential wage adjustment thus generates a negative TFP-hours correlation in my model following a financial crisis, consistent with the data.

For the US and UK, I construct business cycle wedges following Chari, Kehoe & McGratten's (2007) accounting procedure. This exercise shows that the US' recession can be explained mostly through the labour wedge, and the UK's recession through the efficiency wedge. I show that my model is also consistent with this evidence: following a financial shock, the model generates only a labour wedge and no efficiency wedge if wages are fully rigid, and only an efficiency wedge and no labour wedge if wages are fully flexible. The model is thus also able to jointly rationalise the behaviour of the efficiency and labour wedges in the US and UK during the crisis by again appealing to differential wage adjustment.

While financial crises may seem quite different, Reinhart & Rogoff (2009) taught us to look for commonalities under the surface. This paper attempts to make a small contribution in this direction, by identifying a new pattern to how financial crises manifest themselves in the real economy across countries. While the focus here has been on explaining events during the Great Recession, future work should investigate whether this pattern holds in other historical episodes. Additionally, the results of this paper suggest a new tradeoff of monetary policy following a crisis. Attempts to stimulate employment by increasing inflation may come at the expense of lower productivity.

Appendices

2.A Data sources

2.A.1 OECD

Most of the cross-country data are from the OECD dataset "Growth in GDP per capital, productivity and ULC", which I denote OECD1, and is available at `https://stats.oecd.org/Index.aspx?DataSetCode=PDB_GR` at the time of writing. Selected variables for certain countries have been replaced with supplementary series in order to fill in missing data, and I document any changes below. All data are yearly.

- TFP: "Multifactor productivity" series from OECD1.
- Hours: Hours per capita is calculated as "Total hours worked" from OECD1 divided by "Total population; persons; thousands" from the OECD dataset "Level of GDP per capita and productivity", which is available at `https://stats.oecd.org/Index.aspx?DataSetCode=PDB_LV` at the time of writing.
- Wages: The nominal wage is taken as "Labour compensation per hour worked" from OECD1.
  - The wage series for New Zealand is replaced by the wage calculated from preliminary quarterly estimates to extend the series. Specifically, I take Labour Compensation per Employed Person from the OECD dataset "Unit Labout Costs and labour productivity (employment based), total economy" (available at `http://stats.oecd.org/Index.aspx?DataSetCode=ULC_EEQ` at the time of writing) which is then converted to an hourly wage using data on employment and hours from OECD1.
  - The wage series for Australia is replaced using Australian national accounts data. Specifically, I take Compensation of Employees from the national accounts (available at `http://www.abs.gov.au/` at the time of writing) and divide by hours from OECD1 to create a wage series.

- Prices: The price level is taken as the GDP deflator, which is calculated from OECD1 as real GDP ("GDP, constant prices") over nominal GDP ("Gross Domestic Product (GDP); millions").

- Credit intermediation ratio: Taken from the OECD dataset "Financial Indicators – Stocks", which is available at `https://stats.oecd.org/Index.aspx?DataSetCode=FIN_IND_FBS` at the time of writing.

2.A.2 US

Notes: All data are seasonally adjusted. The quarterly national accounts data are presented in yearly rates, and are thus divided by 4 to get quarterly values. The labour series is normalised to have average 1/3 over the sample. Consumption and investment are deflated by the GDP deflator and not their individual deflators. This is standard in the RBC model, since the model does not allow for movements in the relative prices of output, consumption or investment.

- GDP (Y_t): Chained value taken from line 1 of NIPA table 1.1.6. Deflator taken from line 1 of NIPA table 1.1.4.

- Consumption (C_t): non-durables plus services. Nominal, then deflated by gdp deflator. Taken from line 5 and 6 of NIPA table 1.1.6.

- Investment (X_t): Gross domestic private investment (NIPA 1.1.6 line 7) + durable consumption (NIPA 1.1.6 line 4). Both nominal, deflated by GDP deflator. Note that I treat investment from date t in the table as not operational until $t + 1$.

- Initial capital stock (K_0): Year end capital stock constructed from BEA Fixed Asset table 1.1 (yearly data). To match investment data, use Private Fixed Assets (line 3) + Consumer Durables (line 15). These are current cost measures, and I deflate by the yearly GDP deflator. My initial capital stock is the year end value the year before my investment data begins.

- Depreciation (δ): Current cost depreciation data from Fixed Asset table 1.3. Computed for private fixed assets (line 3) + consumer durables (line 15) and again deflated by the GDP deflator. Yearly depreciation rates are computed as $\delta_{y,t} = Dep_t/K_{t-1}$. Implied quarterly depreciation rates are computed as the solution to $\delta_{q,t} + \delta_{q,t}(1-\delta_{q,t}) + \delta_{q,t}(1-\delta_{q,t})^2 + \delta_{q,t}(1-\delta_{q,t})^3 = \delta_{y,t}$. I compute the depreciation rate as the average of these quarterly depreciation rates.

- Capital stock (K_t): Constructed using the perpetual inventory method starting with the initial capital stock and using $K_t = X_t + (1-\delta)K_{t-1}$.

- Hours worked (L_t): Hours worked is the series "Nonfarm Business Sector: Hours of All Persons, Index 2009=100, Quarterly, Seasonally Adjusted", series HOANBS downloaded from FRED.

- Population (N_t): "Civilian Noninstitutional Population, Thousands of Persons, Monthly, Not Seasonally Adjusted" available from the Federal Reserve Economic Data (FRED). I take every third datapoint to construct quarterly data, and then take a one year moving average to deseasonalise.

- Labour share (LS_t): The data for the labour share come from Gross Domestic Income data, since income breakdowns are not available for the GDP data. Since there are small discrepancies between the GDI and GDP data, instead of taking the wage directly from the GDI data, I simply compute the labour share from this data. I can then use the labour share (under the assumption that it is equal in the GDI and GDP data) to back out the implied wage consistent with the GDP data. The nominal GDI data is from NIPA table 1.10 and are constructed along the lines of Karabarbounis (2014). I first construct the unadjusted labour share, LS_t^u as the share of unambiguous labour income to

GDI. This is compensation of employees paid (line 2) divided by GDI (line 1). I then attribute the fraction LS_t^u of ambiguous income to labour. Ambiguous income is "Proprietors' income with inventory valuation and capital consumption adjustments" (line 13) + "Taxes on production and imports" (line 7) - "Less: Subsidies" (line 8). The final labour share is then computed as unambiguous labour income plus the fraction LS_t^u of ambiguous income all divided by GDI.

- Real Wages (w_t): The real wage is calculated as $w_t = LS_t y_t / l_t$

2.A.3 UK

The available UK data which is comparable to the US data is limited, and starts from 1997Q1. In particular, a longer time series for capital is available for the UK, but it does not distinguish between government capital and private capital. In the wedges exercise for the US, following Chari, Kehoe & McGratten (2007), capital is defined as only private capital. For this reason I am restricted to data for the UK from 1997, when a breakdown of the capital stock was first released. Most data are from the UK quarterly national accounts, unless otherwise stated.

- GDP (Y_t): The GDP deflator is series "Implied Deflators: Gross domestic product at market prices" (series YBHA) from table "A1: National Accounts Aggregates". Nominal GDP is the series "Current prices: Gross domestic product at market prices" (series YBGB) from table "A2: National Accounts Aggregates". Real GDP is calculated as nominal deflated by the price deflator.

- Consumption (C_t): non-durables, semi-durables and services. Nominal, then deflated by GDP deflator. Taken from series UTIR, UTIJ and UTIN of table "E2: Household final consumption expenditure (goods and serives) at current prices".

- Investment (X_t): Business investment + Private Sector: Dwellings + Private Sector: Costs of Ownership + Durable Consumption. Investment data from table "G9: Gross Fixed Capital Formation by sector and type of asset". Series: NPEK, L62T and L62U. Durable consumption data from table E2, series UTIB. Nominal, then deflated by GDP deflator.

- Initial capital stock (K_0): Unfortunately, data on the stock of the specific sub-types of investment which I use for my capital series are not available for 1996Q4. I thus experiment with various initial stocks, and the results for the Great Recession episode are not sensitive to the initial choice. The initial choice I settle on is the one that leads to a path for the Solow residual which looks closest to being one with stationary fluctuations around a constant growth rate.

- Depreciation (δ): Again, individual depreciation rates are not available for the sub-types of capital. Since the sub types were chosen to be similar to those for the US, I take the same depreciation rate as the US.

- Capital stock (K_t): Constructed using the perpetual inventory method starting with the initial capital stock and using $K_t = X_t + (1-\delta)K_{t-1}$.

- Hours worked (L_t): Hours worked is the series "HOUR01 Actual weekly hours of work", seasonally adjusted, available from the ONS.

- Population (N_t): Annual "Mid-year population estimates and annual change for the UK mid-1964 onwards" data are interpolated to quarterly data by assuming a constant within-year quarterly growth rate. Data available from the ONS.

- Labour share (LS_t): Data are constructed as for the US, using the table "D: Gross Domestic Product: by category of income". Unambiguous labour income is taken as "UK Total compensation of employees" (DTWM). Ambiguous income is taken as "Other income" (CGBX). Total income is "Income based GVA at factor cost" (CGCB).

- Real Wages (w_t): The real wage is calculated as $w_t = LS_t y_t / l_t$

2.B Cross-country robustness: Outliers

In this section I repeat my cross-country exercise including Ireland, which had been dropped as an outlier. Figure 2.7 repeats the simple cross-section plot (Figure 2.1) with this extra data point. Ireland appears very clearly as an outlier in the top left panel, which plots the TFP-hours relationship. Using the other panels we can ascertain whether Ireland being an outlier in the TFP-hours space is more to do with it having extreme TFP or hours behaviour. The top right panel reveals that Ireland does not appear to be too much of an outlier in the TFP-wage space, compared to the variation of other countries. However, the bottom left panel reveals that Ireland's hours appear unusually low compared to the line of best fit given its wages. Hence perhaps it is something unusual about Ireland's labour market experience during the recession that leads it to be an outlier here.

I also repeat the results of the partial correlations, including Ireland, in Table 2.3. The TFP-hours relationship is still negative, but is now smaller and insignificant. As suggested by the simple correlations above, the TFP-wage relationship is unaffected by the inclusion of Ireland, while the hours-wage correlation is now insignificant.

Table 2.3: Partial correlations

	TFP	l	w	P
TFP	1	-0.0653	0.5316***	-0.1964
l		1	0.1440	0.5833***
w			1	-0.0671
P				1

, **, and * denote significance at the 10%, 5% and 1% level respectively.*

Figure 2.7: Relationships between selected variables across countries.

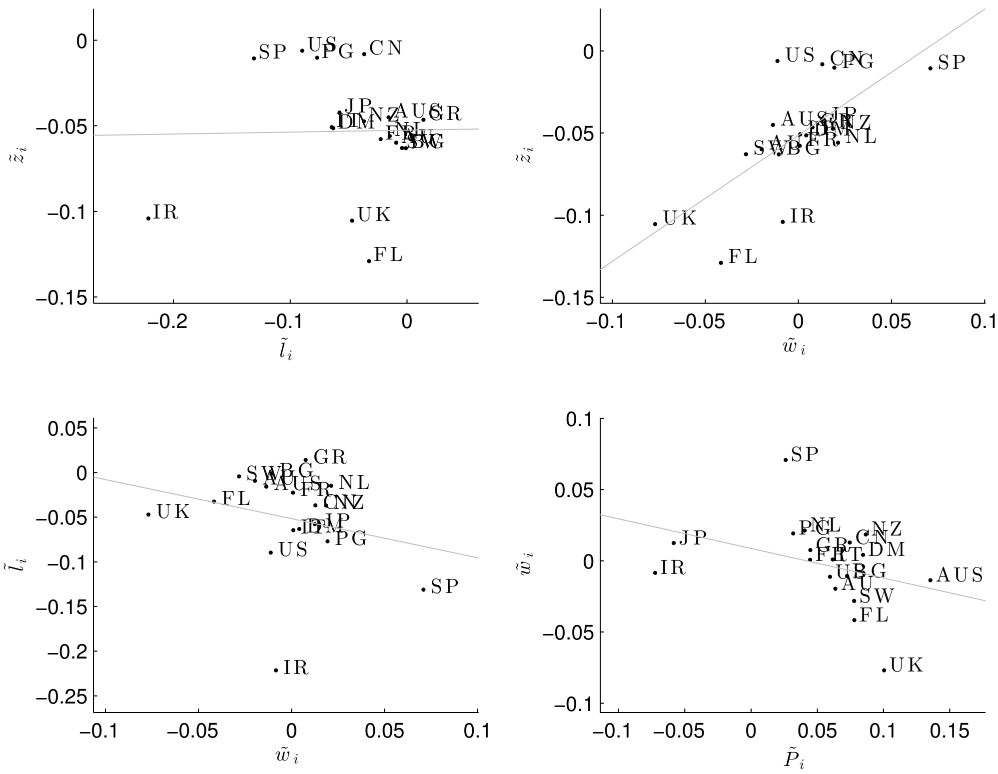

Lines are OLS lines of best fit between the two variables. The country names refer to: AUS = Australia, AU = Austria, BG = Belgium, CN = Canada, DM = Denmark, FL = Finland, FR = France, GR = Germany, IT = Italy, JP = Japan, NL = Netherlands, NZ = New Zealand, PG = Portugal, SP = Spain, SW = Sweden, UK = United Kingdom, US = United States.

2.C Business Cycle Accounting procedure

The prototype economy is relatively simple. There is a representative household with growing population N_t. It has log utility over per-capita consumption, c_t, and convex disutility over hours worked per capita, l_t:

$$U(c_t, l_t) = \log(c_t) - v(l_t) \tag{2.29}$$

The household discounts the future with the discount factor β. The household's budget constraint is:

$$c_t + (1 + g_{N,t})k_t - (1 - \delta)k_{t-1} = (1 - \tau^l_t)w_t l_t + (1 - \tau^x_{t-1})r_t k_{t-1} \tag{2.30}$$

Where $k_{t-1} \equiv K_{t-1}/N_t$ is per-capita capital which is productive at time t. w_t is the hourly wage, and capital depreciates at rate δ. $g_{N,t}$ is the population growth rate between t and $t + 1$. τ^l_t and τ^x_t are percentage taxes on labour and capital income respectively. I assume that the tax on capital income is known at the time the relevant investment decision is made, namely one period in advance. There is a representative firm with Cobb-Douglas production function, $Y_t = Z_t K^\alpha_{t-1} L^{1-\alpha}_t$, whose static optimality conditions equate prices with marginal products. The equilibrium of the economy can be summarised by the following four equations:

$$\tilde{y}_t = e^{\tau^e_t}\tilde{k}^\alpha_{t-1} l^{1-\alpha}_t \tag{2.31}$$

$$v'(l_t)\tilde{c}_t = (1 - \tau^l_t)(1 - \alpha)\frac{\tilde{y}_t}{l_t} \tag{2.32}$$

$$1 = \frac{\beta}{1 + g_z}\mathrm{E}_t\left[\frac{\tilde{c}_t}{\tilde{c}_{t+1}}\left((1 - \tau^x_t)\alpha\frac{\tilde{y}_{t+1}}{\tilde{k}_t} + 1 - \delta\right)\right] \tag{2.33}$$

$$\tilde{c}_t + (1 + g_{N,t})(1 + g_z)\tilde{k}_t - (1 - \delta)\tilde{k}_{t-1} + \tau^g_t\tilde{y}_t = \tilde{y}_t \tag{2.34}$$

Lowercase variables with a tilde refer to per-capital variables which have been de-trended by the average growth rate of TFP, g_z.[16] Each of these equations corresponds to a different wedge, which we can measure using data on output, capital, hours and consumption. The first equation, the aggregate production function, identifies the "efficiency wedge", τ^e_t. This is the log deviation of measured TFP from trend. The second equation, the labour market optimality condition, identifies the "labour wedge", τ^l_t. This is our labour income tax from the model, and the data measures it as

[16] Z_t is measured TFP, and $e^{\tau^e_t}$ is thus the deviation of measured TFP from trend. For any variable x_t, $\tilde{x}_t \equiv x_t/(1 + g_z)^t$ except for capital, for which $\tilde{k}_{t-1} \equiv k_{t-1}/(1 + g_z)^t$.

the required labour tax in order to rationalise the observed level of hours. The third equation, the Euler equation, identifies the "investment wedge". This is our capital income tax, and the data measures it as the tax required to rationalise the level of investment. The final equation, the resource constraint, identifies the "government wedge", τ_t^g. This is measured from the data as the residual in the national accounts after subtracting consumption and investment from output, and hence actually contains both government spending and net exports. It is expressed for convenience as a fraction of output. Define the set of time t wedges as $\tau_t = \{\tau_t^e, \tau_t^l, \tau_t^x, \tau_t^g\}$.

2.C.1 Measuring the wedges

The wedges are measured using time series data on output, consumption, capital and hours and the four model equations. The efficiency, labour, and government wedge can all be backed out from the three static model equations (2.31), (2.32) and (2.34). The investment wedge is measured from the Euler equation, which contains an expectation term. Measuring it thus requires a model of how agents form their expectations. This is achieved by solving the model economy under the assumption of rational expectations. I assume that the wedges follow a VAR(1) process:

$$\tau_{t+1} = \Phi_0 + \Phi_1 \tau_t + Q\varepsilon_{t+1} \tag{2.35}$$

Where $\varepsilon_{t+1} \sim N(0, I_4)$. Denote the whole set of parameters by $\Phi = \{\Phi_0, \Phi_1, Q\}$. I give this process to the agents of the economy, and solve for equilibrium. The state variables for the economy are $s_t = \{\tilde{k}_{t-1}, \tau_t\}$. This leads to a decision rule for capital of the form:

$$\tilde{k}_t = g(\tilde{k}_{t-1}, \tau_t^e, \tau_t^l, \tau_t^x, \tau_t^g; \Phi) \tag{2.36}$$

This allows us to back out the investment wedge from (2.36) given data on $\tilde{k}_t$ and $\tilde{k}_{t-1}$ and the other wedges, which we previously calculated. The only complication here is that the capital policy function, and hence the measured value of the investment wedge, will depend on the assumed process for wedges, Φ. I thus use an iterative procedure to jointly estimate the investment wedge and Φ.[17] I use a linear approximation to the capital decision rule, which is solved for using Dynare.

2.C.2 Procedure to calculate the wedges

The process for wedges is estimated, while some other parameter values are calibrated to standard values. The model is quarterly, so I choose $\beta = 0.99$, and set $\alpha = 1/3$. Depreciation is estimated from the data, as detailed above. For the labour disutility function I choose $v(l) = l^{1+\eta}/(1+\eta)$, and set $\eta = 1$. The results are robust to a large range of values for η. My procedure to calculate the wedges differs from CKM's in a few ways, which I will discuss after outlining the procedure.

1. Back out the labour and government wedges from the data directly using the nonlinear equations (2.32) and (2.34). I also back out the average population growth rate g_N as the sample average of the quarterly growth rates n_{t+1}/n_t.

2. As an initial guess for the investment wedge series, I calculate it assuming perfect foresight. This is backed out from equation (2.33) ignoring the expectations operator.

3. Calculate the log Solow residual, $z_t \equiv \log y_t - \alpha \log k_{t-1} - (1-\alpha)\log l_t$. The efficiency wedge is the deviation from the trend growth in the Solow residual, and both are estimated jointly below.

4. Run a VAR(1) on the four estimated wedge series of the form:

$$\tau_{t+1} = c + \Phi\tau_t + Q\varepsilon_{t+1} \tag{2.37}$$

[17]This amounts to guessing a Φ, using this to calculate a series for τ_t^x, estimating a new Φ, calculating a new series for τ_t^x and so on until convergence. CKM instead use a likelihood procedure to jointly estimate the two.

where $\tau_t = [\tau_t^e, \tau_t^l, \tau_t^x, \tau_t^g]'$, ε_{t+1} is a four by one vector of independent errors with mean zero and unit variance. The elements of c and Φ are estimated by OLS on each equation, and Q as the lower triangular Cholesky decomposition of the estimated covarience matrix of the OLS errors. The regression for the efficiency wedge is run replacing the efficiency wedge with the Solow residual and including a time trend. This allows me to estimate the growth rate of TFP, and the efficiency wedge is then calculated as the deviation from this trend.[18]

5. Calculate the detrended variables $\tilde{y}_t$ etc by detrending by the estimated productivity growth rate: $\tilde{y}_t \equiv y_t/(1+g_z)^t$ and equivalently for all other variables except for $\tilde{k}_{t-1} \equiv k_{t-1}/(1+g_z)^t$.

6. Solve the detrended model using Dynare. There are eight equations to the model: the four detrended model equations, (2.33), (2.32), (2.31) and (2.34), and the four equations of the estimated VAR, (2.37).[19] This gives us the linear capital policy function:
$$\tilde{k}_t = c_0 + c_1\tilde{k}_{t-1} + c_2'\tau_t \tag{2.38}$$

7. Using the estimated policy function, (2.38), back out the implied investment wedge at time t as the value of τ_t^x which leads to the observed choice of $\tilde{k}_t$ given $\tilde{k}_{t-1}$ and the other three wedges.

8. Return to step 4, now running the VAR using the newly calculated investment wedge. Repeat until the calculated investment wedge series has converged.

Note that in this procedure the labour and government wedges are calculated only once in the beginning. During the iteration only the efficiency wedge (if I estimate

[18]I do this for the US. Unfortunately, given the shorter dataset this procedure is unstable for the UK, and there I estimate the productivity trend as the average growth rate during the sample.

[19]Note that in the solution the population growth rate in equation (2.34) is replaced with its sample average. An alternative would be to treat the actual population growth rate as a state variable and estimate a stochastic process for it as we do with the wedges.

the trend jointly with the model), investment wedge, estimated VAR coefficients and the capital policy function change.

2.C.3 Procedure to calculate counterfactual simulations

The counterfactual simulations all start with some initial capital stock, $\tilde{k}_0$, and a sequence of wedges $\{\tau_t\}_1^T$. In the case of the simulations with only one wedge active, I first pick a date range $t = t_0, ..., t_1$. I then take the initial detrended capital stock from the data, and the sequence of wedges also from the data. I switch off all of the wedges except one by setting all of their values to their first period values. The rest of the procedure is then to:

1. At time t, calculate $\tilde{k}_t$ using the linear policy function from Dynare, given the current state $\{\tilde{k}_{t-1}, \tau_t\}$.
2. Calculate the other three variables, $\tilde{c}_t$, $\tilde{y}_t$, and l_t, by solving equations (2.32), (2.31), and (2.34) given $\tilde{k}_t$ and the state. In practice, this is done by combining the three equations into one and solving numerically over l_t.

Notice that, by construction, this procedure will return the data for $\tilde{y}_t$, $\tilde{k}_t$, $\tilde{c}_t$ and l_t if the all of the original wedges are plugged in.

2.D Further discussion of the US vs UK

2.D.1 Decomposing the labour wedge

In this section I perform a decomposition of the labour wedge along the lines of Karabarbounis (2014), which is used to ask whether the labour wedge is caused mostly by distortions on the firm or consumer side. The idea is as follows. Suppose that instead of just considering a prototype economy where the consumer pays a labour income tax, we consider an economy where the consumer pays a labour income tax,

τ_t^{lc}, and the firm also pays a payroll tax, τ_t^{lf}. These are the decomposed labour wedges, which I'll refer to as the "consumer side" and "firm side" labour wedges, respectively. Then if we have data on real wages, the consumer and firm's labour optimality conditions can be rearranged to solve for these two wedges:

$$\tau_t^{lc} = 1 - \frac{v'(l_t)}{w_t u'(c_t)} \tag{2.39}$$

$$\tau_t^{lf} = (1-\alpha)\frac{y_t}{w_t l_t} - 1 \tag{2.40}$$

And it follows trivially from the definition of these wedges that they decompose the overall labour wedge in the following sense: $(1-\tau_t^l) = (1-\tau_t^{lc})/(1+\tau_t^{lf})$. This can be rearranged to construct the overall labour wedge from our two sub wedges:

$$\tau_t^l = 1 - \frac{1-\tau_t^{lc}}{1+\tau_t^{lf}} \tag{2.41}$$

The consumer-side labour wedge measures the wedge between the wage and the consumer's marginal rate of substitution between labour and consumption. It hence measures whether the representative consumer's labour supply is being distorted away from its optimal value. The firm-side labour wedge measures the wedge between the wage and the marginal product of labour. It hence measures whether the representative firm's labour demand is being distorted.

This decomposition thus allows us to see whether the labour wedge we see, especially the wedge in the US during the crisis, arises more from distortions on the firm or consumer side. To do this, I reconstruct a hypothetical labour wedge for the crisis for each country and each sub-wedge using (2.41), and allowing one sub-wedge to vary at a time while holding the other at its initial value. The results of this exercise are presented in the first column of Figure 2.8. For both the US and the UK we see that the cyclical movements in the labour wedge over the crisis are driven almost entirely by the consumer side labour wedge. This is also true of the simulated paths for output and labour: if I use these hypothetical wedges in the simulation exercise it is the consumer-side wedges which deliver movements closest to the data.

Figure 2.8: Decomposing the labour wedge

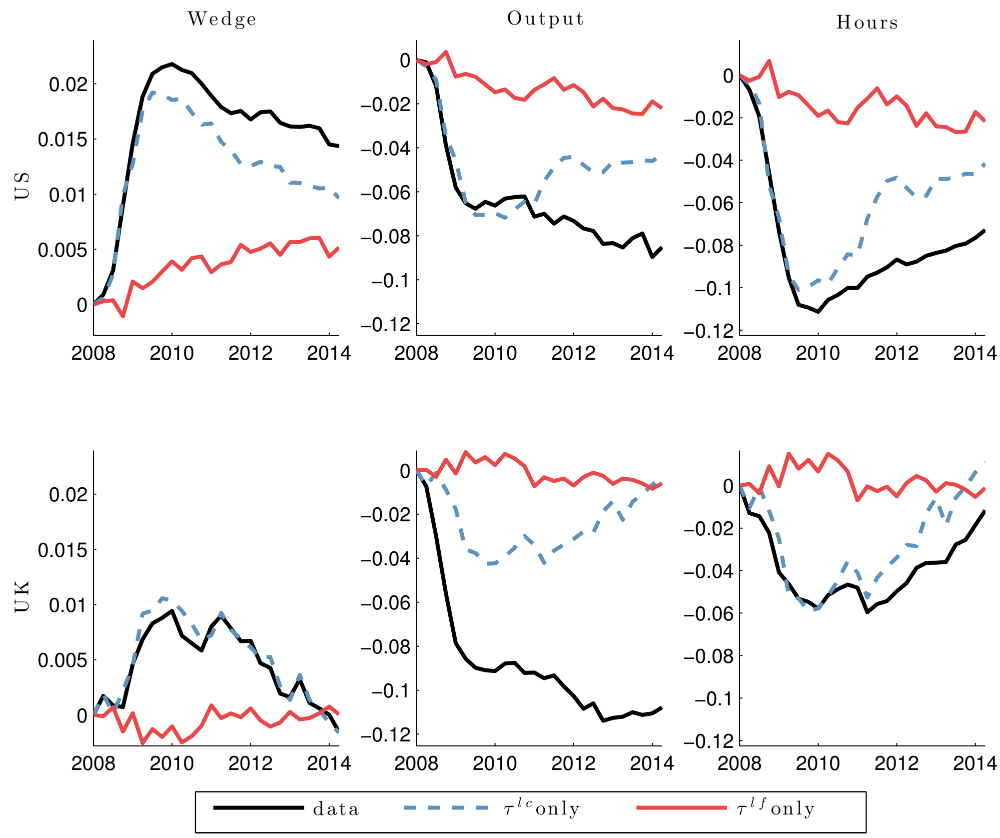

All variables are expressed as a fractional deviation from the value in the initial period. In the first column, "data" refers to the measured value of the overall labour wedge, τ_t^l, and the other two lines compute the value of the labour wedge would take using (2.41) if we allow one of the sub-wedges to vary, and hold the other at its first period value. In the other two columns, the counterfactuals for output and labour are computed using these counterfactual labour wedges. Output is detrended by the estimated trend growth rate for TFP.

In other words, we can think of firms as being roughly on their labour demand curves, and it is consumers whose labour supply decisions are being distorted. Given the sign of the labour wedge, this means that during the crisis we can think of consumers as wanting to work *more* given the current wage, and firms being happy with their employment levels given the current wage. This evidence is consistent with a model of rationing unemployment, where the wage is stuck above the market clearing level. In this scenario, firms are on their demand curves, and hire little. Consumers are rationed and work less than the desired amount, taking them off their supply curves and causing the large consumer-side labour wedge. In my model I assume sticky wages, leading to rationing unemployment of the type discussed above. This evidence is supportive of this assumption, and further implies that my model is able to match not only movements in the labour wedge, but also movements in the two sub wedges.

2.D.2 Controlling for composition effects

We should be concerned that changes in my wage data reflect composition changes, and not within-job wage changes. This could, for example, overemphasise the extent of the wage difference between the US and UK. If people with the lowest wages are fired first, this could put an upwards bias in US aggregate wages, making it look like they didn't fall when they actually did. To address this concern, I review the literature on composition effects in the US and UK during the Great Recession, and apply their corrections to my wage data.

In the UK, Blundell, Crawford & Jin (2014) find that, perhaps surprisingly, composition effects do not play a major role in the change in aggregate wages during the crisis. In particular, while composition effects always affect real wages, there was no change in the size of this effect pre vs post crisis. Hence the decline in real wages over the crisis period is almost entirely driven by falls in wages within jobs. In the US, Daly, Hobijn, & Wiles (2012) find that the size of the composition effect did change during the crisis. In particular, once controlling for composition effects they find average real wage growth of around 2% per year during the 2000s, before the crisis, and growth of around 1% from 2008-2011.

This data is not available from 2011 onwards, so I will restrict myself to looking at the 2008-2011 period. Over this period we need to adjust my detrended real wage growth estimates for the US downwards by around 1% per year. So the first two years of the crisis, where I report that US real wages increased by 2%, now have flat real wages. Over the whole 2008-2011 period, where I report that real wages increased by around 0.5%, corrected real wages now fall by around 2.5%. Overall, while this obviously reduces the gap between wages in the US and the UK, there still remains a sizable difference in the wage responses of the two countries, with the UK's wage response now being around double that of the US, despite the US' hours response being around twice as large.

2.D.3 Real unit labour costs

One might be concerned that wages are not the right measure of labour costs for my exercise. In particular, it is also common to look at "real unit labour cost" (RULC), which is defined as the wage bill divided by output: wL/Y. The argument for this measure is that it measures the labour cost of producing one unit of output, and hence measures labour costs better than the wage. In this section I argue that this is not the right measure for my exercise, and that wages are appropriate.

Firstly, in a simplified setting without composition effects from the labour side, wages are clearly the correct measure of labour costs for my model. Theoretically, the composition of firms is determined by wages, not RULC, as per the productivity threshold defined in (2.19). Additionally, RULC is a very endogenous object, even if wages are fully exogenous. To see this, note that in models with Cobb-Douglas production functions and competitively-priced inputs, RULC is always equal to $1-\alpha$ by firm optimisation. This is because, regardless of the wage, firms adjust hours until $MPL = w$, which leads to RULC equaling $1-\alpha$ in the Cobb-Douglas case. As I showed above, this is also true in my model where we have firm-side composition effects, as per (2.22). It is hard to think about how to interpret movements in a variable which should be constant at $1-\alpha$ in the baseline macroeconomic production function. Of course, RULC does move in the data. However, recent work has shown that these movements are relatively small compared to the size of fluctuations in the labour market. This is the message of the labour wedge decomposition of Karabarbounis (2014). The firm-side labour wedge actually measures movements in RULC, and its movements can explain essentially none of the movements in output or hours at business cycle frequencies. I confirm this result for the Great Recession in the US and UK in Appendix 2.D.1.

Secondly, even if we add labour-side composition effects it is not clear that RULC controls for these. This is the supposed benefit of RULC, since by measuring the total cost of labour per unit of output it controls for workforce composition. Before demonstrating that this claim may not be true, it is worth noting that there are other ways to control for composition effects more directly. This is what I do in my robustness for the US and UK in Appendix 2.D.2 using empirical estimates of the size of the relevant composition effects from worker-level data.

Theoretically, I demonstrate below that, in a simple model, we still recover that result that RULC is constant at $1-\alpha$ even with worker-side composition effects. Suppose there is only a single representative firm who has production function:

$$Y = zK^{\alpha}\left(\int_i z_i l_i di\right)^{1-\alpha} \tag{2.42}$$

This is a standard Cobb-Douglas production function, with heterogeneous labour inputs indexed by i. Labour type i can be thought of as labour of a certain skill level, which has productivity z_i. The firm can choose the number of hours of each skill type to hire, and I assume perfect substitutability between types. Suppose that the firm takes wages of each type, w_i, as given. The firm's FOC for choosing demand for labour from type i is:

$$w_i = (1-\alpha) z_i z K^{\alpha} \left(\int_i z_i l_i di \right)^{-\alpha} \tag{2.43}$$

Where this FOC holds for any labour type which the firm chooses to hire. Note that the perfect substitutes structure places a strong restriction on the relationship between wages of different types. Any labour types which are hired in equilibrium must satisfy $w_i/w_j = z_i/z_j$. In this simple setting I can generate composition effects by assuming that some labour types have wages which are too high relative to their productivity, leading to them not being hired. Define the measured average wage as:

$$w \equiv \frac{\int_{l_i>0} w_i l_i di}{L} \tag{2.44}$$

Where total labour is $L \equiv \int_{l_i>0} l_i di$. Multiplying both sides of the labour FOC by l_i and integrating over all types which are hired yields:

$$\frac{wL}{Y} = (1-\alpha) \tag{2.45}$$

Which says precisely that RULC are always equal to $1-\alpha$, regardless of which labour types are hired. The intuition for this result is exactly the same as that for the standard Cobb-Douglas case with one labour type: RULC are constant because the firm hires each type of labour up to the point where $MPL_i = w_i$. Of course, this model has the very restrictive assumption that labour types are perfect substitutes, and it remains to be seen to what degree RULC would vary in a model with more general substitutability between labour types.

2.E Model appendix

2.E.1 Derivation of Bellman

In this section I prove that the value function is linear in net worth. First, the Bellman for $V_t(n_t, z_{t+1})$ can be expressed as:

$$V_t(n_t, z_{t+1}) = \max_{0 \leq \phi_t \leq 1+\lambda_t} \left\{ \mathrm{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} \left((1-\sigma)n_{t+1} + \sigma V_{t+1}(n_{t+1}, z_{t+2})\right) \right] \right\} \quad (2.46)$$

Such that:

$$n_{t+1} = \left((\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t)\phi_t + r_t\right) n_t \quad (2.47)$$

Taking the first order condition with respect to ϕ_t yields:

$$\mathrm{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} \left(1 - \sigma + \sigma \frac{\partial V_{t+1}(n_{t+1}, z_{t+2})}{\partial n_{t+1}}\right) (\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t) \right] = \lambda_t - \mu_t \quad (2.48)$$

Where λ_t and μ_t are multipliers on the no shorting and borrowing constraints, scaled by n_t. Guessing that $V_t(n_t, z_{t+1}) = v_t(z_{t+1})n_t$ satisfies the Bellman and FOC above along with the policy function in (2.18).

2.E.2 Proof that firms never pay voluntary dividends

Proposition 11. *In the model, for parameters such that aggregate capital is positive in all periods, firms will never pay dividends until they exogenously exit.*

Proof. The proof imposes the linearity of the value function from the start, for simplicity of exposition. If I allow firms to pay dividends, the balance sheet becomes:

$$k_t + e_t n_t = n_t + d_t \quad (2.49)$$

Where $e_t n_t$ is total dividends, and e_t is thus dividends paid per unit of net worth. The assumption that firms cannot raise equity imposes the constraint that dividends must be positive: $e_t \geq 0$. Combining the balance sheet with the transition for net worth, (2.15), gives:

$$\frac{n_{t+1}}{n_t} = (\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t)\phi_t + r_t(1 - e_t) \tag{2.50}$$

The firm's value function is now given by:

$$v_t(z_{t+1}) = \max_{0 \leq \phi_t \leq 1+\lambda_t, e_t \geq 0} \left\{ e_t + \mathrm{E}_t \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} (1 - \sigma + \sigma v_{t+1}(z_{t+2})) \right] ((\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t)\phi_t + r_t \right. \tag{2.51}$$

Taking the first order condition with respect to e_t and using the equilibrium value of r_t from (2.6) reveals that $e_t = 0$ unless:

$$\mathrm{E}\left[v_{t+1}(z_{t+2})\right] \leq 1 \tag{2.52}$$

Note that the ability to pay dividends places a lower bound of one on $v_{t+1}(z_{t+2})$, which only occurs if the firm never makes profit on its investment (i.e. $(\pi(z_{t+1}, w_{t+1}) + 1 - \delta - r_t) \leq 0 \ \forall t$) and hence never invests in capital. However, since firm productivity is i.i.d. this cannot be the case in an equilibrium with positive aggregate capital: positive aggregate capital and i.i.d. shocks imply that any firm can expect to be profitable enough to be one of the investing firms at some point in the future. Hence we must have that $v_t(z_{t+1}) > 1$ for all t and z_{t+1}, implying that $\mathrm{E}\left[v_{t+1}(z_{t+2})\right] > 1$ and $e_t = 0$. $\square$

2.E.3 Aggregating the firm sector

Using firms' individual policy functions and aggregating the firms that produce gives:

$$Y_{t+1} = \int_{n_t} \int_{z_{t+1} \geq \underline{z}(w_{t+1}, r_t)} y_{t+1} f(z_{t+1}) f(n_t) dz_{t+1} dn_t = (1-\alpha)^{\frac{1}{\alpha}-1} w_{t+1}^{1-\frac{1}{\alpha}} (1+\lambda_t) Z(w_{t+1}) N_t \tag{2.53}$$

Where:

$$N_t \equiv \int_{n_t} n_t f(n_t) dn_t \tag{2.54}$$

$$Z(w_{t+1}) = \int_{z_{t+1} \geq \underline{z}(w_{t+1}, r_t)} z_{t+1}^{\frac{1}{\alpha}} f(z_{t+1}) dz_{t+1} \tag{2.55}$$

And I have used that productivity being i.i.d. over time makes n_t and z_{t+1} independent for any firm. Note that $Z'(w) < 0$. Aggregating capital and labour gives:

$$K_t = \int_{n_t} \int_{z_{t+1} \geq \underline{z}(w_{t+1}, r_t)} k_t f(z_{t+1}) f(n_t) dz_{t+1} dn_t = (1 + \lambda_t)(1 - F(\underline{z}(w_{t+1}, r_t))) N_t \tag{2.56}$$

$$L_{t+1} = \int_{n_t} \int_{z_{t+1} \geq \underline{z}(w_{t+1}, r_t)} l_{t+1} f(z_{t+1}) f(n_t) dz_{t+1} dn_t = (1 - \alpha)^{\frac{1}{\alpha}} w_{t+1}^{-\frac{1}{\alpha}} (1 + \lambda_t) Z(w_{t+1}) N_t \tag{2.57}$$

Combining these expressions implies the various aggregated expressions used in the main text. For example, the efficiency wedge can be derived as follows. Start with the definition of the aggregate production function:

$$Y_t = e^{\tau_t^e} K_{t-1}^{\alpha} L_t^{1-\alpha} \tag{2.58}$$

Replacing Y_t, K_{t-1}, and L_t using (2.53), (2.56), and (2.57) yields:

$$\tau_t^e = \log \mathrm{E} \left[z_t^{\frac{1}{\alpha}} | z_t \geq \underline{z}(w_t, r_{t-1}) \right]^{\alpha} \tag{2.59}$$

The aggregated labour supply equation can be derived by taking the ratio $w_t L_t / Y_t$ and again replacing Y_t and L_t with the above values. Doing so reveals that $w_t L_t / Y_t = 1 - \alpha$.

2.E.4 Steady state proofs

Proof of Lemma 1. Labour demand is downwards sloping in the wage ($L_1^d(w; \lambda) < 0$): Taking (2.20) in steady state shows that $\partial N / \partial w < 0$. Combined with $Z'(w) < 0$, this gives us $\partial L / \partial w < 0$ from (2.57) taken in steady state.

Labour demand is increasing in the borrowing limit ($L_2^d(w;\lambda) > 0$): Taking (2.20) in steady state shows that, for a given w, $\partial N/\partial\lambda > 0$. (2.57) in steady state reveals that, for a given wage, $\partial L/\partial N > 0$, which together implies that $\partial L/\partial\lambda > 0$.

Labour supply is decreasing in the borrowing limit ($L_2^s(w;\lambda) < 0$): (2.53) and (2.56) in steady state reveal that, for a given wage, Y and K are proportional to N in steady state. The resource constraint, (2.26), in steady state then gives that c is also proportional to N, and $\partial c/\partial N > 0$ as long as $c > 0$. Combined with $\partial N/\partial\lambda > 0$ this implies that $\partial c/\partial\lambda > 0$. (2.7) gives that $\partial L^s/\partial c < 0$ for a given wage, which overall gives $\partial L^s/\partial\lambda < 0$. $\square$

Proof of Proposition 8. Fall in τ^e: Since $L_2^d(w;\lambda) > 0$ and $L_2^s(w;\lambda) < 0$ the market clearing wage unambiguously falls following a reduction in λ. Since $\underline{z}'_{ss}(w) > 0$ and $\tau^{e\prime}(\underline{z}) > 0$ this leads to a fall in τ^e. τ^l constant at zero: market clearing means that we are on both the labour demand and supply FOCs, so both $w = (1-\alpha)Y/L$ and $w = v'(L)c$, which gives $\tau^l = 0$ from the definition of τ^l. U constant at zero: by definition of market clearing $L^{s*} = L^s = L$, so $U \equiv L^{s*} - L = 0$. $\square$

Proof of Proposition 9. No change in τ^e: Since $\tau^e = \tau^e(\underline{z})$ and $\underline{z} = \underline{z}_{ss}(w)$, with a fully fixed wage there can be no change in τ^e. Increase in τ^l: Since $L_2^d(w;\lambda) > 0$ and $L_2^s(w;\lambda) < 0$, we get a fall in equilibrium L (due to rationing) following a fall in λ. As discussed in the proof of Lemma 1, for a fixed wage we also have $\partial c/\partial\lambda > 0$ so c falls. The definition of τ^l and the fact that firms remain on their labour demand curves (giving $w = (1-\alpha)Y/L$) means that the labour wedge can be expressed as: $\tau^l = 1 - v'(L)c/w$. Since $v''(L) > 0$, reducing L and c leads to an increase in τ^l. Increase in U: Since $L_2^d(w;\lambda) > 0$ and $L_2^s(w;\lambda) < 0$, for a fixed wage $U \equiv L^{s*} - L^d$ rises following the fall in λ. $\square$

Proof of Proposition 10. Consider a decrease in λ such that we have rationing unemployment. Larger fall in w leads to smaller fall in L: Since we are on the labour demand curve, any decrease in w will increase $L = L^d$ since $L_1^d(w;\lambda) < 0$, leading to a smaller fall in L. Larger fall in w leads to larger fall in τ^e: Since $\underline{z}'_{ss}(w) > 0$ and $\tau^{e\prime}(\underline{z}) > 0$, $\partial\tau^e/\partial w > 0$, so a larger fall in w must lead to a larger fall in τ^e. $\square$

Chapter 3

Labour Market Matching, Stock Prices & the Financial Accelerator

3.1 Introduction

It has long been understood that financial frictions can amplify business cycle models when they interact with asset prices. A shock which raises the net worth of productive agents will increase their asset demand if they are financially constrained, pushing up asset prices and again increasing net worth, in a cycle known as the financial accelerator. In this paper I show that this accelerator naturally emerges when we add financially constrained agents to the Diamond-Mortensen-Pissarides labour search and matching model.

Any theory of the financial accelerator requires a theory of how the economy's assets are priced. In the original Kiyotaki & Moore (1997) model the asset is land, which is in fixed supply. The asset price is used to clear the market, by ensuring demand equals this fixed supply. In a model where assets are not in fixed supply it is harder to price assets. If production of the asset in question is competitive then the asset price must equal the marginal cost of producing the asset. Thus simple RBC models where consumption is convertible one-for-one into capital generate an always unit price of capital, and no accelerator if capital is used as collateral. Adding adjustment costs to the model allows the marginal cost of producing, and hence the price of, capital to vary over the cycle, and hence reintroduces a financial accelerator.

I show that labour market frictions introduce an accelerator in an analogous way. The asset in my model is filled vacancies, or "matches" between a vacancies and workers. These can be thought of as equity in firms, and are traded by "experts", who also provide the funding for vacancy posting. In equilibrium the price of existing matches must equal the marginal cost of producing new matches, which depends on labour market tightness (the ratio of job vacancies to unemployment, $\theta \equiv v/u$) in an intuitive way: When many firms are posting vacancies the probability of filling your vacancy is low, meaning the marginal cost of producing a filled match is high. As an example, consider a positive productivity shock. This raises the value of a filled vacancy, increasing vacancy posting, and hence increasing labour market tightness. This increases the marginal cost of producing new matches, and hence pushes up the price of existing matches. This increases the net worth of experts, allowing them to fund more vacancies, and the cycle continues.

In other words, I use the assumption of a frictional labour market to create an upwards sloping supply curve for the economy's assets. This means that increases in demand for assets must lead to asset price increases, creating the financial accelerator. The most natural interpretation of the accelerator is as a feedback between the stock market and the labour market. In one direction, changes in stock prices affect expert net worth and hence the funds available for vacancy posting. In the other, changes in vacancy posting feed back into stock prices by changing the marginal cost of producing new firms, via labour market tightness.

I present a model where the experts are firms owned by the representative household, which could be thought of as banks. However, it is not crucial that they are banks to understand the story. The crucial requirement for my story is that whoever it is who benefits from increases in stock prices is financially constrained, and is the same person who provides funds for new vacancies. This sets the scene for the accelerator because increases in stock prices benefit agents who then reinvest that money in creating new vacancies, which further pushes up stock prices. One could instead model the experts as a separate species and get similar results.

My main result is thus demonstrating the existence of a financial accelerator in the search and matching model, which operates through labour market tightness. Secondly, I derive an arbitrage equation in my model between existing matches and vacancies which is identical to the standard free entry condition, with match value replaced with match price. The free entry condition, equating the value of a filled vacancy with the cost of producing one, holds in the standard matching model for both match value and price, since they are identical.

I show that it also holds in my model even after the introduction of financial frictions, for match prices but *not* value. This equation links market tightness and the price of a filled match through the marginal cost arguments made above, and implies a tight link in the model between the volatility of tightness and asset prices: for the standard matching elasticity of one half it implies that the volatility of market tightness must be twice the volatility of asset prices in the model. I construct measures of these volatilities, and show that this implies that if my model is calibrated to match the volatility of asset prices, it can explain 82% of the volatility of market tightness in the data.

Furthermore, this holds for any model which shares this arbitrage equation (including the standard matching model) and thus implies that, regardless of the source, any model which can match the volatility in asset prices will do equally well at matching the volatility of tightness. This suggests a potential avenue for work in the matching literature, focusing on improving the asset pricing abilities of these models. This result is inspired by the recent work of Winkler (2015), who argues that the key to generating sufficient amplification from financial frictions models is generating sufficient asset price volatility.

I show that wage stickiness interacts with the financial accelerator. Increasing the degree of wage stickiness in the model increases the gap between the volatilities of the models with and without financial frictions, showing that wage stickiness boosts the amplification given by the financial accelerator. Sticky wages make the stock market more volatile, which boosts the financial accelerator since stock prices feed back into expert net worth and hence vacancy posting.

Finally, I examine the role played by market incompleteness in my model. My baseline model assumes, as is common in the financial frictions literature, that agents can trade only in a bond which is not contingent on aggregate shocks. I also solve a version of the model which still contains financial frictions caused by a moral-hazard problem, but where agents trade a contingent bond, and show that this version does not deliver any amplification relative to the model without financial frictions. This highlights the key role that market incompleteness plays in generating the financial accelerator. The combination of asset values which are state dependent due to price movements and fixed liabilities generates the volatility of net worth required to deliver volatility in the real side of the economy. With complete markets expert liabilities also become state contingent, and I show that this can completely undo the financial accelerator. To avoid confusion, in the rest of the paper reference to a model "with financial frictions" refers to the model with financial frictions and incomplete markets, unless otherwise noted.

The remainder of the paper is structured as follows. In section 3.2 I review related literature. In section 3.3 I set up the baseline model with incomplete markets, and in section 3.4 I set up the models without financial frictions and with complete markets. In section 3.5 I analyse the differences between the models via their key equations, and in section 3.6 I present both steady state and dynamic analytical results. Section 3.7 contains numerical results and robustness checks, and section 3.8 concludes.

3.2 Related literature

My paper is related to several broad strands of literature. It builds on the labour market matching models of Diamond, Mortensen & Pissarides, summarised for example in Pissarides (2000). Within this literature it is also related to papers on the ability of the matching model to quantitatively replicate the data, such as Shimer (2005) and Hagedorn & Manovskii (2008). I contribute to this literature by showing that financial frictions can help resolve the Shimer critique, and by demonstrating the key relationship between asset-price volatility and volatility in market tightness.

I also build on the large financial frictions literature. Within this literature my work is closest to those papers which emphasise the interplay between asset prices and net worth, such as the early contribution by Kiyotaki & Moore (1997). My contribution is to show that a financial accelerator naturally arises in my model because of the matching market, even without assumptions on varying marginal products of agents, or adjustment costs on capital. Bernanke, Gertler & Gilchrist (1999) provide a model where adjustment costs on capital provide the movements in the price of capital, as well as an extensive review of the early literature.

My paper is not the first to investigate the intersection between financial and labour market frictions. My contribution here is that my paper is the first, to my knowledge, to use the asset price implications of labour market frictions to generate a financial accelerator. By putting financial frictions on the people who own firms, rather than within the firms themselves, firms' stock market prices affect expert net worth, and hence the funds that experts have to reinvest in their firms for vacancy creation. This feature is absent from the existing literature. Christiano, Trabandt & Walentin (2011) combine matching unemployment and financial frictions in a small open economy framework. Mumtaz & Zanetti (2013) add labour market frictions to the Bernanke, Gertler & Gilchrist (1999) framework and discuss how labor market frictions amplify or dampen the response of the model to different shocks. Petrosky-Nadeau (2013) introduces financial frictions into the search and matching framework and notes changes in amplification and propagation as well as effects on wage bargaining positions. Quadrini & Sun (2015) argue that firms can improve their bargaining position versus workers by taking on more debt, and estimate the effect this has on hiring in a structural model. Schoefer (2015) argues that wage stickiness affects hiring by making the net worth of firms more volatile, impacting the resources they can put towards paying hiring costs.

Other papers study the interaction of labour and finance along other dimensions. Favilukis & Lin (2015) show how sticky wages help explain the equity premium (by making profits more volatile and hence equity riskier) and several other asset pricing facts. Petrosky-Nadeau, Kuehn & Zhang (2013) show how an appropriately calibrated matching model endogenously generates rare disasters, and hence again helps explain the equity premium. Caggese & Cunat (2008) study how financially constrained firms choose between hiring workers on fixed-term and permanent contracts. The result that the financial accelerator relies on incomplete markets has been explored in the existing literature, for example by Carlstrom, Fuerst, Ortiz & Paustian (2014) and Dmitriev & Hoddenbagh (2014). Finally, I exploit the tight link in my model between labour market tightness and stock prices. This feature is present in existing matching models with a free entry condition on vacancy creation. Farmer (2012a) and Hall (2014) present evidence on the tight link between unemployment and the stock market, and Farmer (2012b) exploits this link theoretically in a model of multiple equilibria.

3.3 Model

The model combines elements of Getler & Karadi's (2011) financial frictions model with the standard search and matching model. As in the standard matching model, I abstract from capital, and instead have experts trade in the equity of firms. Firms and workers must match according to a matching technology, and vacancy posting costs must be paid in order to maintain vacancies. Gertler & Karadi (2011) choose to model experts as a separate species from households, but I instead choose to model them as intermediary firms owned by the households. The name "experts" is retained for consistency with the literature. The model features a representative household, which supplies labour and saves using a risk free bond. Experts post vacancies, trade existing matches in a spot market, and borrow using the risk free bond.

Firms in the model are simply matches between vacancies and workers, and do not face any significant optimisation problem. Note that experts own the equity of all the firms in the economy, which enables me to combine the expert and firm sectors and consider experts directly posting vacancies. My financial structure is thus quite stylised. In particular, all firms are funded with equity from experts, and all experts are funded with risk free debt from the household. The household is unable to directly invest in the equity of firms. Time is discrete and the horizon infinite.

3.3.1 Individual expert's problem

Experts are owned by the representative households. They are restricted severely in their equity issuance: they cannot raise money via equity, and must purchase assets using retained earnings or debt. Experts exit exogenously each period with probability $(1 - \sigma)$ and new experts are created each period so that the mass of experts is constant. New experts receive an exogenous equity injection from the representative consumer. If the value of being an expert exceeds one (as it does in a neighbourhood of the steady state) experts will not pay out dividends until they exogenously shut down. In this case, an expert's balance sheet gives us:

$$Q(s)k'_o + \kappa z v = d' + n \tag{3.1}$$

n is beginning of period net worth, which is a state variable from the expert's point of view, and s is the aggregate state. d' is borrowing, which is combined with net worth to purchase assets. All aggregate variables are indexed by the aggregate state, s. Individual level variables, such as d', will be denoted without reference to state variables before they are optimised, and $d' = d(s,n)$ will refer to their optimised, equilibrium values. On the left hand (asset) side, the expert has two choices. Firstly, she can buy an existing match on the spot market for price $Q(s)$. The number of existing matches she wishes to purchase is denoted by k'_o. These matches produce next period, and then a fraction ρ_x exogenously separate. Those that don't separate can be resold tomorrow for price $Q(s')$. Alternatively, the expert can decide to set up some new matches herself by issuing vacancies, v. She pays a flow vacancy cost κz per vacancy, where z is aggregate productivity, and a fraction $q(\theta(s))$ are successful. q denotes the vacancy filling probability and θ market tightness, both of which individuals take as given. If a vacancy is successful today then it produces for sure tomorrow, and then a fraction ρ_x exogenously separate. Hence notice that buying a existing match today or setting up a match yourself yield the same payoff tomorrow.

I assume away the idiosyncratic risk that an expert's vacancies don't match by assuming that the expert issues a continuum of vacancies. Thus if an expert posts v vacancies today then it gets for sure $q(\theta(s))v$ successful matches, and we can think of the expert as directly choosing the number of successful matches, k'_n, as opposed to the number of vacancies. So if we define $k'_n \equiv q(\theta)v$ we can rewrite the balance sheet as:

$$Q(s)k'_o + \frac{\kappa z}{q(\theta(s))}k'_n = d' + n \tag{3.2}$$

Note that experts can only post non-negative vacancies, so $v \geq 0$ is a constraint for the expert. This implies the equivalent constraint $k'_n \geq 0$ as long as the probability of a successful match is non zero ($q > 0$). Expert net worth next period is the return on assets less the repayment of debt:

$$n' = \left(z' - w(s') + (1-\rho_x)Q(s')\right)\left(k'_o + k'_n\right) - r(s)d' \tag{3.3}$$

Where $w(s)$ is the wage, which depends on the aggregate state, and $r(s)$ is the interest rate on debt. Combining this with the balance sheet equation gives:

$$n' = \left(z' - w(s') + (1-\rho_x)Q(s') - r(s)Q(s)\right)k'_o +$$

$$\left(z' - w(s') + (1-\rho_x)Q(s') - r(s)\frac{\kappa z}{q(\theta)}\right)k'_n + r(s)n \tag{3.4}$$

I derive a constraint on borrowing using Gertler & Karadi's (2011) limited commitment problem. Within this period but after raising funds, experts can abscond with an amount of resources equal to a fraction Λ of the value of the assets they invested in. The remaining fraction $1-\Lambda$ is exogenously destroyed, leaving nothing for the lender to recover. If experts abscond they lose the franchise value of being an expert, but gain the stolen resources. Since this is a within-period problem, lenders can anticipate exactly when an expert will abscond with their resources and they will restrict the amount they lend to make sure this doesn't happen. Define the value function conditional on a choice of (k'_o, k'_o) as $V^*(n, s; k'_o, k'_n)$. Then this limited commitment problem gives the constraint:

$$\Lambda\left(Q(s)k'_o + \frac{\kappa z}{q(\theta(s))}k'_n\right) \leq V^*(n, s; k'_o, k'_n) \tag{3.5}$$

This requires that the value of the expert must exceed the value of the assets she has the potential to steal, in order to guarantee that the expert does not have an incentive to abscond with them. The conditional value function is given by:

$$V^*(n, s; k'_o, k'_n) = \mathrm{E}\left[\Omega(s', s)\left((1-\sigma)n' + \sigma V(n', s')\right)\middle|\, s\right] \tag{3.6}$$

Where $\Omega(s', s) \equiv \beta u'(c(s'))/u'(c(s))$ is the consumer's stochastic discount factor (SDF), and where n' is replaced with the value implied by (3.4). $V(n', s')$ is the overall maximised value next period, and today's value is given by the maximisation:

$$V(n,s)=\max_{(k_o',k_n')} V^*(n,s;k_o',k_n') \tag{3.7}$$

Subject to (3.4), (3.5) and $k_n' \geq 0$. The following lemma summarises the solution to the expert's problem. I focus on the case where the non-negativity constraint on vacancies never binds, since my expert sector will aggregate and this case is thus consistent with the observation that total vacancies are always positive in the data.

Lemma 2. *If the non-negative vacancies constraint isn't binding and prices are such that the expert cannot acquire infinite value, then the solution to the individual expert's problem requires that capital price and tightness satisfy:*

$$Q(s)=\frac{\kappa z}{q\left(\theta(s)\right)} \tag{3.8}$$

This implies that old and new matches yield the same return, and individual experts are indifferent between the two and optimise over the sum $k' \equiv k_o' + k_n'$. Defining leverage as $\phi \equiv Q(s)k'/n$, expert net worth evolves as:

$$n'=\left(\left(r_k(s',s)-r(s)\right)\phi+r(s)\right)n \tag{3.9}$$

Where $r_k(s',s)$ is the return on investing in a match, given by:

$$r_k(s',s)\equiv\frac{z'-w(s')+(1-\rho_x)Q(s')}{Q(s)} \tag{3.10}$$

Optimal leverage is independent of expert net worth, and optimal k', d' and V are linear in net worth. Expert value is given by $V(n,s)=\nu(s)n$, where $\nu(s)$ is defined recursively by:

$$\nu(s;\phi)=\mathrm{E}\left[\Omega(s',s)\left(1-\sigma+\sigma\nu(s')\right)\left(\left(r_k(s',s)-r(s)\right)\phi+r(s)\right)\middle|\, s\right] \tag{3.11}$$

and $\nu(s) = \max_\phi \nu(s;\phi)$ *subject to the moral hazard constraint* $\Lambda\phi \leq \nu(s;\phi)$. *Equilibrium leverage,* $\phi = \phi(s)$, *is given by the value that solves that maximisation. Total match demand,* $k' = k(s,n)$, *and debt,* $d' = d(s,n)$, *are given by* $Q(s)k(s,n) = \phi(s)n$ *and* $d(s,n) = (\phi(s)-1)n$. *If the moral hazard constraint binds then value and leverage are jointly determined by:*

$$\nu(s) = \mathrm{E}\left[\Omega(s',s)\left(1-\sigma+\sigma\nu(s')\right)\left(\left(r_k(s',s)-r(s)\right)\phi(s)+r(s)\right)\middle| s\right] \tag{3.12}$$

$$\Lambda\phi(s) = \nu(s) \tag{3.13}$$

If the moral hazard constraint isn't binding then prices must satisfy:

$$\mathrm{E}\left[\Omega(s',s)\left(1-\sigma+\sigma\nu(s')\right)\left(r_k(s',s)-r(s)\right)\middle| s\right] = 0$$

The expert is then indifferent about her leverage, and expert value is given by:

$$\nu(s) = \mathrm{E}\left[\Omega(s',s)\left(1-\sigma+\sigma\nu(s')\right)r(s)\middle| s\right]$$

The proof of this lemma is left to the appendix, but the intuition and results are straightforward. The idea is that as long as an expert wants to post vacancies ($v > 0$) then she must be indifferent between posting vacancies and purchasing existing matches on the stock market. This is because the two assets have identical payoffs tomorrow, and the moral hazard constraint does not restrict the expert from performing arbitrage between them. For the expert to be indifferent it must be that the cost of purchasing an existing match, $Q(s)$, is equal to the cost of producing a new match, $\kappa z/q\left(\theta(s)\right)$, as stated in equation (3.8). Notice that this equation is identical to the free entry condition found in the standard matching model (if the value of a match is replaced with the price of a match), even though the derivation is different.

3.3.2 Aggregating the experts

Experts as a whole enter the period with total undepreciated matches K, and a total value of debt to be repaid D. Define N_c as net worth from continuing experts and N_e as net worth from new experts. Total expert net worth at the beginning of the period is thus:

$$N(s) = N_c(s) + N_e(s)$$

Where each of the components is given by:

$$N_e(s) = (1 - \sigma)\, w_e z$$

$$N_c(s) = \sigma \left((z - w(s) + (1 - \rho_x) Q(s))\, K - D \right)$$

New experts get net worth proportional to aggregate productivity. The net worth of continuing experts comes from the output and resale value of their total capital less their debt repayment. Thus overall expert net worth evolves according to:

$$N(s) = \sigma \left((z - w(s) + (1 - \rho_x) Q(s))\, K - D \right) + (1 - \sigma)\, w_e z \tag{3.14}$$

The transitions for K and D can be found by aggregating the individual policy functions:

$$K'(s) = \frac{\phi(s) N(s)}{Q(s)} \tag{3.15}$$

$$D'(s) = r(s)\, (\phi(s) - 1)\, N(s) \tag{3.16}$$

Note that the definition of D is slightly different from that of d: D is defined to contain the interest rate, for convenience when it is used as a state.

3.3.3 The labour market

The structure of my labour market is standard. The total mass of workers within the household is normalised to one, so unemployment at the beginning of the period is given by:

$$u(s) = 1 - K \tag{3.17}$$

Tightness is defined as usual as the ratio of total vacancies to unemployment:

$$\theta(s) = \frac{v(s)}{u(s)} \tag{3.18}$$

The matching function is assumed to take the constant returns to scale Cobb Douglas form:

$$m(s) = \psi_0 u(s)^{\psi_1} v(s)^{1-\psi_1} \tag{3.19}$$

This allows us to express the vacancy filling probability as a function only of tightness:

$$q(s) = \frac{m(s)}{v(s)} = \psi_0 \theta(s)^{-\psi_1} \tag{3.20}$$

Total employment next period is the sum of new and undepreciated matches:

$$K'(s) = m(s) + (1 - \rho_x)K \tag{3.21}$$

In this paper I take a reduced form approach to wage determination, rather than modelling the mechanisms more explicitly. While this approach is somewhat unsatisfactory, I note that the focus of the present paper is to understand the links between financial frictions and the matching model. Hence as long as the wage process is chosen carefully to realistically match the data on wages one would expect my results to also hold in a model with alternative wage determination mechanisms which were also able to match wage data. Following Michaillat (2012), I assume that the wage is an exogenous function of productivity:

$$w = \bar{w} z^{\gamma} \tag{3.22}$$

$\bar{w}$ thus controls the average wage and γ the degree of wage rigidity. $\gamma = 0$ corresponds to wages which are completely rigid, and $\gamma = 1$ corresponds to wages which move one-for-one with productivity. This is the equilibrium outcome of fully flexible Nash-wages in the standard DMP model when vacancy posting costs are proportional to productivity, unemployment income is proportional to wages, and the utility function is logarithmic. Hence I refer to low values of γ as generating relatively rigid wage, and higher values generating flexible wages.

3.3.4 Goods market and household problem

In the baseline incomplete markets model the household lends to the expert using a risk free, one period bond. Household optimality requires that the interest rate satisfies the standard Euler equation:

$$r(s) = \frac{u'(c(s))}{\beta \mathrm{E}[u'(c(s'))|s]} \tag{3.23}$$

Where $c(s)$ is consumption. Since labour of each household member is indivisible, there is no labour supply choice. I assume that the household always chooses for all of its unemployed members to search for a job, and hence the Euler equation above is the only household optimality condition. Goods market clearing requires that all output is either consumed by the household, or used to pay vacancy posting costs:

$$zK = c(s) + \kappa z v(s) \tag{3.24}$$

Finally, productivity follows a stationary AR(1) process:

$$\log z' = (1-\rho)\log \bar{z} + \rho \log z + \sigma_z \varepsilon' \tag{3.25}$$

Where ε is an independent and identically distributed standard normal. $\bar{z}$ controls the mean of productivity, ρ its autocorrelation, and σ_z the standard deviation of productivity innovations.

3.3.5 Definition of equilibrium

The model can be solved with three state variables: $s = (z, K, D)$. Productivity, z, and employment, K, the state variables in the standard matching model, are augmented with the debt repayment made by experts to the household, D. Since I will be linearising around a steady state where the financial friction binds, I define equilibrium under the assumption that the financial friction always binds:

Definition 1. *Incomplete markets equilibrium (IME) is a sequence of quantities and prices v, ϕ, N, D, K, Q, θ, m, q, u, r, c, z, w, r_k, and v such that:*

1. *Households optimise taking prices as given: (3.23)*
2. *Experts optimise taking prices as given: (3.8) (3.10) (3.12) (3.13) (3.14) (3.15) (3.16)*
3. *The goods market clears: (3.24)*
4. *The labour market evolves according to the matching function: (3.18) (3.19) (3.20) (3.17) (3.21)*
5. *The wage is given by the wage rule: (3.22)*
6. *Productivity evolves according to: (3.25)*

In the next section I set up two comparison models with different financial structures. Competitive equilibrium is defined similarly to the above for these economies, and the definitions are excluded for brevity.

3.4 Comparison models

3.4.1 Model without financial frictions

The model without any financial frictions (i.e. no limited-commitment problem, and market completeness) dispenses with most of the expert equations, and instead has matches valued simply using the consumer's SDF. The derivation is standard and is omitted. Given the definition of r_k optimality can be compactly stated as:

$$\mathrm{E}\left[\Omega(s', s) r_k(s', s) \middle| \, s\right] = 1 \tag{3.26}$$

This is simply another way of writing the usual recursion for job value. This model is simpler than the model with experts because we do not have to worry about financial variables. The definition of equilibrium, which I label a Standard Equilibrium (SE), is similar to the definition of an IME, replacing the expert optimality equations with just (3.26).

3.4.2 Complete markets model

I now consider the model where experts and consumers trade state contingent securities instead of just risk free debt, but we retain the limited-commitment problem. As in the incomplete markets model, arbitrage between vacancies and existing matches allows me to combine them into a single asset, which I impose from the start for simplicity of exposition. An individual expert's balance sheet is now:

$$Q(s)k' = n + \int_{s'} d(s') \, \mathrm{d}s' \tag{3.27}$$

Where $d(s')$ donates the quantity of securities purchased that are payable if next period's state is s', and $\mathbf{d} \equiv \{d(s')\}$ denotes the collection. Net worth evolves according to:

$$n' = (z' - w(s') + (1 - \rho_x)Q(s'))\, k' - r(s')d(s') \tag{3.28}$$

The interest rate is state contigent and derived from the representative household's optimality:

$$r(s') = \frac{u'(c(s))}{\beta u'(c(s'))p(s'|s)} \tag{3.29}$$

Where $p(s'|s)$ is the marginal density of the state s' conditional on today's state, s. Expert value conditional on a choice of k' and $\mathbf{d}$ is given by:

$$V^*(n, s; k', \mathbf{d}) = \mathrm{E}\left[\Omega(s', s)\left((1 - \sigma)n' + \sigma V(n', s')\right)\middle| s\right] \tag{3.30}$$

Where it is understood that n' is replaced using (3.28). Experts maximise overall value:

$$V(n, s) = \max_{k', \mathbf{d}} V^*(n, s; k', \mathbf{d}) \tag{3.31}$$

Subject to (3.27), (3.28), and the moral hazard constraint:

$$\Lambda Q(s)k' \leq V^*(n, s; k', \mathbf{d}) \tag{3.32}$$

Note that now the maximisation is also over all the state contingent securities, $\mathbf{d}$. The following proposition establishes the central result of the complete markets model:

Lemma 3. *Assuming that the moral hazard constraint is always binding, there is a solution to the individual expert's problem in the complete markets model featuring constant leverage, $\bar{\phi}$. Expert value is linear in net worth, with a constant first derivative: $V(n, s) = \bar{\nu}n$. Leverage satisfies $\Lambda\bar{\phi} = \bar{\nu}$, and the solution requires that:*

$$\mathrm{E}\left[\Omega(s', s)r_k(s', s)\middle| s\right] = \frac{\bar{\nu} - (1 - \bar{\phi})\left(1 - \sigma + \sigma\bar{\nu}\right)}{\left(1 - \sigma + \sigma\bar{\nu}\right)\bar{\phi}} \tag{3.33}$$

The proof is relegated to the appendix, but the intuition for constant leverage and marginal value is relatively simple. With complete markets the expert is able to use contingent debt to allocate resources across future states of the world. She actually has a lot of freedom to do this, because the moral hazard constraint limits her overall borrowing, not how she allocates debt across the contingent states. This is why marginal value, $\bar{\nu}$ has to be constant, because if it was not the expert would use contingent debt to borrow in states with low value, and transfer those resources to states with higher value. Since she is unconstrained in doing this and value is linear in net worth, she would take infinitely large positions, and achieve infinite overall value. This would violate that the moral hazard constraint is binding (which I assumed) since with infinite value she can always borrow more. Given constant marginal value, constant leverage follows trivially from the binding borrowing constraint.

Given the constant marginal value and leverage, equation (3.33) delivers the main result of the complete markets model. Notice that the right hand side is constant, and that apart from this the equation is identical to the first order condition of the standard equilibrium, (3.26). In the limiting case of $\bar{\nu} = 1$ the two equations are exactly identical, and the standard and complete markets models deliver identical equilibria. In general, the two models are identical up to this "wedge" due to the moral hazard constraint, and we will see in later sections that they deliver very similar dynamics.

3.5 Discussion of key equations

Having set up the three models, in this section I discuss the differences between them using two key model equations: the free entry or arbitrage equation, and the discounted sum pricing matches. This thus serves as an introduction to the models before moving on to more explicit analytical and numerical results in later sections. One key idea from this section is the tight link that the matching model imposes between the volatility of market tightness and asset prices, which is an idea also taken up in Hall (2014).

3.5.1 Comparing equations: The free entry condition

The first thing to note is that both the financial frictions models and the standard matching model contain the familiar equation:

$$Q(s) = \frac{\kappa z}{q\left(\theta(s)\right)} = \frac{\kappa z}{\psi_0}\theta(s)^{\psi_1}$$

The interpretation is slightly different in the three models. In the standard matching model this is the free entry condition, stating that the value of posting a vacancy should be equal to zero. Q, the value of a filled vacancy, should be equal to the cost of posting a vacancy, adjusted for the probability of success, leaving no surplus. Notice that in the standard matching model Q is both the value of a filled vancancy, and the price a filled vacancy would trade on the market.

However, in the financial frictions models we must be careful because Q must be interpreted as the market price of a filled match. This will be different from the value of a match to an expert due to the financial friction. In the standard model the value of posting a vacancy must be equal to zero in equilibrium, but this is not true if the financial friction binds: Experts would like to post another vacancy, they just don't have the funds. If this is the case, how do we still recover an equation identical to the free entry condition of the standard model?

This is because the equation is in fact a no-arbitrage equation for the experts, which says that they must be indifferent between creating a new match themselves (by posting $1/q$ vacancies) and purchasing an existing match on the spot market. Because I have assumed away vacancy risk, these two choices represent assets which give identical payoffs in the future: if you buy a match today or create one, you have a match tomorrow in either case. Hence if they had different prices experts could make infinite profit by going long in one and short in the other, which cannot happen in equilibrium. This is not prevented by the financial friction. In other words, matches are priced at marginal cost, as we would expect as the spot market for matches is competitive.

A complete proof can be found in the appendix, and it relies on aggregate vacancy posting being positive, which I assume. In the appendix I also provide another way of deriving the free entry condition in my model, by assuming the existence of competitive "match producing firms" who create matches and sell them to experts.

Another way of putting this is to note that this equation is the key to understanding where stock market value derives from in this model. Since there is no physical capital in the economy, matches between workers and firms are the only physical asset which can be owned. But these matches only have value to the extent that they can't be replicated costlessly: as discussed above, they are priced at marginal cost. In the limit where matching frictions disappear ($\kappa \to 0$), equation (3.8) implies that $Q = 0$ at all times. This is because current matches must be a worthless asset in order for this equation to hold, since they can be costlessly replicated by posting enough (free) vacancies.

At this point it is worth noting that the free entry or arbitrage conditions place a very strong link in the model between the volatilities of asset prices and market tightness. To see this, take logarithms of (3.8) to get:

$$\log \frac{Q(s)}{z} = \log \frac{\kappa}{\psi_0} + \psi_1 \log \theta(s)$$

This implies a very strong link between the standard deviations of labour market tightness and asset prices:

$$\sigma\left(\log \theta\right) = \frac{1}{\psi_1} \sigma \left(\log \frac{Q}{z}\right) \tag{3.34}$$

Where $\sigma(x)$ refers to the standard deviation of variable x. This implies that the volatility of log tightness, a key moment in the search and matching literature, is pinned down exactly by the volatility of log asset prices (scaled by labour productivity). Given the standard value of $\psi_1 = 1/2$ the above equation implies that the model will always generate a volatility of tightness twice that of asset prices. To the extent that the introduction of financial frictions can increase the volatility of asset prices we should thus expect them to increase the volatility of the labour market as well via this arbitrage equation.

3.5.2 Comparing equations: The discounted sum

Since the models with and without financial frictions both have the same free entry condition, where is the substantive difference between them? In this section I show that much of the difference between the models can be understood via the recursion for the price of a match. In the standard matching model the value of a filled match to its owner can be expressed recursively as:

$$Q(s) = \mathrm{E}\left[\Omega(s', s)\left(z' - w(s') + (1 - \rho_x)Q(s')\right)\middle| s\right]$$

As previously explained, this is also equal to the market price of a filled match since there are no financial frictions. However, in the appendix I derive the following expression for the price of a match in the model with financial frictions and incomplete markets:

$$Q(s) = \frac{\mathrm{E}\left[\tilde{\Omega}(s', s)\left(z' - w(s') + (1 - \rho_x)Q(s')\right)\middle| s\right]}{\mathrm{E}\left[\tilde{\Omega}(s', s)\middle| s\right] / \mathrm{E}\left[\Omega(s', s)\middle| s\right] + \frac{\lambda(s)\Lambda}{1+\lambda(s)}}$$

Where $\tilde{\Omega}(s', s) \equiv \Omega(s', s)\left(1 - \sigma + \sigma\nu(s')\right)$ is the expert's SDF, which is the household's SDF "twisted" by the fact that, due to the financial frictions, the value of funds might be higher inside the intermediary than in the hands of the household . The two recursions are similar except that in the incomplete markets recursion: 1) the numerator uses the expert's SDF instead of the consumer's and 2) the recursion is divided by some extra terms.[1] $\lambda(s)$ is the equilibrium lagrange multiplier on an expert's limited commitment constraint. The equation shows us that, ceteris paribus, $Q(s)$ is decreasing in $\lambda(s)$. In other words, the more the financial constraint binds, the lower is the asset price, which is intuitive since we are discussing asset demand, and a tighter borrowing constraint reduces the funds available to purchase assets, pushing down prices. The arbitrage equation, (3.8) reveals that a lower price must also mean lower vacancy posting. This means that we can tell a rough story where the model with financial frictions is either a dampened or amplified version of the standard matching model, depending on the cyclical behaviour of the financial friction. If the financial friction binds less in booms ($\lambda(s)$ countercyclical) then the model will be an amplified version of the standard matching model. This is because in a boom not only is the value of a match higher, but also now the experts are less constrained and can fund more matches. This can happen if asset prices are sufficiently procyclical so that expert net worth increases enough in booms to relax their borrowing constraints. On the other hand, if the financial friction binds more in booms ($\lambda(s)$ procyclical) then the model will be a dampened version of the standard matching model. This can happen if asset prices are not sufficiently procyclical, so that in a boom expert net worth does not increase enough. In this case experts will feel more constrained in a boom, because they want to invest to take advantage of higher productivity but do not have sufficient net worth. While both cases are possible, depending on how you calibrate the model, we will see in the numerical section that a model calibrated to match the volatility of asset prices in the data will deliver amplification.

Finally, the complete markets financial frictions model delivers a discounted sum which is similar to the standard matching model up to a constant wedge. Since this

[1] We can verify that this equation reduces to the standard discounted sum from a normal matching model if we remove the financial friction by setting $\Lambda = 0$ so that the experts can't steal anything, and hence aren't constrained in equilibrium: If $\Lambda = 0$ then the expert is never constrained, so $\lambda = 0$. We can verify in this case that $\nu \equiv 1$, which means that $\tilde{\Omega} = \Omega$ (i.e. the expert's SDF is just the consumer's SDF). Finally this means that the denominator is equal to one, leaving: $Q(s) = \mathrm{E}\left[\Omega(s', s)\left(z' - w(s') + (1 - \rho_x)Q(s')\right)\middle| s\right]$

Figure 3.1: Graphical solution to the steady state of the financial frictions model

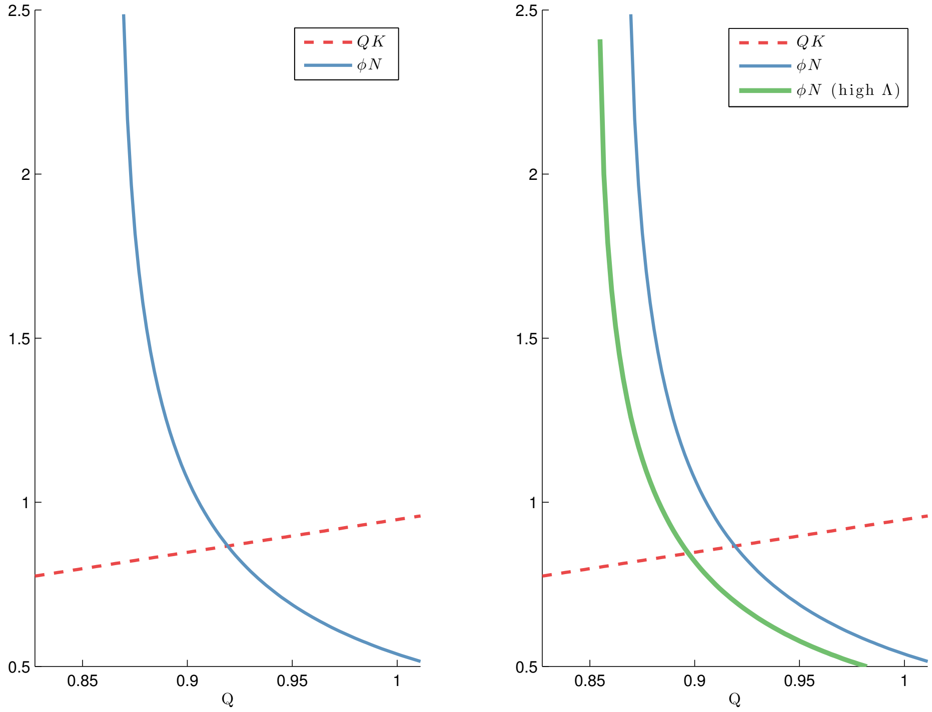

The left panel plots $L(Q)$ in dashed red and $R(Q)$ in solid blue. Their crossing gives the steady state value for Q. The parameterisation is the one used in the numerical work below. The right panel plots the same, and the thick green line plots $R(Q)$ for a value of Λ 20% higher than the baseline value, leading $R(Q)$ to shift to the left and steady state Q to be lower.

wedge is constant, we should not expect any drastic cyclical differences between the complete markets and standard model.

3.6 Analytical results

In this section I present analytical results for the steady states of the models, as well as for a special case of the fully dynamic model. These serve to illustrate the key mechanisms of the models in a sharp and transparent manner before I move on to the numerical results. In particular, I am able to prove that financial frictions must both increase unemployment in steady state, and increase its volatility. Using a sequence of proofs I show how it is crucially the transmission of net worth, and experts' inability to insure against it, which causes the divergence between the models with and without financial frictions.

3.6.1 Steady state results

In this section I compare how the steady states are determined in the models with and without financial frictions. I focus on the non-stochastic steady states, which means that the models with and without complete markets become identical. I also abstract from wage setting and focus on steady states conditional on a given, fixed wage. I denote steady state variables by omitting the explicit depending on the state, s. The determination of the steady state in the model without financial frictions can be summarised in the following three equations:

$$Q = \frac{\beta(z-w)}{1-\beta(1-\rho_x)}$$

$$\downarrow$$

$$\theta = \left(\frac{\psi_0 Q}{\kappa z}\right)^{\frac{1}{\psi_1}} \tag{3.35}$$

$$\downarrow$$

$$K = \frac{\psi_0\theta^{1-\psi_1}}{\psi_0\theta^{1-\psi_1}+\rho_x} \tag{3.36}$$

As indicated by the arrows, we can solve the equations sequentially. The discounted match surplus is given on the top, which gives us the tightness required by free entry, which gives us the steady state level of employment.

Solving for the steady state in the model with financial friction is more complicated, but while it is hard to get analytical solutions we can characterise the equilibrium graphically. The definition of expert leverage gives us $QK = \phi N$, which we can interpret as the intersection of the supply and demand for matches. The right hand side gives us the total resources experts are putting towards buying old and new matches: net worth multiplied by leverage. The left hand side tells us that this must be spent on the total value of matches in the economy: their price multiplied by their quantity. To solve for equilibrium I note that we can express both the left and right hand sides solely as functions of Q. The left hand side, which I interpret as match supply, uses equations (3.35) and (3.36), which are common with the standard model:

$$L(Q) \equiv QK = Q\frac{\psi_0\left(\frac{\psi_0 Q}{\kappa z}\right)^{\frac{1-\psi_1}{\psi_1}}}{\psi_0\left(\frac{\psi_0 Q}{\kappa z}\right)^{\frac{1-\psi_1}{\psi_1}} + \rho_x}$$

Notice that $L'(Q) > 0$, so our supply curve is upwards sloping. The demand curve can be shown to be downwards sloping because both leverage (ϕ) and net worth (N) are decreasing in Q. The proof is left to the appendix, but the intuition is simple. The steady state return on investing is $r_k = (z-w)/Q + 1 - \rho_x$, which is decreasing in Q. A lower return reduces expert value and hence the maximum leverage allowed by the borrowing constraint, hence $\phi'(Q) < 0$. Lowever returns and leverage both reduce expert earnings and hence steady state net worth, so $N'(Q) < 0$. Hence $R(Q) \equiv \phi N$, with $R'(Q) < 0$.[2]

As shown in Figure 3.1, equilibrium Q is at the intersection of the supply and demand curves. Once we have Q, equilibrium tightness and employment can be calculated as in the model without financial frictions. We can also use the graph to prove some results using comparative statics, and illustrate how the financial friction

[2]One very important issue here is uniqueness of equilibrium. Gertler & Karadi (2011) do not discuss this, but it is actually possible to for their model to feature two steady states for some parameterisations, and thus admit the possibility of multiple equilibria selected by sunspots. To see this, note that in steady state we solve for ϕ and ν from the expert equations (3.37) and (3.38) for a given value of r_k. Combining the two equations gives a quadratic equation in ϕ, giving two different solutions. Why is there multiplicity here? The intuition is simple. Leverage is limited by expert value, but expert value is higher when you're allowed more leverage. This multiplicity is a potentially interesting source of fluctuations, however for my baseline calibration there is actually only a unique steady state, around which I linearise.

affects the economy in steady state. Firstly, it is worth asking how the steady state of the financial frictions economy compares to the steady state of the standard matching model when they are given the same parameter values:

Proposition 12. *In a steady state where the financial friction binds, employment is strictly lower than in the model without financial frictions.*

Proof. If the financial frictions model had employment weakly greater than the model without financial frictions then tightness would also be weakly greater, and by (3.8) so too would be the steady state match price, Q. But since $r_k = (\bar{z} - \bar{w})/Q + 1 - \rho_x$ and $r_k = r = 1/\beta$ in the model without frictions, this would imply $r_k \leq r$ in the financial frictions model, in which case the financial friction does not strictly bind. □

This result is perhaps to be expected. The intuition is quite simple. In the steady state of the standard economy, the return on capital is equal to the interest rate: $r_k = r$. If the financial frictions economy had the same level of employment as the standard economy it would have to be the case that $r_k = r$ in the financial frictions economy too. However, if $r_k = r$ then the financial friction does not strictly bind, because experts do not make positive profits on lending, and are hence indifferent about lending more. The next proposition establishes some comparative statics within the financial frictions models:

Proposition 13. *In a steady state where the financial friction binds, an increase in the amount that experts can expropriate, Λ, or a reduction of the equity injection given to new experts, w_e, reduces the steady state match price, Q, and hence employment, K.*

Proof. Both of these changes leave the supply curve, $L(Q)$, unchanged, while shifting the demand curve, $R(Q)$, to the left. For a given Q, reducing w_e shifts $R(Q)$ to the left by reducing N while leaving ϕ unchanged, while increasing Λ reduces ϕ and hence consequently also reduces N. □

Both of these changes reduce the funds that experts can allocate to purchasing matches. Increasing Λ allows experts to steal more, and hence requires them to have lower leverage, and reducing w_e directly reduces expert net worth. Both of these shift the demand curve to the left, as shown in Figure 3.1, reducing steady state employment.

3.6.2 Dynamic results

In this section I analytically explore some features of the dynamic equilibria in a special case of the models with log utility and wages proportional to productivity. This specialisation is useful because it implies a particularly simple equilibrium in the model without financial frictions: unemployment is constant over the business cycle in response to productivity shocks. This stark result allows us to characterise what elements of the financial frictions model bring to the table, because under certain conditions we can replicate this result in the financial frictions economy. The following proposition establishes the initial result:

Proposition 14. *If wages are proportional to productivity and the household has log utility, then the model without financial frictions has an equilibrium where unemployment is constant in response to productivity shocks. The price of a match is proportional to current productivity:* $Q(s) = \bar{Q}z$.

This result is standard, and also relies on my assumption that vacancy posting costs are proportional to current productivity. Following a positive productivity shock the linear wage ensures that the value of a match rises proportionally to current productivity, as does the posting cost. Hence there is no incentive to change vacancy posting. The stock market does move though, and stock prices are proportional to current productivity.

Does adding financial frictions break this result? In the end it will, but I show that the simultaneous presence of several elements is required. Firstly a volatile stock market is not enough on its own. After all, the model without financial frictions generates movements in stock prices. The second element we need is an interaction between stock prices and expert net worth.

We can see this by considering versions of the financial frictions model which explicitly shut down the interaction between net worth and stock prices. The first version I consider is one where experts pay out all of their net worth as dividends each periods: $\sigma = 0$. This means that expert net worth each period is simply the net worth of new experts, $w_e z$, which is assumed proportional to productivity. In this case we can prove the following proposition:

Proposition 15. *If wages are proportional to productivity, the household has log utility, and experts pay out all of their net worth as dividends each period ($\sigma = 0$), then the model with financial frictions has an equilibrium where unemployment is constant in response to productivity shocks. The price of a match is proportional to current productivity: $Q(s) = \bar{Q}z$.*

This model features a moral hazard problem which restricts leverage, and incomplete markets since experts can only borrow risk free. However, we still recover the result that unemployment is constant. Why is this? As in the model without financial frictions, the match price being proportional to productivity allows vacancy posting to be constant since posting costs are also proportional. However, we now need to understand why the match price being proportional is allowable even with the financial friction. Several elements come together to make this possible. Firstly, experts optimally choose constant leverage in this model. This is because the experts excess return on lending, once discounted with log utility and consumption which is (in equilibrium) proportional to z, becomes constant. With constant leverage, net worth being proportional to productivity means that even though positive productivity shocks make experts richer, they end up spending this on the increased match price and posting costs, which are also proportional. In other words given the increase in asset prices, experts do not have any left over cash to spend on increasing total matches.

The key to the result is really that net worth is proportional to z, and leverage is constant. Since net worth is only as volatile as the productivity shock there is no financial accelerator. The same result emerges from the model with complete markets, even if $\sigma > 0$, because the contingency of debt leads agents to optimally make net worth proportional to z:

Proposition 16. *If wages are proportional to productivity, then the model with financial frictions and state contingent debt has an equilibrium where unemployment is constant in response to productivity shocks. The price of a match is proportional to current productivity:* $Q(s) = \bar{Q}z$.

In this model experts receive a higher, leveraged return when there are good productivity shocks, leading to the possibility that net worth is more volatile than productivity. However, the contracts they choose offset this in equilibrium, since they choose to structure their contingent claims to repay more in good states than bad, leading to net worth again only being as volatile as productivity.

The final result considers the case of incomplete markets with $\sigma > 0$ and shows that, in contrast to the cases above, it is not possible to generate an equilibrium with constant unemployment:

Proposition 17. *The model with financial frictions with* $\sigma > 0$ *does not have an equilibrium where unemployment is constant in response to productivity shocks when wages are proportional to productivity and utility is log.*

Proof. The proof of this proposition is a simple disproof: we conjecture that the model does have an equilibrium with constant unemployment and show this violates one of the equilibrium conditions. Constant unemployment requires constant market tightness, which requires, via (3.8), that Q is proportional to z: $Q(s) = \bar{Q}z$. As in the other models, this means that experts optimally choose constant leverage, $\bar{\phi}$. Denote by $\bar{K}$ the constant level of employment, and (3.15) requires that $N = \bar{Q}\bar{K}z/\bar{\phi}$.

In other words, for experts to have the right amount of net worth, on aggregate, to purchase the stock of matches requires that net worth be proportional to productivity. However, we can easily show that this leads to a contradiction since debt is not state contingent. Plugging in $Q(s) = \bar{Q}z$ to the equation governing total expert net worth, (3.14), gives:

$$N = \sigma \left(\left(z - \bar{w}z + (1 - \rho_x)\bar{Q}z \right) \bar{K} - D \right) + (1 - \sigma)\, w_e z$$

This is not proportional to z due to the fixed stock of debt, D, which does not vary with z. Hence it cannot be the case that $N(s) = \bar{N}z$. □

Since the model without financial frictions and the model with complete markets have zero volatility of unemployment over the cycle, and we have proven that the incomplete markets model must have positive volatility, I have proven that the incomplete markets model is more volatile in this special case. I have also shown the crucial role of net worth in this mechanism. Of course, one should be sceptical of analytical results derived from special cases, and this is certainly true here. In general we know that financial frictions can deliver either amplification or dampening, as I discussed in the previous section. To this end, I present calibrated numerical results in the next section.

3.7 Numerical results

In this section I present perturbation numerical result to analyse the quantitative significance of the ideas presented in the precious sections. In particular, I calibrate the model to assess whether it is able to match key features of the data. To test the model, I will compare its ability to generate volatility in market tightness and unemployment to the data, once the model is calibrated to match other moments of the data. I will be interested in comparing the ability of the models with and without financial fractions to generate volatility in unemployment. Thus the calibration of the financial frictions parameters will be important as they will determine how powerful financial frictions can be in a quantitative sense. For this, I take an approach similar in spirit to Winkler (2015). He chooses certain parameters of his model in order to match properties of asset prices in the data, and I do the same here. In particular, I will choose the parameters governing financial frictions to match certain asset price moments.

Another key issue in assessing the quantitative performance of my model is the current controversy over how to calibrate wages in the search and matching model. As discussed further below, the average level of wages is important in determining the volatility of unemployment. This is true in the baseline search and matching model, and is also true in my extension with financial frictions. I thus perform robustness checks for different values of this parameter, as well as various financial frictions parameters. I solve the model using first order log-linearisation in Dynare. The model is solved and simulated at a monthly frequency, and I take simple averages to compute quarterly statistics.

3.7.1 Data moments

In this section I describe the data I aim to test my model against. Table 3.1 presents the covariances and autocorrelations for seven key US time series. The data is quarterly, and covers the period 1951Q4 to 2014Q2. All data are seasonally adjusted, logged, and HP-filtered with smoothing parameter 10^5. Any data which is collected with monthly frequency are converted to quarterly figures by a simple average.

Non-financial moments

My measure of unemployment, u, is the Civilian Unemployment Rate in percent from the Current Population survey. Vacancies, v, is the composite Help Wanted Index of

Barnichon (2010), available from the author's website.[3] Market tightness, θ, is calculated as the ratio of the Help Wanted Index and Total Unemployment (thousands) from the Current Population Survey. Real wages, w, are calculated as total labour compensation per employee from the national accounts. To measure this I first construct the labour share (as detailed in my second chapter) and then measure wages per employee as the labour share multiplied by output over employment. Output, y, is chained real GDP taken from Line 1 of Table 1.1.6 of the National Income and Product Accounts from the Bureau of Economic Analysis. Labour productivity, z, is my measure of output divided by total employment. Total employment is measured as Total Nonfarm Employees from the Current Employment Statistics survey.

Constructing a measure of asset prices

Finally I need to construct a measure of asset prices. Ideally, this should be as close as possible to the definition of the asset price in my model, Q, which is the price of the entire equity stake in a firm with a single worker. One issue that arises is in the treatment of firm assets in the data, which the model abstracts from. I will discuss this in more detail below. I use two measures of equity prices, and I opt to measure in both the data and my model the quantity Q/z, which is thus the price of equity in a single worker firm scaled by labour productivity. This can be conveniently measured in the data as the ratio of the total real value of equity in the economy to real GDP.[4] I use two different measures of the total nominal value of equity. The first is the closing price of the S&P 500 index, collected from Yahoo finance.[5] The second is a measure of total market capitalisation of the US economy from the Flow of Funds accounts.

[3] At the time of writing, the data is available at `https://sites.google.com/site/regisbarnichon/research`

[4] To see this, note that $TE/y = (TE/n)/(y/n) = \tilde{Q}/z$ where TE is a measure of total equity value, n is employment, and $\tilde{Q}$ is the average equity value per worker in the data.

[5] Series S&P500 (ˆGSPC). At the time of writing this is available at `https://finance.yahoo.com/q/hp?s=%5EGSPC+Historical+Prices`

I use Nonfinancial Corporate Business; Corporate Equities; Liability.[6] Both nominal values of equity are deflated using the GDP deflator, taken from Line 1 of Table 1.1.4 of the NIPA accounts. I report the results only for my second measure of equity, but the moments of the two series are remarkably similar: the log standard deviations of the HP-filtered market capitalisation and S&P 500 series are 0.1516 and 0.1511 respectively. Their similarity is heartening, especially since one measure contains financial firms and one does not, and one might think that it would be appropriate to strip out financial firms from my measure of equity given that I split out experts from the rest of the economy in my model. Stripping out the value of firm assets from the data is more challenging, and I do not undertake this task here. Instead I choose to use the raw measures of asset prices as my primary data, and discuss the effects and challenges of attempting to split out the value of firm assets from the data in my robustness section.

Table 3.1: Data moments

	u	v	θ	w	y	z	Q/z
Standard deviation	0.195	0.188	0.371	0.016	0.025	0.015	0.152
Autocorrelation	0.947	0.941	0.948	0.930	0.939	0.907	0.845
Correlation	1	-0.889	-0.973	-0.237	-0.864	-0.193	-0.293
	–	1	0.970	0.097	0.803	0.202	0.270
	–	–	1	0.177	0.856	0.211	0.297
	–	–	–	1	0.526	0.687	0.088
	–	–	–	–	1	0.511	0.217
	–	–	–	–	–	1	0.142
	–	–	–	–	–	–	1

All data are quarterly, logged and then HP-filtered with smoothing parameter 10^5.

[6] Series id: FL103164103.Q. Note that when using this measure, my measure Q/z is actually the Market Capitalisation to GDP ratio popularised by Warren Buffett.

3.7.2 Baseline calibration

In order to test my model I first calibrate the parameters. Some parameters are calibrated to steady state targets, while others are calibrated to match certain moments of the data. After solving the model I compute moments which are comparable to the moments calculated in the data. Specifically, I simulate the model for a length of time equal to the length of my data one hundred times, and calculate the means and standard deviations of each moment across the repetitions. A summary table of calibration tables can be found in the appendix.

Starting with the household, I choose a standard CRRA utility function, $u(c) = c^{1-\sigma_c}/(1-\sigma_c)$, and specialise to log utility. I choose the discount factor $\beta = 0.9966$ to match an annual risk free rate of 4.17%. The parameters of the productivity process are chosen so that, once log HP-filtered, the means of the standard deviation and autocorrelation of the log HP-filtered series match the data I presented in the previous section. I choose $\sigma_e = 0.0043$ to match the standard deviation of 0.0146 in the data, and $\rho_z = 0.98975$ to match the autocorrelation of 0.9068. I normalise steady state productivity, $\bar{z}$ to one.

The labour market is parameterised following the calibration of Den Haan & Kaltenbrunner (2009), who report data giving a monthly job finding probability of $\lambda_w = 45.4\%$, vacancy filling probability of $\lambda_f = 33.8\%$ and unemployment rate of $u_s s = 5.7\%$. This allows me to pin down steady state tightness as $\theta_{ss} = \lambda_w/\lambda_f$. I assume a standard value of $\psi_1 = 0.5$ (Petrongolo & Pissarides, 2001) for the matching function elasticity. This allows me to pin down match efficiency as $\psi_0 = \lambda_f \theta_{ss}^{\psi_1} = 0.3917$. The job separation rate is picked to equate the flows of workers in and out of unemployment in steady state, giving $\rho_x = \lambda_w u_{ss}/(1-u_{ss}) = 0.0274$.

Real wage flexibility is set to $\gamma = 0.7$ following the empirical discussion in Michaillat (2012). This corresponds to an elasticity of wages to productivity of 0.7, consistent with the empirical evidence from job movers of Haefke, Sonntag & Van Rens (2007). I also set the steady state real wage, $\bar{w}$ following Michaillat (2012). Based on empirical estimates, he requires that the steady state recruiting cost, κ, is equal to 0.32 of a worker's steady state wage. This allows me to jointly solve for $\bar{w}$ and κ from equations (3.8) and (3.10), given a value of the steady state return on matches, $r_{k,ss}$, which I detail below. This gives values $\bar{w} = 0.9709$ and $\kappa = 0.3107$.

The expert parameters are calibrated to match asset price moments. There are three parameters to choose: the fraction of experts who survive each period, σ, the fraction of assets the experts can steal, Λ, and the equity injections given to new experts, w_e. These are jointly chosen to match three asset pricing moments: the equity premium and the standard deviation and autocorrelation of asset prices. I target the values for the standard deviation and autocorrelation of asset prices from the data in Table 3.1. For the equity premium I instead target a value lower than that found in the data, targeting a 1% premium of yearly equity returns over the risk free rate in steady state. This value corresponds to the premium in Gertler & Karadi (2011). The presence of financial frictions means that my model generates an equity premium even in the non-stochastic steady state. The equity premium in the data presumably reflects this wedge, as well as compensation for risk. Since there is no risk in my non stochastic steady state I do not want to attribute this part of the data to this moment, and hence choose a lower value. I use a numerical minimisation routine to find the values of the parameters which achieve these values of the moments, leading me to choose $\Lambda = 0.4854$, $\sigma = 0.9770$, and $w_e = 0.3026$. These values correspond to experts surviving 3.62 years on average, and having steady state leverage of 2.23. The expert sector pays out a fraction $1 - \sigma = 0.0230$ of its net worth as equity per month.

The model without financial frictions is calibrated using the same procedure as above, but without the financial frictions components. This means that there is a slight difference difference in the calibrated values of $\bar{w}$ and κ between the two models. The procedure to choose the values of these two parameters is exactly the same as for the financial frictions model, imposing $r^k_{ss} = 1/\beta$. For the complete markets model I take the values of all of the parameters as the calibrated values from the incomplete markets model. Note that I am thus not calibrating the three models to the same targets: the incomplete markets model is calibrated to match asset price moments, whereas the other two models are not calibrated to match these moments. In this sense I am not providing a test across the three models. I am only testing the ability of the incomplete markets model to match the volatility of unemployment once it is properly calibrated. The other two models are not tested, and their solutions are only provided to serve as references against which to compare the incomplete markets model.

3.7.3 Model evaluation

Moments

Table 3.2 reports moments calculated from simulating the incomplete markets model which are comparable to the empirical moments presented in Table 3.1. I simulate the model 100 times for a length of time equal to the length of the data sample, and calculate the same moments I calculated in the data for each of these samples. I then report the means and standard deviations (in parentheses) of these moments. The volatility and autocorrelation of productivity and asset prices is the same as the data by construction since these were calibrated to fit the data.

Given that the model was not calibrated to match the volatility of unemployment, the performance is surprisingly good. The model generates an average standard deviation of labour market tightness of 0.307, which is 82% of that observed in the data. Similarly for unemployment, the model is able to generate 72% of the volatility observed in the data. In the current calibration this represents a significant improvement over the model without financial frictions, which only delivers 61% of the volatility of tightness from the data. Moments for the model without financial frictions are available in Table 3.3 in the appendix.

Table 3.2: Simulated moments of the incomplete markets financial frictions model

	u	v	θ	w	y	z	Q/z
Standard deviation	0.138	0.180	0.303	0.010	0.023	0.015	0.152
	(0.017)	(0.017)	(0.036)	(0.002)	(0.003)	(0.002)	(0.018)
Autocorrelation	0.890	0.709	0.845	0.907	0.914	0.907	0.845
	(0.025)	(0.056)	(0.034)	(0.026)	(0.023)	(0.026)	(0.034)
Correlation	1	-0.812	-0.938	-0.952	-0.980	-0.952	-0.938
		(0.001)	(0.001)	(0.001)	(1×10^{-4})	(0.001)	(0.001)
	–	1	0.964	0.879	0.864	0.879	0.964
			(3×10^{-4})	(0.001)	(0.001)	(0.001)	(3×10^{-4})
	–	–	1	0.956	0.960	0.956	1.000
				(0.001)	(0.001)	(0.001)	(0)
	–	–	–	1	0.994	1.000	0.956
					(2×10^{-4})	(0)	(0.001)
	–	–	–	–	1	0.994	0.960
						(2×10^{-4})	(0.001)
	–	–	–	–	–	1	0.956
							(0.001)
	–	–	–	–	–	–	1

I simulate the model 100 times for 526 months, corresponding to the length of the data sample. I convert the data to quarterly frequency as per Michaillat (2012). All data are quarterly, logged and then HP-filtered with smoothing parameter 10^5. *The numbers in parentheses are the standard errors of the sample statistics.*

My calibration features a relatively high (on average 97% of productivity) and sticky wage which explains why the model without financial frictions performs relatively well compared to Shimer's (2005) calibration. However, as I discuss in my robustness, given my calibration strategy this does not actually impact on the ability of my model with financial frictions to match the data significantly. The moments and IRFs for the complete markets model are not reported since they are very similar to the model without financial frictions. For example, the average standard deviation of tightness is 0.2255 in the model without financial frictions, and 0.2212 in the complete markets model. Where the model does not perform well is the correlation of stock prices with real variables. The correlations in the model are all very high, whereas in the data they tend to be much lower. This likely reflects the absence of other shocks in the model which could introduce independent volatility into both variables.

Impulse response functions

Figure 3.2 plots the impulse response functions of both the financial frictions model and the standard model for comparison. That financial frictions give amplification is clear from the tightness and unemployment panels: the peak response of tightness is almost twice as high in the financial frictions model. The impulse responses reveal the mechanisms behind the model's financial accelerator, which can be traced out as follows. Recall that experts are financial constrained, and would like to fund extra vacancies on the margin but cannot afford to. A positive productivity shock allows them to do so, providing an initial increase in vacancies, and hence an increase in market tightness. An increase in market tightness makes it harder for firms to hire workers, which increases the value of being an existing firm, hence pushing up the stock price of existing matches. This is exactly the arbitrage argument implied by equation (3.8). Since existing matches are owned by the expert sector this pushes up experts' net worth, which allows them to fund even more vacancies, and the cycle

Figure 3.2: Impulse responses to a one standard deviation innovation to productivity

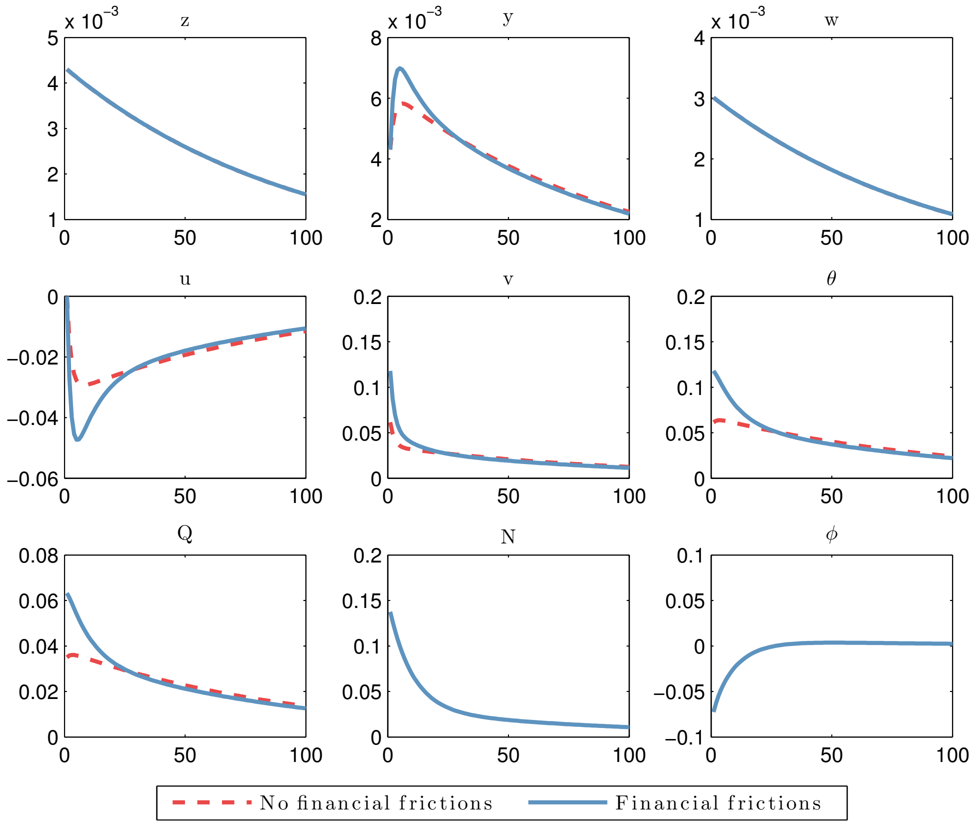

Impulse responses are log deviations from steady state. One period corresponds to one month.

continues.

3.7.4 Robustness

Robustness to alternative parameterisations

In this section I perform robustness checks on several parameters. This serves firstly as a check on my results, but also highlights the key role of asset price volatility in my model. Recall that arbitrage between old matches and vacancies implied the following relationship between the volatilities of market tightness and asset prices:

$$\sigma\left(\log\theta\right)=\frac{1}{\psi_1}\sigma\left(\log\frac{Q}{z}\right)$$

This implies that any model in which this arbitrage equation holds will generate the same standard deviation of tightness as long they share the same volatility of stock prices, and hence that any model calibrated to match the data on the volatility of asset prices will generate 82% of the volatility in tightness in the data, as does my baseline model. This has strong implications for my robustness checks, since it implies that if even if I change a parameter, as long as I recalibrate the financial sector to match the volatility of asset prices the model will still perform just as well at explaining the volatility in tightness.

I illustrate this by varying three parameters. I first explore the effects of varying the average level and stickiness of the real wage. These are key parameters for which there is still debate on how to calibrate. I also explore different values of the steady state equity premium. I do this because my calibration strategy involved only targeting an arbitrary fraction of the equity premium from the data, and I show below that the results are robust to alternative values. After perturbing these parameters, I report the solution to two variants of the model. The first variant holds all other parameters at their original values. In this solution, perturbing one parameter will thus affect the volatility of asset prices and tightness. In the second variant, I perturb the parameter but also adjust the financial sector parameters in order to maintain the volatility of asset prices at its original level.[7]

Figure 3.3 plots the results of this exercise for the average wage, $\bar{w}$. The dashed red line shows that, without recalibrating the financial parameters, a higher average wage pushes up the volatility of tightness, unemployment, and asset prices as we

[7]Specifically, I keep σ at its original level, and adjust Λ and w_e in order to maintain the volatility of asset prices and the steady state equity premium at their original levels. In principle I could also adjust σ so that I also maintain the autocorrelation of asset prices. This exercise is slightly more computationally demanding, and the results are similar.

Figure 3.3: Robustness: $\bar{w}$

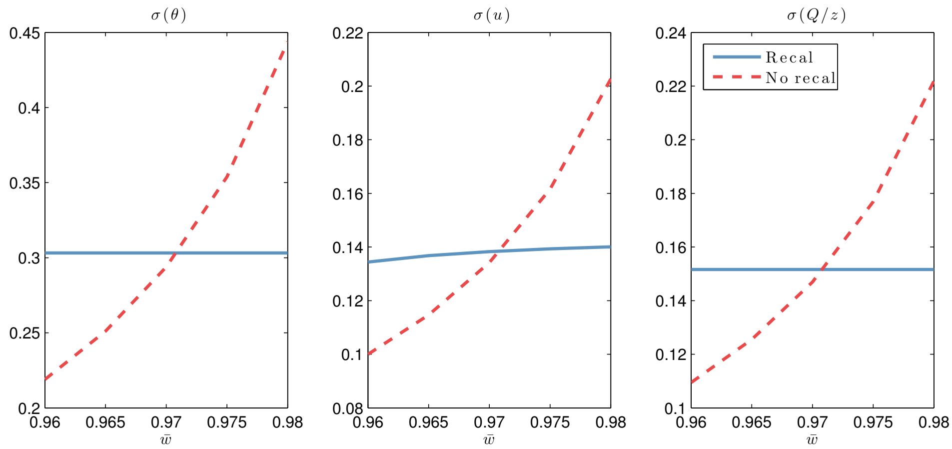

Standard deviations are computed as in the original model: these are the average standard deviations across 100 replications for different parameter values. Variables are logged and HP-filtered. The solid blue line gives volatilities when the financial variables are recalibrated, the dashed red line when they are not.

would expect. However, once we recalibrate to keep the volatility of asset prices constant, this effect disappears. The volatility of both asset prices and tightness remain constant at their original level, the former by construction and the latter due to the strong arbitrage arguments made above. The effect on unemployment volatility is virtually, but not entirely removed since unemployment is a stock and its volatility depends also on the volatility of its lagged values.

The story is the same for the changes in wage stickiness, γ, which are plotted in Figure 3.4. Increasing γ reduces wage stickiness and hence reduces the volatility of labour market variables. Interestingly, we can view this from a financial angle as well. Reducing wage stickiness reduces the responsiveness of profit to shocks, which reduces the sensitivity of net worth to shocks and hence dampens the financial accelerator. This exercise reveals an interesting interaction between wage stickiness and the financial accelerator: wage stickiness makes the financial accelerator more severe. We can see this by comparing the difference between the volatilities of tightness in the models

Figure 3.4: Robustness: γ

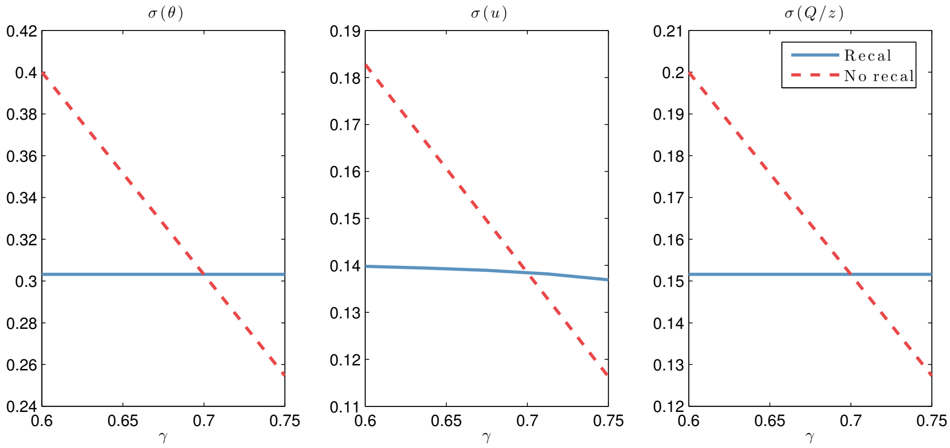

Standard deviations are computed as in the original model: these are the average standard deviations across 100 replications for different parameter values. Variables are logged and HP-filtered. The solid blue line gives volatilities when the financial variables are recalibrated, the dashed red line when they are not.

with and without financial frictions for different values of γ. With the baseline value of $\gamma = 0.7$ this difference is 0.078, but if stickiness is increased by setting $\gamma = 0.6$ the difference widens to 0.107, showing that increases in wage stickiness generate extra volatility in the financial frictions model above the model without financial frictions.[8] The same argument holds for increases in the average wage: these increase the gap between the volatility of the models with and without financial frictions. As before, once we recalibrate the financial parameters, the effects disappears for θ and Q/z, and is severely diminished for u.

Finally, I perform robustness for my assumed value of the steady state equity premium. In this exercise the recalibrated model is calibrated to give the original volatility of asset prices while matching the new equity premium. This exercise reveals

[8]This argument is similar to the argument of Schoefer (2015) within a financial accelerator context. He makes the argument within firms: sticky wages make firm cash-flow more volatile leading to volatility in the cash available for hiring. I make the argument that sticky wages make asset prices more volatile, further impacting the volatility of the net worth of experts.

Figure 3.5: Robustness: steady state r_k

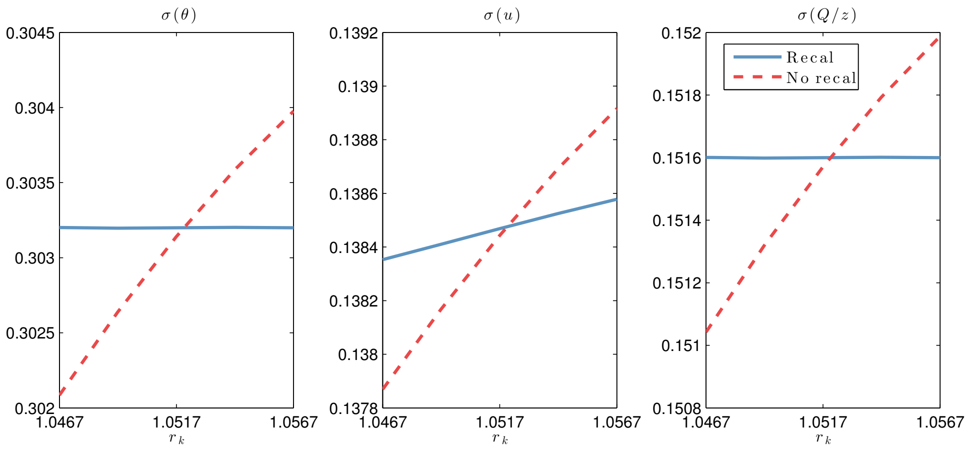

Standard deviations are computed as in the original model: these are the average standard deviations across 100 replications for different parameter values. Variables are logged and HP-filtered. The solid blue line gives volatilities when the financial variables are recalibrated, the dashed red line when they are not. The horizontal axis gives the steady state yearly value of r_k.

that my choice of equity premium target is not particularly important, and that varying the target within a one percentage point range has very small effects on the volatilities of asset prices and tightness even if I don't recalibrate the other financial parameters.

Robustness of asset price data

Checking the robustness of my measure of asset prices is an important exercise, since, as I pointed out above, the ability of my calibrated model to match the volatility of tightness depends crucially on the volatility of asset prices. I have already discussed the robustness of my data to the inclusion or exclusion of financial firms, and in this section I attempt to address another concern: the effect of controlling for the value of firms' assets.

Stripping out the value of firms' assets from the data is challenging. To see why it is important, remember that my model abstracts from capital, and firms' only assets are their relationships with employees. The productivity process is calibrated to therefore implicitly include capital movements, and firms don't own any capital. In the data the total equity value of firms would contain both the value of their relationships with workers and asset ownership, as well as any other sources of value such as tax shields or intangibles. Ultimately splitting out these various sources reliably is a challenging feat, and data limitations place bounds on our ability to do this. For example, Hall (2014) points out that attempts to do so lead to large periods when the stock market value of firms falls far below the measured value of firms' plants and equipment. These concerns aside, I make an attempt here to investigate how robust my results are to doing so. Since we would expect the value of firms' assets to be procyclical, due to both price and quantity effects, controlling for this could reduce the volatility of asset prices compared to the baseline measure.

Measuring the individual components of firm value is hard, so I first present an example showing how mismeasurement does not necessarily reduce my measure of the volatility of the worker-relationships component of firm value. One might expect this to be the case, especially if both relationship value and firms' asset values are positively correlated over the cycle: an increase in asset values should reduce the amount of an increase in total firm value we ascribe to worker relationships, and hence reduce its true volatility. However, since I am working with log volatilities this is not the case. Suppose we can cleanly split total firm equity value, E, into a component deriving from worker relationships, W, and a component deriving from the value of owned assets, K, giving $E = W + K$. I should calibrate my model to movements in W, but only measure E. To see how this mismeasurement need not reduce the true volatility of W, consider the special case where W and K are perfectly correlated such that we can write them both as loaded onto a common factor: $W = wX$ and $K = kX$. In this case we have that $\sigma(\log W) = \sigma(\log K) = \sigma(\log E) = \sigma(\log X)$. Since we would expect that both W and K should be procyclical over the business cycle, to the extent that they will be highly correlated the above example suggests that the log-volatility of total firm value should be a good proxy for the log-volatility of the value of worker relationships.

In the case where the different components of firm value are not perfectly correlated the above does not exactly hold and we need to try and measure the components individually. I attempt to strip out the value of firm assets from my measure of firm equity by subtracting the net worth of non-financial firms from the market value of their equity. Their net worth is a measure of the value of their assets net of their liabilities, and is thus a measure of the replacement cost of a firm. My measure of firm net worth, NW, is Nonfinancial Corporate Business; Net Worth, Level from the Flow of Funds.[9] If I simply subtract this from the market value of firm equity (which is the measure I use in the calibration, and is constructed for the same sample as net worth) then I run in to the same problem as indicated by Hall (2014): for much of the sample, this leaves negative value. Specifically for over 85% of the quarters in my sample. Regardless of whether this is correct or reflecting of measurement issues, this leaves me with the immediate problem that I cannot take logarithms of the adjusted series to compare its logged, HP-filtered volatility to my original series. Given this issue, my first check is to simply compare HP-filtered volatilities without taking logs. To do this I compute $\sigma(E/Y)$ and $\sigma(\tilde{E}/Y)$ where E refers to the market value of firm equity in the data, and $\tilde{E} = E - NW$ is the equity value series minus the net worth series. Recall that measuring this ratio gives a series comparable to Q/z in my model. If I do this, I find $\sigma(E/Y) = 0.1191$ and $\sigma(\tilde{E}/Y) = 0.1318$. In other words, rather than reducing the volatility of measured asset prices, as we would expect, the correction actually increases it slightly. This is surprising since we would expect net worth to be procyclical. From the 90s onwards this is certainly the case, however earlier in the sample the series (once HP-filtered) displays a slight countercyclicality, which could be due to measurement issues.

[9] Series id: FL102090005.Q

The second correction I do aims to take seriously the issue of the many negative values of firm value once the value of assets is stripped out. In particular, it is plausible that the way assets are market to market and valued makes the series not easily comparable. To try and adjust for this, I construct the series $\tilde{E}_\alpha = E - \alpha NW$ for different values of $\alpha \in (0, 1]$, the idea being that the scaling of the two series might not be comparable. $\alpha = 0$ corresponds to my original series, and $\alpha = 1$ corresponds to the corrected series above. This generates a minimum value of $\sigma(\tilde{E}_\alpha/Y) = 0.1191$ and a maximum of $\sigma(\tilde{E}_\alpha/Y) = 0.1318$. Additionally α can be chosen to match average values of the series. For example, $\alpha = 0.485$ generates a series for $\tilde{E}_\alpha$ which implies that on average 20% of total equity value derives from sources other than firms' assets, and generates a value of $\sigma(\tilde{E}_\alpha/Y) = 0.1217$. Overall, this exercise, while imperfect, does not immediately suggest that any large overstatement of the correct asset price for my model is induced by ignoring firms' assets.

Other sources are less easy to account for. Intangible capital and the value of tax shields are both hard to measure, and would require model based frameworks in order to estimate their contribution to firm equity value. Ultimately, the importance of addressing these concerns cannot be overstated, but I leave the exercise to future work.

3.8 Conclusion

In summary, I introduce financial frictions into the labour market matching model, and study interactions between the two frictions. I demonstrate a feedback between asset and labour markets which amplifies the model's response to exogenous shocks. Shocks which increase expert net worth allow experts to fund more vacancies, raising market tightness and lowering the ease with which firms can hire workers. This increases the value of being an existing firm, causing stock prices to appreciate. Since experts own firm stocks, this increases expert net worth further, amplifying the initial shock in a classic financial accelerator mechanism. I show how sticky wages, by making the stock market more volatile, amplify this financial accelerator, and how incomplete markets are required to generate the necessary volatility in expert net worth.

I derive an arbitrage equation in my model between equity prices and market tightness similar to the standard free entry condition. I show that as long as a matching model which shares this arbitrage condition is calibrated to match the volatility in asset prices in the data, it will always be able to generate 82% of the volatility in market tightness, and hence do a reasonable job at describing the volatility of the labour market. This is true in the standard matching model, and any variants where at least one agent is free to perform arbitrage between vacancies and existing matches.

This holds regardless of the underlying source of shocks or the fractions of the volatility caused by sticky wages or financial frictions. Does this mean that I am simply assuming the result by calibrating my model to match the volatility of asset prices? In a sense I am, although it is worth remembering that there is is no ex ante guarantee that calibrating to asset prices will make the search and matching model work well. Indeed, it is actually very good news for the matching model that one of its key equations, the free entry condition, holds up so well against the data.

However, the key limitation of this approach is that while I have shown that a model with financial frictions can do a good job at explaining the data, I have not presented any direct evidence that financial frictions are the *only* mechanism which can do so. One could imagine that augmenting the model instead with other mechanisms to introduce volatility into asset demand and hence asset prices would achieve the same end.

Learning, habit formation, and non-time separable preferences have all been shown to improve asset pricing behaviour, and could all potentially serve as alternative explanations to financial frictions. Ultimately more work is needed to help disentangle which of these forces is responsible for the volatility in asset prices. However, if my paper has shown anything it is that once the correct source has been identified, the matching model has the potential to utilise it to generate a meaningful fraction of the unemployment volatility in the data.

Appendices

3.A Equations and unknowns, incomplete markets model

3.A.1 Experts $[\phi, \nu, r_k, \lambda, K, D, N]$

Individual problem:

$$\nu(s) = \Lambda\phi(s)$$

$$\nu(s) = \mathrm{E}\left[\Omega(s', s)\left(1 - \sigma + \sigma\nu(s')\right)\left(\left(r_k(s', s) - r(s)\right)\phi(s) + r(s)\right)\middle| s\right]$$

$$r_k(s', s) \equiv \frac{z' - w(s') + (1 - \rho_x)Q(s')}{Q(s)}$$

$$Q(s) = \frac{\mathrm{E}\left[\Omega(s', s)\left(1 - \sigma + \sigma\nu(s')\right)\left(z' - w(s') + (1 - \rho_x)Q(s')\right)\middle| s\right]}{\mathrm{E}\left[\Omega(s', s)\left(1 - \sigma + \sigma\nu(s')\right)\middle| s\right] r(s) + \frac{\lambda(s)\Lambda}{1+\lambda(s)}}$$

Aggregation:

$$Q(s)K'(s) = \phi(s)N(s)$$

$$D'(s) = r(s)(\phi(s) - 1)N(s)$$

$$N(s) = \sigma\left(\left(z - w(s) + (1 - \rho_x)Q(s)\right)K - D\right) + (1 - \sigma)\, w_e z$$

3.A.2 Matching $[Q, \theta, v, u, m, q, w]$

$$\theta(s) \equiv v(s)/u(s)$$

$$m(s) = \psi_0 u(s)^{\psi_1} v(s)^{1-\psi_1} = \psi_0\theta(s)^{1-\psi_1} u$$

$$q(s) = m(s)/v(s) = \psi_0\theta(s)^{-\psi_1}$$

$$u(s) = 1 - K$$

$$Q(s) = \frac{\kappa z}{q(\theta(s))} = \frac{\kappa z}{\psi_0}\theta(s)^{\psi_1}$$

$$K'(s) = m(s) + (1 - \rho_x)K$$

$$w(s) = \bar{w}z^{\gamma}$$

3.A.3 Goods $[r, c]$

$$r(s) = \frac{u'(c(s))}{\beta \mathrm{E}[u'(c(s'))|s]}$$

$$zK = c(s) + \kappa z v(s) = c(s) + \kappa z \theta(s) u(s)$$

3.B Steady state equations, incomplete markets model

3.B.1 Experts $[\phi, \nu, r_k, \lambda, K, D, N]$

$$\nu = \Lambda \phi \tag{3.37}$$

$$\nu = \frac{\beta (1-\sigma) ((r_k - r) \phi + r)}{1 - \beta\sigma ((r_k - r) \phi + r)} \tag{3.38}$$

$$r_k \equiv \frac{z - w(z)}{Q} + 1 - \rho_x$$

$$Q = \frac{\beta (1 - \sigma + \sigma\nu) (z - w + (1-\rho_x)Q)}{\beta (1 - \sigma + \sigma\nu) r + \frac{\lambda\Lambda}{1+\lambda}}$$

$$QK = \phi N$$

$$D = r(\phi - 1)N$$

$$N = \sigma ((z - w(z) + (1-\rho_x)Q) K - D) + (1-\sigma) w_e z$$

3.B.2 Matching $[Q, \theta, v, u, m, q, w]$

$$\theta \equiv v/u$$

$$m = \psi_0 \theta^{1-\psi_1}(1 - K)$$

$$q = m/v = \psi_0 \theta^{-\psi_1}$$

$$K + u = 1$$

$$Q = \frac{\kappa z}{\psi_0}\theta^{\psi_1}$$

$$m = \rho_x K$$

$$w = \bar{w}$$

3.B.3 Goods $[r, c]$

$$r = 1/\beta$$

$$zK = c + \kappa z\theta(1 - K)$$

3.C Proofs

Proof of Lemma 2. We can set up the lagrangian:

$$\begin{aligned}\mathcal{L} = {} & \mathrm{E}\left[\Omega(s', s)\left((1-\sigma)n' + \sigma V(n', s')\right)\middle|\, s\right] \\ & + \lambda\left(\mathrm{E}\left[\Omega(s', s)\left((1-\sigma)n' + \sigma V(n', s')\right)\middle|\, s\right] - \Lambda\left(Q(s)k'_o + \frac{\kappa z}{q(\theta)}k'_n\right)\right) \\ & + \mu k'_n \end{aligned} \tag{3.39}$$

Where it is understood that any n' are replaced using (3.4). The FOCs wrt k'_n and k'_o are:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial k'_o} = {} & (1{+}\lambda)\mathrm{E}\left[\Omega(s', s)\left((1-\sigma) + \sigma V_1(n', s')\right)\left(z' - w(s') + (1-\rho_x)Q(s') - r(s)Q(s)\right)\middle|\, s\right] \\ & - \lambda\Lambda Q(s) = 0 \end{aligned} \tag{3.40}$$

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial k'_n} = {} & (1{+}\lambda)\mathrm{E}\left[\Omega(s', s)\left((1-\sigma) + \sigma V_1(n', s')\right)\left(z' - w(s') + (1-\rho_x)Q(s') - r(s)\frac{\kappa z}{q(\theta)}\right)\middle|\, s\right] \\ & - \lambda\Lambda\frac{\kappa z}{q(\theta)} + \mu = 0 \end{aligned} \tag{3.41}$$

Defining $\tilde{\Omega} \equiv \Omega(s', s)\left((1-\sigma) + \sigma V_1(n', s')\right)$ and rearranging:

$$\frac{\partial \mathcal{L}}{\partial k'_o} \Rightarrow Q(s) = \frac{(1+\lambda)\mathrm{E}\left[\tilde{\Omega}\left(z' - w(s') + (1-\rho_x)Q(s')\right)\Big|\, s\right]}{(1+\lambda)\mathrm{E}\left[\tilde{\Omega}\Big|\, s\right] r(s) + \lambda\Lambda} \tag{3.42}$$

$$\frac{\partial \mathcal{L}}{\partial k'_n} \Rightarrow \frac{\kappa z}{q(\theta)} = \frac{(1+\lambda)\mathrm{E}\left[\tilde{\Omega}\left(z' - w(s') + (1-\rho_x)Q(s')\right)\Big|\, s\right] + \mu}{(1+\lambda)\mathrm{E}\left[\tilde{\Omega}\Big|\, s\right] r(s) + \lambda\Lambda} \tag{3.43}$$

From this we see that if $\mu = 0$ then the FOCs require that:

$$Q(s) = \frac{\kappa z}{q\left(\theta(s)\right)} \tag{3.44}$$

Since μ was the multiplier on the non-negative vacancies constraint, this means that if experts are happy to post vacancies in equilibrium, then $Q = \kappa z/q$. Since aggregate vacancies are typically positive in the data I'll restrict attention to the region where $\mu = 0$. This allows us to impose the condition $Q = \kappa z/q$ and treat existing matches and vacancies as the same from the expert's point of view. Defining $k' \equiv k'_o + k'_n$, we can re-express (3.4) as

$$n' = \left(r_k(s', s) - r(s)\right) Q(s)k' + r(s)n \tag{3.45}$$

Where $r_k(s', s) \equiv (z' - w(s') + (1-\rho_x)Q(s'))/Q(s)$. The limited commitment constraint (3.5) becomes

$$\Lambda Q(s)k' \leq V^{\diamond}(n, s; k') \tag{3.46}$$

Where the new conditional value function is:

$$V^{\diamond}(n, s; k') = \mathrm{E}\left[\Omega(s', s)\left((1-\sigma)n' + \sigma V(n', s')\right)\right| s] \tag{3.47}$$

With n' replaced with the value implied by (3.45). Experts maximise value:

$$V(n,s) = \max_{k'} V^{\diamond}(n,s;k') \tag{3.48}$$

Subject to (3.45) and (3.46). I ignore the $v \geq 0$ constraint since I have assumed it is not binding. It is possible to show that the expert's problem is linear in n, which allows us to aggregate. If this is true the conditional value function is given by:

$$V^{\diamond}(n,s;k') = \mathrm{E}\left[\Omega(s',s)\left(1-\sigma+\sigma\nu(s')\right)n' \middle| \, s\right] \tag{3.49}$$

Where $V(n,s) = \nu(s)n$. We can define $\phi \equiv Q(s)k'/n$ and rewrite the flow BC as:

$$n' = \left(\left(r_k(s',s) - r(s)\right)\phi + r(s)\right)n \tag{3.50}$$

Thus n' is a function only of ϕ and n, and not k'. Hence we can rewrite the conditional value function as $\nu(s;\phi)n = V^{\diamond}(n,s;k')$:

$$\nu(s;\phi)n = \mathrm{E}\left[\Omega(s',s)\left(1-\sigma+\sigma\nu(s')\right)n' \middle| \, s\right] \tag{3.51}$$

Hence the overall maximisation can be written as

$$\nu(s)n = \max_{\phi} \nu(s;\phi)n \tag{3.52}$$

Subject to (3.9) and to

$$\Lambda Q(s)k' \leq \nu(s;\phi)n \Rightarrow \Lambda\phi \leq \nu(s;\phi) \tag{3.53}$$

This gives us the policy and value functions:

$$Q(s)k(s,n) = \phi(s)n$$

$$d(s, n) = (\phi(s) - 1)n$$

$$V(n, s) = \nu(s)n$$

$\square$

Proof of Lemma 3. Setting up the lagrangian:

$$L(n, s) = (1+\lambda)\mathrm{E}\left[\Omega(s', s)\left((1-\sigma)n' + \sigma V(n', s')\right)\middle| s\right] + \mu\left(n + \int_{s'} d(s')\,\mathrm{d}s' - Q(s)k'\right)$$
$$- \lambda\Lambda Q(s)k' \tag{3.54}$$

Where I am implicitly assuming that vacancies are positive, λ is the multiplier on the moral hazard constraint, μ is the multiplier on the balance sheet, and it is understood that all n' are replaced using (3.28). The first order condition with respect to a generic $d(s')$ gives:

$$(1+\lambda)p(s'|s)\Omega(s', s)\left(1 - \sigma + \sigma V_1(n', s')\right) r(s') = \mu \tag{3.55}$$

Using the consumer's first order condition, (3.29), to remove $r(s')$:

$$(1+\lambda)\left(1 - \sigma + \sigma V_1(n', s')\right) = \mu \tag{3.56}$$

Since λ and μ are common for all s', this shows that the expert chooses state contingent debt to equalise the marginal value of net worth, $V_1(n', s')$ across states next period. It is possible (proof omitted) to prove that as usual expert value is linear in net worth: $V(n, s) = \nu(s)n$, which implies that $V_1(n, s) = \nu(s)$. Combining this with (3.56), this implies that the value of $\nu(s')$ doesn't depend on next period's shock. Using this, we can guess and verify an equilibrium where $\nu(s)$ is constant across states and time: $\nu(s) = \bar{\nu}\ \forall s$. Notice that if the moral hazard constraint is always binding then this implies that leverage is also constant: $\bar{\phi} = \bar{\nu}/\Lambda$.

We can actually completely characterise asset prices and the labour market using the recursive definition of expert value, and the fact the ν and ϕ are constant. Using

(3.11), (3.27) and (3.28):

$$\bar{\nu} = \mathrm{E}\left[\Omega(s', s)\left(1 - \sigma + \sigma\bar{\nu}\right)\left(r_k(s', s)\bar{\phi} + (1 - \bar{\phi})r(s')\right)\middle|\, s\right] \tag{3.57}$$

Using (3.29) to remove $r(s')$ leaves:

$$\bar{\nu} = \left(1 - \sigma + \sigma\bar{\nu}\right)\mathrm{E}\left[\Omega(s', s)r_k(s', s)|\, s\right]\bar{\phi} + (1 - \bar{\phi})\left(1 - \sigma + \sigma\bar{\nu}\right) \tag{3.58}$$

Or:

$$\mathrm{E}\left[\Omega(s', s)r_k(s', s)|\, s\right] = \frac{\bar{\nu} - (1 - \bar{\phi})\left(1 - \sigma + \sigma\bar{\nu}\right)}{\left(1 - \sigma + \sigma\bar{\nu}\right)\bar{\phi}} \tag{3.59}$$

Notice that this is exactly the same as the definition of job value in the standard economy, except for the term on the right hand side, which would be one in that case. Notice as well that the term on the right hand side is constant over states and time, hence the interpretation is that the complete markets model is similar to the model without financial frictions, apart from a steady state "wedge".

It is easy to verify that the other equilibrium conditions are satisfied, verifying our guess. The model only has two state variables, z and K, as in the standard model. Net worth is not a state variable because of the state contingent contracts. Net worth is calculated as the required net worth for experts to purchase the capital stock, and we back out the required past debt choice each period to make this hold. □

Proof of Proposition 12. $R(Q) = \phi N$. I prove that $R(Q)$ is decreasing in Q by showing that both ϕ and N are. First note that Q only affects ϕ and N via r_k, and that r_k is decreasing in Q. ϕ is increasing in r_k because higher r_k increases expert value, allowing higher ϕ via $\phi = \nu/\Lambda$.[10] Hence ϕ is decreasing in Q. Steady state N is given by:

$$N = \frac{(1 - \sigma)w_e z}{1 - \sigma(r_k - r)\phi - \sigma r}$$

[10]We can show that ϕ is increasing r_k by considering perturbations to (3.37) and (3.38). Increasing r_k increases ν in (3.38), which allows higher ϕ in (3.37) which feeds back into higher ν in (3.38) and so on.

Steady state net worth is increasing in r_k as long as $N > 0$, and increasing in ϕ as long as $r_k > r$, which is required in a steady state where the financial friction binds. Since ϕ is also increasing in r_k, N is increasing in r_k and hence decreasing in Q. Therefore, $R(Q)$ is decreasing in Q. $\square$

Proof of Proposition 14. The model equations are satisfied by constant values for u, θ, m, q, and K, and values for c, w, and Q which are proportional to z. In particular, the free entry condition

$$Q(s) = \frac{\kappa z}{\psi_0} \theta(s)^{\psi_1}$$

is satisfied by constant θ and $Q = \bar{Q}z$. The recursive value of a match, with the definition of r_k plugged in, becomes

$$\mathrm{E}\left[\beta \frac{c(s)}{c(s')} \frac{z' - \bar{w}z' + (1-\rho_x)Q(s')}{Q(s)} \middle| s \right] = 1$$

which is satisfied by $Q(s) = \bar{Q}z$ and $c = \bar{c}z$. $\square$

Proof of Proposition 15. The model equations are satisfied by constant values for u, θ, m, q, ϕ, v and K, and values for c, w, and Q which are proportional to z. To see this, if we guess a constant value for ϕ then (3.13) implies a constant value for ν. We can verify that this satisfies the recursion for expert value with the definitions of Ω and r_k substituted in, and r replaced using (3.23):

$$\nu(s) = \mathrm{E}\left[\beta \frac{c(s)}{c(s')} (1 - \sigma + \sigma\nu(s')) \left(\frac{z' - \bar{w}z' + (1-\rho_x)Q(s')}{Q(s)} \phi(s) + (1-\phi) \frac{c(s')}{\beta c(s)} \right) \middle| s \right]$$

This is satisfied for any z if ϕ and ν are constant and $Q(s) = \bar{Q}z$ and $c(s) = \bar{c}z$. If $\sigma = 0$ then (3.14) implies that N is proportional to z: $N(s) = w_e z$. This implies constant employment from (3.15):

$$K'(s) = \frac{\phi N(s)}{Q(s)}$$

Both N and Q are proportional to z, leaving K' constant. The value of D' at any time can be backed out from (3.16). $\square$

Proof of Proposition 16. In the proof of Lemma 3 I proved that there exists an equilibrium with constant leverage and expert value. I now show that under the assumptions of Proposition 16 the rest of the model equations can be satisfied with constant employment. As previously discussed, constant employment requires constant tightness, which implies that $Q(s) = \bar{Q}z$ via (3.8). This satisfies the main complete markets equation, (3.33) with the definitions of Λ and r_k substituted in:

$$\mathrm{E}\left[\beta \frac{c(s)}{c(s')} \frac{z' - \bar{w}z' + (1-\rho_x)Q(s')}{Q(s)} \middle| s\right] = \frac{\bar{\nu} - (1-\bar{\phi})(1-\sigma+\sigma\bar{\nu})}{(1-\sigma+\sigma\bar{\nu})\,\bar{\phi}}$$

This is satisfied with $Q(s) = \bar{Q}z$ and $c(s) = \bar{c}z$. (3.15) requires that $N(s) = \bar{N}z$, which is always feasible for any path of shocks by picking the right sequence of contingent claims. $\square$

3.D Alternative setup: competitive match producing firms

The model in the paper is equivalent to a model where experts only trade in completed matches, and there exists a perfectly competitive "match producing sector". The match producing sector pays vacancy posting costs and sells any completed matches on the spot market to experts. The match producing sector's problem is thus a static profit maximisation problem. Profit is given by:

$$\pi = Q(s)k'_n - \kappa z v$$

Where $k'_n = q(\theta(s))v$ is the number of successful matches a match producing firm produces if it posts v vacancies. Plugging this in and taking the FOC with respect to v (or imposing zero profit) gives us the arbitrage equation from the main model, $Q(s) = \frac{\kappa z}{q(\theta(s))}$. The boundary case with no vacancy posting is also supported here, since match producing firms must produce positive vacancies.

3.E Figures and tables

Table 3.3: Simulated moments of the model without financial frictions

	u	v	θ	w	y	z	Q/z
Standard deviation	0.104	0.127	0.226	0.010	0.021	0.015	0.113
	(0.016)	(0.017)	(0.034)	(0.002)	(0.003)	(0.002)	(0.017)
Autocorrelation	0.936	0.831	0.913	0.907	0.927	0.907	0.913
	(0.018)	(0.045)	(0.025)	(0.026)	(0.021)	(0.026)	(0.025)
Correlation	1	-0.886	-0.965	-0.959	-0.980	-0.959	-0.965
		(0.007)	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)
	–	1	0.977	0.981	0.961	0.981	0.977
			(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
	–	–	1	1.000	0.998	1.000	1.000
				($1e^{-5}$)	($1e^{-4}$)	($1e^{-5}$)	(0)
	–	–	–	1	0.996	1.000	1.000
					($2e^{-4}$)	($1e^{-16}$)	($1e^{-5}$)
	–	–	–	–	1	0.996	0.998
						($2e^{-4}$)	($1e^{-4}$)
	–	–	–	–	–	1	1.000
							($1e^{-5}$)
	–	–	–	–	–	–	1

I simulate the model 100 times for 526 months, corresponding to the length of the data sample. I convert the data to quarterly frequency as per Michaillat (2012). All data are quarterly, logged and then HP-filtered with smoothing parameter 10^5. *The numbers in parentheses are the standard errors of the sample statistics.*

Table 3.4: Calibration

	Interpretation	Value	Source
β	Discount factor	0.9966	4.17% annual interest rate
σ_c	Risk aversion	1	Standard value
ρ_x	Job destruction	0.0274	Steady state transition probabilities
ψ_0	Match efficiency	0.3917	Steady state transition probabilities
ψ_1	Matching elasticity	0.5	Petrongolo & Pissarides (2001)
κ	Recruiting cost	0.3107	$0.32 \times$ steady state wage
$\bar{w}$	Steady state real wage	0.9709	Steady state unemployment
γ	Real wage flexibility	0.7	Michaillat (2012)
σ	Expert exit prob.	0.9770	Asset price moments
w_e	New expert equity	0.3026	Asset price moments
Λ	Fraction of divertable capital	0.4854	Asset price moments
$\bar{z}$	Steady state productivity	1	Normalisation
ρ	Autocorrelation of productivity	0.98975	Quarterly (log HP-filtered) autocorrelation
σ_e	Standard deviation of ε	0.0043	Quarterly (log HP-filtered) std.

Bibliography

[1] AAF (American Action Forum), 2015. *The Growth Consequences of Dodd-Frank.* Available at http://americanactionforum.org/uploads/files/research/The_Growth_Consequences_of_Dodd_5-6-15.pdf. May 6, 2015.

[2] Acemoglu, Daron, 2002. *Directed Technical Change.* The Review of Economic Studies, Vol. 69, No. 4 (Oct., 2002), pp. 781-809.

[3] Acharya, Viral V, Tanju Yorulmazer, 2007. *Too many to fail – An analysis of time-inconsistency in bank closure policies.* Journal of Financial Intermediation, Volume 16, p1-31.

[4] Adjemian, Stephane, Houtan Bastani, Michel Juillard, Frederic Karame, Ferhat Mihoubi, George Perendia, Johannes Pfeifer, Marco Ratto & Sebastien Villemot (2011). *Dynare: Reference Manual, Version 4.* Dynare Working Papers, 1, CEPREMAP.

[5] Admati, Anat & Hellwig, Martin, 2013. *The Bankers' New Clothes: What's Wrong With Banking and What to Do About It.* Princeton University Press, February 24, 2013.

[6] Arellano, Cristina, Bai, Yan, & Kehoe, Patrick J, 2012. *Financial Frictions and Fluctuations in Volatility.* Federal Reserve Bank of Minneapolis Research Department Staff Report 466. July 2012.

[7] Bank for International Settlements, 2010. *An assessment of the long-term economic impact of stronger capital and liquidity requirements.* Basel Committee on Banking Supervision, August 2010.

[8] Barnichon, Regis, 2010. *Building a composite Help-Wanted Index.* Economics Letters, Elsevier, vol. 109(3), pages 175-178, December.

[9] Bernanke, Ben, 2012. *The Federal Reserve's Response to the Financial Crisis, Lecture 3.* George Washington University School of Business, March 27, 2012. Available at http://www.federalreserve.gov/newsevents/files/chairman-bernanke-lecture3-20120327.pdf. Last accessed: 21/4/2015.

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

[11] Blundell, Richard, Crawford, Claire & Jin, Wenchao, 2014. *What Can Wages and Employment Tell Us about the UK's Productivity Puzzle?* The Economic Journal, May 2014. Volume 124, Issue 576, Pages 377-407.

[12] Brunnermeier, Markus K, and Yuliy Sannikov, 2014a. *A Macroeconomic Model With A Financial Sector.* American Economic Review 104.2: 379-421.

[13] Brunnermeier, Markus K, and Yuliy Sannikov, 2014b. *International Credit Flows And Pecuniary Externalities.* Working paper.

[14] Buera, Francisco, Fattal-Jaef, Roberto, & Shin, Yongseok, 2015. *Anatomy of a Credit Crunch: From Capital to Labor Markets.* Review of Economic Dynamics, Elsevier for the Society for Economic Dynamics, vol. 18(1), pages 101-117, January.

[15] Caggese, Andrea & Vicente Cunat, 2008. *Financing Constraints and Fixed-term Employment Contracts.* Economic Journal, Royal Economic Society, vol. 118(533), pages 2013-2046, November.

[16] Carlstrom, Charles T & Fuerst, Timothy S, 1997. *Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis.* American Economic Review, American Economic Association, vol. 87(5), pages 893-910, December.

[17] Carlstrom, Charles T, Fuerst, Timothy S, Ortiz, Alberto, Paustian, Matthias, 2014. *Estimating Contract Indexation in a Financial Accelerator Model.* Journal of Economic Dynamics and Control, Volume 46, September 2014, Pages 130-149.

[18] Chadha, Jagjit S., & Warren, James 2013. *Accounting For The Great Recession In The UK: Real Business Cycles And Financial Frictions.* Manchester School, University of Manchester, vol. 81, pages 43-64, October.

[19] Chari, V. V., Kehoe, Patrick J., & McGrattan, Ellen R., 2007. *Business Cycle Accounting.* Econometrica, Econometric Society, vol. 75(3), pages 781-836, 05.

[20] Cheng, Ing-Haw, Sahil Raina, and Wei Xiong, 2014. *Wall Street and the Housing Bubble.* American Economic Review, 104(9): 2797-2829.

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

[22] Claessens, Stijn, & M. Ayhan Kose, 2013. *Financial Crises: Explanations, Types, and Implications.* IMF Working Paper WP/13/28. January 2013.

[23] Daly, Mary C. & Hobijn, Bart, 2014. *Downward Nominal Wage Rigidities Bend the Phillips Curve.* Federal Reserve Bank of San Francisco Working Paper 2013-08, January 2014.

[24] Daly, Mary C., Hobijn, Bart & Wiles, Theodore S., 2012. *Dissecting Aggregate Real Wage Fluctuations: Individual Wage Growth and the Composition Effect.* Federal Reserve Bank of San Francisco Working Paper Series. Working Paper 2011-23, May 2012.

[25] Den Haan, Wouter J. & Kaltenbrunner, Georg, 2009. *Anticipated growth and business cycles in matching models.* Journal of Monetary Economics, Elsevier, vol. 56(3), pages 309-327, April.

[26] Den Haan, Wouter, Rendahl, Pontus, Riegler, Markus, 2015. *Unemployment (Fears), Precautionary Savings, and Aggregate Demand.* Mimeo.

[27] Diamond, D., and Dybvig, P., 1983. *Bank Runs, Deposit Insurance, and Liquidity.* Journal of Political Economy 91, 401-419.

[28] Dmitriev, Mikhail & Hoddenbagh, Jonathan, 2013. *The Financial Accelerator & the Optimal Lending Contract.* Mimeo.

[29] Druant, M., Fabiani, S., Kezdi, Gabor, Lamo, Ana, Martins, Fernando, & Sabbatini, R., 2009. *How are Firms Wages and Prices Linked: Survey Evidence in Europe.* Working Papers w200918, Banco de Portugal, Economics and Research Department.

[30] Duchin, Ran, Denis Sosyura, 2014. *Safer Ratios, Riskier Portfolios: Banks' Response to Government Aid.* Journal of Financial Economics (JFE), 113, 1-28.

[31] Dwenger, Nadja,. Fossen, Frank M, & Simmler, Martin, 2015. *From financial to real economic crisis: Evidence from individual firmbank relationships in Germany.* Mimeo.

[32] Eggertsson, Gauti B. & Paul Krugman, 2012. *Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach.* The Quarterly Journal of Economics, Oxford University Press, vol. 127(3), pages 1469-1513.

[33] Ennis, Huberto and Keister, Todd, 2003. *Economic Growth, Liquidity and Bank Runs.* Federal Reserve Bank of Richmond Working Paper No. 03-01.

[34] Farhi, Emmanuel, and Jean Tirole. 2012. *Collective Moral Hazard, Maturity Mismatch, and Systemic Bailouts.* American Economic Review, 102(1): 60-93.

[35] Farmer, Roger, 2012a. *The Stock Market Crash of 2008 Caused the Great Recession.* 2012 Meeting Papers 145, Society for Economic Dynamics.

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

[37] Favilukis, Jack & Xiaoji Lin, 2015. *Wage Rigidity: A Quantitative Solution to Several Asset Pricing Puzzles.* Mimeo, available at `https://sites.google.com/site/jackfavilukis/research`. Last accessed 27/03/2015.

[38] Financial Times, 2007. *Citigroup chief stays bullish on buy-outs.* Available at http://www.ft.com/cms/s/0/80e2987a-2e50-11dc-821c-0000779fd2ac.html#axzz3NknxFGvy. Published: July 9, 2007, 10:08 pm. Last accessed: 4/1/14.

[39] Financial Times, 2011. *Osborne to set out bank reform plans.* Available at http://www.ft.com/cms/s/0/54e40b28-0fb5-11e1-a468-00144feabdc0.html#axzz3h5ueHe1y. Published: November 15, 2011, 7:38pm. Last accessed: 27/7/15.

[40] Foote, Christopher L, Kristopher S. Gerardi and Paul S. Willen, 2012. *Why Did So Many People Make So Many Ex Post Bad Decisions? The Causes of the Foreclosure Crisis.* Federal Reserve Bank of Atlanta Working Paper 2012-7.

[41] Galií, Jordi, 2011. *Unemployment Fluctuations and Stabilization Policies.* Cambridge, Mass. MIT Press.

[42] Gertler, Mark & Karadi, Peter, 2011. *A model of unconventional monetary policy.* Journal of Monetary Economics, Elsevier, vol. 58(1), pages 17-34, January.

[43] Gertler, Mark & Nobuhiro Kiyotaki, 2010. *Financial Intermediation and Credit Policy in Business Cycle Analysis.* Handbook of Monetary Economics, Volume 3, 2010, Pages 547-599.

[44] Gertler, Mark & Nobuhiro Kiyotaki, 2013. *Banking, Liquidity and Bank Runs in an Infinite-Horizon Economy.* NBER Working Papers 19129, National Bureau of Economic Research, Inc.

[45] Giroud, Xavier, Mueller, Holger M, 2015. *Firm Leverage and Unemployment during the Great Recession.* NBER Working Paper No. 21076. Issued in April 2015.

[46] Gregg, Paul, Machin, Stephen & Fernández-Salgado, Mariña, 2014. *Real Wages and Unemployment in the Big Squeeze.* The Economic Journal, May 2014. Volume 124, Issue 576, Pages 408-432.

[47] Guardian, 2014. *The next crisis could be here soon.* Economics blog, Larry Elliott. http://www.theguardian.com/business/economics-blog/2014/jan/20/next-financial-crisis-here-soon-history. Published: 20/01/2014. Last accessed: 11/10/14

[48] Guerrieri, Veronica and Lorenzoni, Guido, 2011. *Credit Crises, Precautionary Savings and the Liquidity Trap.* Working paper.

[49] Haefke, Christian, Marcus Sonntag, and Thijs van Rens. 2007. *Wage Rigidity and Job Creation.* Institute for the Study of Labor (IZA) Discussion Paper 3714.

[50] Hagedorn, Marcus, Karahan, Fatih, Manovskii, Iourii, & Mitman, Kurt, 2015. *Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects.* NBER Working Paper 19499, January 2015.

[51] Hagedorn, Marcus & Iourii Manovskii, 2008. *The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited.* American Economic Review, American Economic Association, vol. 98(4), pages 1692-1706, September.

[52] Hagedorn, Marcus, & Manovskii, Iourii, 2015. *The Impact of Unemployment Benefit Extensions on Employment: The 2014 Employment Miracle?* NBER Working Paper 20884, January 2015.

[53] Hall, Robert, 2014. *High Discounts & High Unemployment.* Mimeo.

[54] Hanlon, W. Walker, 2015. *Necessity Is the Mother of Invention: Input Supplies and Directed Technical Change.* Econometrica, Volume 83, Issue 1, pages 67-100, January 2015.

[55] Imbs, Jean, 2007. *Growth and volatility.* Journal of Monetary Economics, Volume 54, Issue 7, October 2007, Pages 1848-1862.

[56] Jermann, Urban, and Vincenzo Quadrini. 2012. *Macroeconomic Effects of Financial Shocks.* American Economic Review, 102(1): 238-71.

[57] Karabarbounis, Loukas, 2014. *The Labor Wedge: MRS vs. MPN.* Review of Economic Dynamics, Elsevier for the Society for Economic Dynamics, vol. 17(2), pages 206-223, April.

[58] Kaur, Supreet, 2014. *Nominal Wage Rigidity in Village Labor Markets.* Mimeo.

[59] Khan, Aubhik, Julia K. Thomas, 2013. *Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity.* Journal of Political Economy, University of Chicago Press, vol. 121(6), pages 1055 - 1107.

[60] Kilenthong, Weerachart T, & Townsend, Robert M, 2014. *A Market Based Solution to Price Externalities: A Generalized Framework.* NBER Working Paper No. 20275.

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

[62] Kocherlakota, Narayana R. 2009. *Bursting Bubbles: Consequences and Cures.* Manuscript, Federal Reserve Bank of Minneapolis.

[63] Koudijs, Peter and Hans-Joachim Voth, 2014. *Leverage and Beliefs: Personal Experience and Risk Taking in Margin Lending.* NBER Working Paper No. 19957, March 2014.

[64] KPMG, 2013. *Moving On: The Scope for Better Regulation.* Available at `https://www.kpmg.com/Global/en/IssuesAndInsights/ArticlesPublications/better-regulation-in-banking/Documents/better-regulation-in-banking.pdf`. May, 2013.

[65] Mailath, George J, Loretta J. Mester, 1994. *A Positive Analysis of Bank Closure.* Journal of Financial Intermediation, Volume 3, Issue 3, June 1994, Pages 272-299.

[66] Martin, Alberto and Jaume Ventura, 2014. *Managing Credit Bubbles.* NBER Working Paper No. 19960, March 2014.

[67] McDonald, Robert & Paulson, Anna (2015). *AIG In Hindsight* NBER Working Paper No. 21108, April 2015.

[68] Mian, Atif, & Sufi, Amir, 2014. *What Explains the 20072009 Drop in Employment?* Econometrica, Volume 82, Issue 6, pages 21972223, November 2014.

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

[70] Mulligan, Casey, 2015. *Fiscal Policies and the Prices of Labor: A Comparison of the U.K. and U.S.* NBER Working Paper No. 21358, July 2015.

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

[72] Ohanian, Lee E. 2010. *The Economic Crisis from a Neoclassical Perspective.* Journal of Economic Perspectives, 24(4): 45-66.

[73] Olivei, Giovanni & Tenreyro, Silvana, 2007. *The Timing of Monetary Policy Shocks.* American Economic Review, American Economic Association, vol. 97(3), pages 636-663, June.

[74] Olivei, Giovanni & Tenreyro, Silvana, 2010. *Wage-setting patterns and monetary policy: International evidence.* Journal of Monetary Economics, Elsevier, vol. 57(7), pages 785-802, October.

[75] Osilia, Una Okonkwo, Anna Paulson, 2014. *Crises and confidence: Systemic banking crises and depositor behavior.* Journal of Financial Economics, Volume 111, Issue 3, March 2014, Pages 646-660.

[76] Perri, Fabrizio, & Quadrini, Vincenzo , 2014. *International Recessions.* Working paper.

[77] Pessoa, Joo Paulo & Van Reenen, John, 2014. *The UK Productivity and Jobs Puzzle: Does the Answer Lie in Wage Flexibility?* The Economic Journal, May 2014. Volume 124, Issue 576, Pages 289-699.

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

[79] Petrosky-Nadeau, Nicolas, 2013. *TFP during a credit crunch.* Journal of Economic Theory, Elsevier, vol. 148(3), pages 1150-1178.

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

[81] Petrosky-Nadeau, Nicolas, Kuehn, Lars-Alexander & Zhang, Lu. *Endogenous Disasters and Asset Prices.* Charles A. Dice Center Working Paper No. 2012-1.

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

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

[84] Ramey, Garey and Valerie A. Ramey, 1995. *Cross-Country Evidence on the Link Between Volatility and Growth.* The American Economic Review, Vol. 85, No. 5 (Dec., 1995), pp. 1138-1151.

[85] Reinhart, Carmen, & Rogoff, Kenneth, 2009. *This Time is Different. Eight Centuries of Financial Folly.* Princeton University Press.

[86] Reinhart, Carmen, & Rogoff, Kenneth, 2012. *Sorry, U.S. Recoveries Really Aren't Different.* Bloomberg. http://www.bloomberg.com/news/articles/2012-10-15/sorry-u-s-recoveries-really-aren-t-different. Published: 10/15/2012. Last accessed: 14/05/2015.

[87] Romer, Paul M, 1986. *Increasing Returns and Long-run Growth.* Journal of Political Economy, University of Chicago Press, vol. 94(5), pages 1002-37, October.

[88] Schmitt-Grohé, Stephanie & Uribe, Martin, 2015. *Downward Nominal Wage Rigidity, Currency Pegs, and Involuntary Unemployment.* Mimeo.

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

[90] Schroth, Enrique, Gustavo A. Suarez, Lucian A. Taylor, 2014. *Dynamic debt runs and financial fragility: Evidence from the 2007 ABCP crisis.* Journal of Financial Economics, Volume 112, Issue 2, May 2014, Pages 164-189.

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

[92] Siemer, Michael, 2014. *Firm Entry and Employment Dynamics in the Great Recession.* Finance and Economics Discussion Series 2014-56, Board of Governors of the Federal Reserve System.

[93] Wall Street Journal, 2012. *Sowing Seeds of the Next Major Crisis.* Current Account, Francesco Guerrera. http://online.wsj.com/news/articles/SB10001424052702304818404577347641050572260. Published: 16/04/2012. Last accessed: 11/10/14.

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