ESSAYS ON THE MACROECONOMICS OF THE GREAT RECESSION

Gregory Thwaites
LONDON SCHOOL OF ECONOMICS

THESIS SUBMITTED FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY
April 4, 2016

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.

Statement of conjoint work

I confirm that Chapter 2 was jointly co-authored with Professor Silvana Tenreyro of LSE and I contributed 80% of this work.

I confirm that Chapter 4 was jointly co-authored with Jeremy Franklin and May Rostom and I contributed 50% of this work.

Contents

List of Figures

List of Tables

Preface

I would like to thank my employer the Bank of England for enabling me to undertake my doctoral studies, my teachers at LSE for giving me the toolkit to conduct this research, and especially my supervisor Silvana Tenreyro for her patience and sage advice throughout. I would also like to thank her in her capacity as co-author of Chapter 2 of this thesis, and similarly to thank Jeremy Franklin and May Rostom for their collaboration on Chapter 4. The views expressed herein are mine and not necessarily those of the Bank of England and, with my coauthors where appropriate, I accept full responsibility for any errors it may contain.

I dedicate this thesis to the countless ordinary people whose lives were damaged by the Great Recession, ultimately by the mistakes of the economists and policymakers that they trusted to run things properly. We must do better in the future.

Greg Thwaites
LSE
February 2015

Chapter 1

Introduction

The Great Recession that began around the world in 2008 caused hardship for millions. It has also prompted economists to re-evaluate what they thought they knew about macroeconomics, and policymakers to question the same in regard to policy. This thesis contains three papers, each of which investigates a question prompted or made more salient by the events of the Great Recession.

Throughout the crisis and across the industrialised world, policymakers placed great emphasis on the ability of monetary policy to stimulate demand and close the large negative output gaps that the crisis opened up. This emphasis only increased as fiscal policy started to tighten after 2010. For this policy mix to work, however, it is necessary for monetary policy to have some effect when the economy is weak. The first substantive chapter of this thesis investigates whether this is the case. We estimate the impulse response of key US macro series to the monetary policy shocks identified by Romer and Romer (2004), allowing the response to depend flexibly on the state of the business cycle. We find strong evidence that the effects of monetary policy on real and nominal variables are more powerful in expansions than in recessions. The magnitude of the difference is particularly large in durables expenditure and business investment. The asymmetry relates to how fast the economy is growing, rather than to the level of resource utilisation. There is some evidence that fiscal policy has counteracted monetary policy more in recessions than in booms. We also find evidence that contractionary policy shocks have more powerful effects than expansionary shocks.

But contractionary shocks have not been more common in booms, so this asymmetry cannot explain our main finding.

The second paper also deals with interest rates, but with causes rather than consequences, and with long horizons rather than the business cycle. Over the past four decades, real interest rates have risen then fallen across the industrialised world. Over the same period, nominal investment rates are down, while house prices and household debt are up. In the second substantive chapter of this thesis, I explain these four trends with a fifth - the widespread fall in the relative price of investment goods. I present a simple closed-economy OLG model in which households finance retirement in part by selling claims on the corporate sector (capital goods) accumulated over their working lives. As capital goods prices fall, the interest rate must fall to reflect capital losses. And in the long run, a given quantity of saving buys more capital goods. This has ambiguous effects on interest rates in the long run: if the production function is inelastic, in line with most estimates in the literature, interest rates stay low even after relative prices have stopped falling. Lower interest rates reduce the user cost of housing, raising house prices and, given that housing is bought early in life, increasing household debt. I extend the model to allow for a heterogeneous bequest motive, and show that wealth inequality rises but consumption inequality falls. I test the model on cross-country data and find support for its assumptions and predictions. The analysis in this paper shows recent debates on macroeconomic imbalances and household and government indebtedness in a new light. In particular, low real interest rates may be the new normal. The debt of the young provides an alternative outlet for the retirement savings of the old; preventing the accumulation of debt, for example through macroprudential policy, leads to a bigger fall in interest rates.

Real interest rates have fallen still further since the onset of the financial crisis of 2008, which was also associated with falls in corporate lending, business investment, labour productivity and real wages in the United Kingdom. The third substantive chapter of this thesis uses a large firm-level dataset of UK companies and information on their pre-crisis lending relationships to identify the causal links from changes in credit supply to the real economy following the 2008 financial crisis. Controlling for demand in the product market and conditional upon survival, it finds that the contraction in credit supply reduced labour productivity, wages and the capital intensity of production at the firm level. Furthermore, firms experiencing adverse credit shocks were more likely to fail, other things equal. The paper shows that these effects are robust, statistically significant and economically large, but only when instruments based on pre-crisis banking relationships are used.

Taken together, what can we conclude from this work? Chapter 2 shows that nominal interest rates may need to be cut further than first thought in recessions, because a cut of a given size has less effect. The average level of nominal interest rates may therefore need to be higher to afford the monetary authorities more space to cut them. In extremis, monetary policy may not be sufficient to mitigate large recessions by itself, and may need to be supplemented by additional tools of demand management, such as fiscal policy. Chapter 3 suggests that the long-run real interest rate may have fallen for reasons independent of the business cycle, such that the inflation rate will need to be higher to achieve a given steady-state nominal interest rate. Macroprudential limits on household debt may exacerbate the tendency for real interest rates to fall, and it may be preferable instead to target a higher ratio of public debt to GDP. Between them, Chapters 2 and 3 provide arguments that the inflation target should be increased to keep the normal level of nominal interest rates a safe distance above zero. Chapter 4 suggests that adverse credit shocks may have very large negative effects on the productive efficiency of the corporate sector. This chapter attests to the large aggregate costs associated with financial crises, and accordingly indicates that measures to make them less frequent will be worthwhile even if these measures themselves entail appreciable gross costs. Taken together, these three papers suggest that there is much more to macroeconomic stabilisation than monetary policy, and that stabilisation is about the long run as well as the business cycle: fiscal and macroprudential policy should perhaps be used in conjunction with monetary policy to stabilise the business cycle, but furthermore to avoid a dichotomy between low real interest rates and high private debt, and stabilise the supply of credit.

Chapter 2

Pushing on a string: US monetary policy is less powerful in recessions

2.1 Introduction

Is monetary policy effective in recessions? In recent years this perennial question took centre stage in the public policy debate, as central banks in the United States and Europe faced the deepest post-war crisis. A priori, whether monetary policy is more powerful in recessions or expansions is unclear. Expenditure could be more or less sensitive to real interest rates at different points in the business cycle. Imperfections in the financial system might magnify or dampen the transmission of policy at different times. Prices might be more or less sticky. And the systematic component of monetary policy itself might behave differently. Previous work has studied this question, and adjacent ones, finding mixed results.

We investigate this question anew on US data, and find strong evidence that monetary policy shocks typically have much more powerful effects on output and inflation in an expansion than in a recession. In order to allow impulse response functions to depend on the state of the business cycle, we adapt the local projection method of Jordà (2005) and combine it with the smooth transition regression method of Granger and Terasvirta (1994).[1]

[1] Auerbach and Gorodnichenko (2011) and Ramey and Zubairy (2014) use a similar

We investigate the state-dependence of monetary policy impulse response functions in this framework, examining the response of a range of real and nominal variables to monetary policy shocks identified in the manner of Romer and Romer (2004).

The main result from our investigation is that shocks to the federal funds rate are more powerful in expansions than in recessions. Nearly all of the effect we observe on average in the data is attributable to the effect in good times, and in particular to the response of durable consumption and business and household investment. In an expansion, output and then inflation fall in response to a negative monetary shock in the textbook fashion. Within this, and in line with previous findings, business investment and consumer expenditure on durable goods and housing are substantially more sensitive than other expenditures, whereas the responses of durables and nondurables prices are much closer together. In a recession, in contrast, the response of output and inflation to monetary policy interventions is negligible and generally insignificantly different from zero. These differences are not attributable to differences in the amplification afforded by the response of credit prices or quantities. We find that contractionary shocks are more powerful than expansionary shocks - in line with Angrist, Jordà, and Kuersteiner (2013), who employ a different method. But given that they are equally common in both expansions and recessions, this cannot be the source of asymmetry across the business cycle. We study different indicators of the state of the economy and find that measures of the growth rate of activity such as GDP growth or NBER recession dates are the relevant determinants of monetary policy effectiveness, whereas measures of the level of resource utilization such as the output gap or the unemployment rate do not as clearly distinguish regimes. We find that fiscal policy seems to have counteracted monetary policy more strongly in recessions than in expansions, which provides one explanation for our results.

These findings are relevant for the design of stabilisation policy and the models used to analyse it. If changes in the policy rate have limited impact in a recession, central banks will be more likely to need to resort to other (unconventional) monetary policy measures to achieve the desired

procedure to study the effect of fiscal policy, though the method has never been applied to the analysis of monetary policy.

expansionary effect. Policymakers may also need to rely more heavily on fiscal or financial policies to stabilize the economy in a deep or protracted slump. On the modelling side, the findings call for macroeconomic models that generate a higher sensitivity in the response of the economy (and in particular, the durable-good sector) during expansions.

The remainder of this paper is structured as follows. Section 2.2 reviews the literature. Section 2.3 explains the empirical method and describes the dataset. Section 2.4 sets out the main results. We conduct sensitivity analysis in section 2.5. Section 2.6 concludes with some thoughts for future research.

2.2 Literature

There is a small empirical literature on how the impact of monetary policy varies with the business cycle, mostly written a decade or more ago. Previous research produced mixed results and, perhaps as a result, the mainstream monetary policy literature, both theoretical and empirical, has largely ignored the potential for asymmetries and their policy implications. See for example Christiano, Eichenbaum, and Evans (2005), Woodford (2011) and Gali (2008). Our paper makes use of important subsequent methodological innovations in the estimation of impulse response functions in regime-switching environments.

The closest paper to ours in terms of implementation is Weise (1999). Weise (1999) estimates regime-dependency with a smooth-transition technique (Granger and Terasvirta (1994)), as do we, but applies this to a VAR rather than a local projection model. The set of variables in the VAR is small: industrial production, consumer prices, and M1, detrended in complicated piecewise fashion over 1960Q2-1995Q2. Monetary shocks are identified with Choleski orthogonalisation, putting money last. The regime is indicated by the first lag of quarterly GDP growth, such that high-frequency shifts in regime are possible. As with other VAR-based regime-switching models (and in contrast to the local projection model we employ), the researcher must decide how to account for the possibility that a shock causes a shift in regime. In this case, impulse response functions are calculated as the difference between two stochastic simulations with different initial conditions for output.

Taken together, the results in this paper are difficult to interpret. In his linear model, a positive shock to the growth rate of M1 reduces output over a three-year horizon, against the weight of empirical evidence on this matter. The response of output in a high growth regime is similar to the linear model - i.e. a positive shock to money growth reduces output, whereas the response in a low-growth regime is almost nonexistent. The price level responds more positively in booms than in recessions. So the paper implies that monetary policy is virtually ineffective in a low-growth regime, and actually contractionary in a high-growth regime, a result that is hard to reconcile with the standard empirical result that, on average, monetary policy is expansionary.

Garcia (2002) studies the response of quarterly industrial production growth to monetary policy in the US from 1955:2 to 1993:1. The business cycle is identified with a two-state Markov switching regime and estimate

$$\Delta y_t - \mu_0 - S_t\mu_1 = \Sigma_{i=1}^{r}\phi_i\left(\Delta y_{t-i} - \mu_0 - S_{t-i}\mu_1\right) + \beta_{iq}X_{t-i} + S_{t-i}\beta_{ip}X_{t-i} + \epsilon_t$$

where X_t is the interest rate in period t and $S_t = 1$ if the economy is in an expansion at time t. The procedure strongly rejects the null[2] that monetary policy, measured either as the simple level of Fed Funds rate or as Choleski innovations to a standard three-variable VAR, is equally powerful in both regimes, in favour of the alternative that they are more powerful in recessions. This method assumes, among other things, that the intrinsic persistence and other stochastic properties of GDP are the same in booms and recessions. There is substantial evidence that this assumption does not hold (see, for example, Acemoglu and Scott (1997) and references therein).

Smets and Peersman (2001) study the response of quarterly industrial production growth to monetary policy in seven Euro-area countries. First, they identify the business cycle with a two-state Markov switching regime with fixed autoregressive coefficients but state-dependent means μ_{i,s_t} for each country i at time t in state s

[2]i.e.the hypothesis that $\Sigma_{i=1}^{r}\beta_{ip} = 0$ for $r = 4$

$$\Delta y_{i,t} - \mu_{i,s_t} = \phi_1 \left(\Delta y_{i,t-1} - \mu_{i,s_{t-1}}\right) + \phi_2 \left(\Delta y_{i,t-2} - \mu_{i,s_{t-2}}\right) + \epsilon_{i,t}$$

They then separately identify monetary policy shocks with a linear VAR and use the historical contribution to the time-t policy rate in this VAR as the measure of the shock. They add the first lag of monetary policy shocks (the contribution of historical shocks to the current interest rate) to the AR(2)

$$\Delta y_{i,t} - \mu_{i,s_t} = \phi_1 \left(\Delta y_{i,t-1} - \mu_{i,s_{t-1}}\right) + \phi_2 \left(\Delta y_{i,t-2} - \mu_{i,s_{t-2}}\right) + \beta_{s_{t-1}} MP_{t-1} + \epsilon_{i,t}$$

imposing that the state of the economy is the same across the countries in the sample. They find that β is more negative in recessions than in booms - essentially the opposite of our finding.

This method imposes strong assumptions on the dynamics of output. Firstly, it assumes that past monetary policy shocks can be aggregated across time in a linear model when the underlying environment may be nonlinear. Secondly, it assumes that the propagation of a given monetary shock (the ϕ coefficients) is the same in different regimes; in other words, all of the difference in the impact of monetary policy is apparent in the single β coefficient.

Lo and Piger (2005) estimate the following equation

$$\phi(L)\, y_t^T = \gamma_0(L)\, x_t + \gamma_0(L)\, x_t S_t + \epsilon_t$$

where y_t^T is the transitory component of log quarterly industrial production, and x_t is a monetary policy shock identified from a three-variable structural VAR. S_t is a two-state Markov-switching process, in which the probabilities of transition from boom to recession is a function of state variables z_t. The authors find that putting a constant and two lags of an NBER recession date indicator in z_t yields very strong evidence of asymmetry in the response of output to monetary policy. They calculate impulse response functions to a monetary policy shock in the four possible combinations of realisations of the state variable $\{S_t, S_{t+1}\}$ and find that monetary policy is most powerful when the economy is in a recession either in period t or $t+1$. Accordingly, in calculating the impulse response, they do not allow the future state of the economy to change, either exogenously or in response to a monetary policy shock. Given that the aim of the exercise is to assess the impact of monetary policy on output - the state variable - this approach is difficult to defend.

In results, though not in method, our paper is closer to Thoma (1994), who estimates a non-linear VAR in output and monetary variables, allowing some of the coefficients to depend linearly on the deviation of output growth from trend. Like us, he finds that monetary shocks (especially contractionary ones) have more powerful effects in expansions than recessions. Unlike the approach we follow, however, his approach requires the researcher to make a number of discretionary decisions on the econometric specification. In contrast to this and other papers discussed above—and importantly for understanding the transmission menchanism-our paper stresses the difference in the response during booms of durables and business investment on the one hand and non-durables on the other, a dimension glossed over in this literature.

In summary, the general form of empirical model employed in the studies above is

$$(y_t - \overline{y}_t) = \alpha(L)\left(y_{t-1} - \overline{y}_{t-1}\right) + \beta x_t + \varepsilon_t$$

where x is the policy shock shock and y is the set of outcome variables. These studies typically allow only a proper subset of $\{\alpha(L), \beta, \overline{y}\}$ to depend on the state of the cycle. They must also take a stand on how the policy shock alters the transmission between regimes.

In contrast to the methods used previously, a local projection model (Jordà (2005)) has a number of advantages relative to a VAR. First, it does not impose the dynamic restrictions implicit in a VAR - the true model can take any form. Secondly, one can economize on parameters and, in some circumstances, increase the available degrees of freedom. In particular, one loses observations from the need to use leads as dependent variables. But the number of variables on the right-hand side need only be enough to ensure that the shocks ε_t are exogenous; none are needed to describe the dynamics of the endogenous variable conditional on the shock. If the VAR representation involves a large number of variables and lags, the net result will be an increase in the available degrees of freedom.

Thirdly - and most importantly for the present study - with a regime-switching local projection model one does not need to take a stand on **how** the economy switches from one regime to another. More specifically, a regime-switching local projection model takes the form

$$y_{t+h} = F\left(z_t\right)\left(\beta_b^h x_t + \gamma_b' z_t\right) + \left(1 - F\left(z_t\right)\right)\left(\beta_r^h x_t + \gamma_r' z_t\right) + u_t$$

where $F(z_t)$ is an indicator of the regime. The coefficents β_h^j measure the average effect of a shock as a function of the state of the economy when the shock hits, and therefore encompasses the average effect of the shock on the future change in the economy's state. In contrast, when using a regime-switching VAR model, the impulse response of the VAR implicitly assumes no change in the state of the economy, an assumption that is difficult to defend when we are considering shocks with large real effects. Ramey and Zubairy (2014) finds that this can have an important bearing on the results when estimating the state-dependence of US fiscal policy. It may explain the difference between our findings and some of those in the previous literature on state-dependent monetary policy summarised above.

Overall the theoretical literature has not had much to say about the state-dependent impact of macroeconomic policy across the cycle. One notable exception is Vavra (2013), who in recent work argues that recessions are often characterized by high realized volatility, and thus frequent price changes, which leads to a steep Phillips curve and ineffective monetary policy. He estimates a New Keynesian Phillips Curve on US data and finds support for this hypothesis. Berger and Vavra (2012) simulate a model of durables expenditure in the presence of adjustment costs and show that durables purchases are less sensitive to subsidies when output is low. They also show that the conditional variance of an ARCH process describing durables expenditure is higher during booms than in recessions, suggesting that either aggregate shocks are larger in booms, or that durables expenditure is more sensitive to shocks of a given size. They supply additional evidence against the former possibility, suggesting that durables expenditure is more sensitive to aggregate shocks - including monetary shocks - during booms. Our findings support the implication of Berger and Vavra (2012)'s model that monetary policy interventions are more effective during expansions and that most of the effect results from the response of durables and business investment.

2.3 Econometric method

In this section we first set out the specification of the econometric model used in this study. Then we explain our approach to statistical inference. Finally we describe our data sources, our state variables, and our identified policy shocks.

2.3.1 Specification

Our econometric model closely resembles the smooth transition - local projection model (STLPM) employed in Auerbach and Gorodnichenko (2011) and Ramey and Zubairy (2014) to analyze fiscal policy. The impulse response of variable y_t at horizon $h \in \{0, H\}$ in state $j \in \{b, r\}$[3] to a shock ε_t is estimated as the coefficient β_h^j in the following regression

$$y_{t+h} = \tau t + \left(\alpha_h^b + \beta_h^b \varepsilon_t + \gamma^{b\prime} x_t\right) F(z_t) + \left(\alpha_h^r + \beta_h^r \varepsilon_t + \gamma^{r\prime} x_t\right)(1 - F(z_t)) + u_t \tag{2.1}$$

where τ is a linear time trend, α_h^j is a constant and x_t are controls.[4] $F(z_t)$ is a smooth increasing function of an indicator of the state of the economy z_t. Following Granger and Terasvirta (1994) we employ the logistic function

$$F(z_t) = \frac{\exp\left(-\theta \frac{(z_t - c)}{\sigma_z}\right)}{1 + \exp\left(-\theta \frac{(z_t - c)}{\sigma_z}\right)},$$

where c is a parameter that controls what proportion of the sample the economy spends in either state and σ_z is the standard deviation of the state variable z. The parameter θ determines how violently the economy switches from expansion to recession when z_t changes.

In this paper, for each variable we estimate the $H+1$ equations of the IRF at horizon $0, ..., H$ as a system of seemingly unrelated regression equations. By Kruskal's theorem, this yields the same point estimates of the regression

[3] b denotes an expansion, r denotes a recession

[4] In the baseline specification, $x_t = y_{t-1}$

coefficients as equation-by-equation OLS, because the explanatory variables are the same in each equation. But it enables us to calculate the distribution of functions of parameters at different horizons, such as the smoothed IRFs presented in the figures below.

2.3.2 Inference

We employ two different approaches to conducting inference on our estimated impulse response functions. In order to conduct inference on cumulative impulse responses, moving averages and other functions of response variables at different horizons, each of these approaches needs to calculate the correlation of parameter estimates between equations. The first is to calculate standard errors analytically, allowing for the possibility of serially correlated residuals within equations and across equations. To capture this we follow Ramey and Zubairy (2014) and use the Driscoll and Kraay (1998) method to adjust standard errors for the possibility of correlation in the residuals across dates t and horizons h. This amounts to estimating the parameters of the equations separately, as above, and then averaging the moment conditions across horizons h when calculating Newey-West standard errors. Following Jordà (2005), we set the maximum autocorrelation lag $L = h + 1$ where h is the maximum horizon of the impulse response function.

The second approach is to bootstrap the key statistics of interest, namely the sign of $\beta_h^b - \beta_h^r$. This will allow not only for various forms of dependence among the residuals, but will also account for the fact that we scale the IRFs with estimated parameters (i.e. the impact effect of a policy shock on policy rates). Montiel Olea, Stock, and Watson (2012) argue that proper inference should take account of this, but also that standard large-sample 2SLS statistics can be misleading. To perform the bootstrap we construct 10,000 samples with replacement of size T and calculate the fraction of cases in which our null hypothesis does not hold. We transform the resulting p-value into a t-statistic for comparability with the other measures.

Inference on the above families of t-statistics - $H + 1$ for each response variable - will generate a 'multiple testing problem': if we test n true null hypotheses at significance level α, we will on average reject αn of them. Methods such as Holm (1979) exist to deal with this issue. However, in our setting there are no strong **a priori** grounds for specifying at what horizon the effects of monetary policy shocks depend on the state of the business cycle, rendering these methods inapplicable in the present study. It turns out that the t-statistics we present in section 2.4 are strongly correlated at adjacent horizons, alleviating the practical concern of this problem. But to deal further with this concern, we calculate and do inference on cumulative impulse response functions at discrete horizons.

2.3.3 Data

We work predominantly with chain-linked US National Accounts data downloaded from the website of the Philadelphia Fed.[5] Where our aggregates do not correspond directly with published data, we construct our own approximations to the chain-linked aggregates with Tornqvist indices (Whelan (2000)). We work with log levels of volume indices, and quarterly annualised log differences of implied deflators.

Our sample period (after the effects of the leads and lags described below are taken into account) runs from (shocks occurring in) 1969Q1-2002Q4, with the response variables measured up to five years later. Our sample runs therefore over the four decades leading up to the collapse of Lehman brothers, but does not include the ensuing major financial crisis, when the impact of monetary policy could have been different to a 'normal' recession.[6]

2.3.4 The state variable and the shocks

We define z_t as a seven-quarter moving average of real quarterly GDP growth. Following Ramey and Zubairy (2014), and in contrast to Auerbach and Gorodnichenko (2011), our moving average term z_t is a lagging rather than centred moving average, so that future values of response variables do not appear on the right-hand side of the regression. Higher values of θ mean that $F(z_t)$ spends more time close to the $\{0, 1\}$ bounds of the process, moving the model closer to a discrete regime-switching setup. Smaller values of θ mean that more of the observations are taken to contain some information about behaviour in both regimes. We follow Auerbach and Gorodnichenko

[5]We use the latest vintage of the data rather than real-time estimates

[6]Using end-quarter data - i.e. the shock in the final month of the quarter - yielded qualitatively similar results to those below.

(2011) and set $\theta = 3$ to give an intermediate degree of intensity to the regime switching, and also follow them in defining a recession as the worst 20 per cent of the periods in our sample, setting c to make this so. The robustness of our results to each of these choices is investigated below.

Our monetary policy shocks ε_t are quarterly averages of the monetary policy shocks identified by Romer and Romer (2004), which we have extended through 2008 in the manner of Coibion, Gorodnichenko, Kueng, and Silvia (2012). For our baseline estimates of impulse responses, we use the same regression specification as Romer and Romer (2004). When estimated over a common sample, we replicate the results exactly; when extending the sample the regression coefficents and hence the residuals change slightly, such that our extended series has a correlation of .987 with the original. To test the robustness of our results, we also run a non-linear, state-dependent analogue of the regression, thereby allowing for the reaction function of the Fed to depend on the state of the economy. These are discussed in subection 2.4.3.

The original Romer shocks, the shocks identified with the same method on a larger sample, and the non-linearly identified shocks are all shown in Figure 2.1, which also shows our transformed state variable $F(z_t)$ at the baseline parameter values. The figure shows, inter alia, that the monetary policy shocks associated with the early part of Paul Volcker's Chairmanship of the Federal Reserve - the period of greatest variability in the shocks - took place on the whole at a time of relatively weak economic activity. Section 2.5 examines the robustness of our findings to alternative choices of z_t, θ (the intensity of regime-switching), c (the proportion of the sample we call a recession), the length and phase-shift of the moving average state variable, and the identification scheme of the monetary policy shocks.

2.4 Results

In this section of the paper, we first set out our baseline results. We then explore whether the asymmetry we find is due to a different pattern of shocks across the business cycle.

2.4.1 Baseline results

The first four columns of Figure 2.2 show the smoothed impulse responses of the volume of GDP, the inflation rate of the GDP deflator and the Federal Funds rate to an identified monetary-policy shock that generates an initial 1 percentage point rise in the Federal Funds Rate - i.e. h is on the x-axis, β^h is on the y-axis[7]. The first column displays the central estimate of the impulse response in expansions (dashed lines), recessions (dotted lines) and a linear model (solid lines, where we restrict the coefficient to be constant across regimes). The second to fourth columns display central tendencies and 90% confidence intervals for the linear model, expansions and recessions respectively. The charts in the fifth column represent our two estimates of the t-statistic of the null hypothesis that $\left(\beta_h^b - \beta_h^r\right)$, with the area between ± 1.65 shaded blue. So, for example, if the solid line in the fourth columns falls below the lower extreme of the blue area at some horizon h we can reject the null that the IRFs at that horizon are equal in favour of the alternative that they are more negative in expansions at a 5% significance level. The IRFs are scaled so that the shock results in a 1 percentage point increase in the Fed Funds rate in all three regimes.

Figure 2.2 shows that the linear model delivers a familiar picture. Following a contractionary monetary policy shock, the level of output starts to fall, reaching a minimum of about half a percent below baseline two to three years after the shock, before beginning to recover. The price level is initially sticky, but eventually falls by about one per cent, flattening off by the end of the horizon. The policy rate is persistent but reverts towards and eventually passes through the conditional mean.

The difference between expansions and recessions is seen most clearly in the left-hand column. Output responds almost an order of magnitude more strongly in an expansion than in a recession, with the maximum fall about 1 per cent in an expansion. The price level also falls much more sharply, by about 3 per cent in an expansion against 0.5 per cent in a recession. In a recession, the responses of output and prices are mostly statistically

[7] In this case 'smoothed' means three-period centered moving averages of the IRFs, except at the endpoints of the function. The standard errors of these moving averages are calculated taking account of the covariance between the estimates at different points estimated above.

insignificantly different from zero. In an expansion, the nominal policy rate falls sharply below the conditional mean about two years after the shock, perhaps because of the systematic component of policy responding to the contraction the previous shock has created. The path the policy rate takes in a recession is on average higher than in an expansion. It is therefore clear from the figures that the larger response of nominal and real variables in an expansion is not attributable to a bigger rise in long-term nominal interest rates.

Table 2.1 cumulates the impulse response functions for GDP and inflation, and shows two alternative estimates of how significant the difference between them is, as set out above - Driscoll-Kraay standard errors and bootstrapped significance tests. The cumulative effect of a monetary policy shock is significantly larger at standard levels, with the precise horizon depending on how the standard errors are calculated.

Figure 2.3 plots the impulse response of the volumes of three expenditure aggregates to the same shock as before. In line with the response of aggregate output, all the volume indices respond much more in an expansion than in a recession, with the difference already significant one quarter after the shock. The top row - corresponding to an index of durable household expenditure - responds roughly an order of magnitude more than nondurable consumption, both in an expansion and in the linear model. In a recession, the response of all three kinds of expenditure is insignificant. Figure 2.4 shows that the response of prices is also larger in an expansion, although the estimated differences are generally less significant.

Figure 2.5 plots the impulse responses of four other variables often implicated in the transmission of monetary policy shocks. The first two rows show the response of real government consumption and net tax revenues (as a share of GDP) respectively. The first row shows that there is weak evidence that real government consumption responds positively on average to a tightening of monetary policy. Why this should be so is not clear. One possibility is that spending is set in nominal terms, such that real spending increases because the price level falls. However, the behaviour of government consumption is response to monetary policy at different points in the cycle does not support this explanation. In an expansion, when the disinflationary effects of policy are at their strongest, there is no evidence that real government consumption increases in response to a monetary policy tightening, whereas there is a significant increase in a recession. It could be that the larger fall in the price level we see in an expansion raises the real burden of public debt, such that spending cuts become necessary. Whatever the reason for this asymmetry, to the extent that increases in government consumption are expansionary, they will be offsetting the contractionary impulse provided by monetary policy: government spending seems to have been 'working against' monetary policy during recessions but not during expansions.

There is weaker evidence for the same on the tax side. The second line of the figure shows that, after the first few quarters, the tax-GDP ratio seems to rise more sharply in response to a monetary tightening in an expansion than in a recession. This again may be due to the stronger response of the price level in an expansion, and its effect on the government's intertemporal budget constraint. To the extent that tax rises are contractionary, this will reinforce the effect of monetary policy. However, much of the government debt that is being revalued by the disinflation is held by the US private sector, who will therefore enjoy a positive wealth shock that will offset much of the extra taxation needed to service the increase in debt, offsetting the contractionary effect of tax rises somewhat.

The third row of the table shows a measure of the external finance premium - the Gilchrist-Zakrajsek bond spread (Gilchrist and Zakrajsek (2012a)). Monetary policy could be more powerful in a boom if the external finance premium is more increasing in interest rates in good times than in bad, such that the rates at which households and firms can borrow move by more than the policy rate suggests. However, there no evidence of an effect in this direction and by the end of the sample if anything the opposite appears to be the case: the external finance premium counteracts the effect of a monetary shock in an expansion. In a recession, the premium amplifies the shock. So the response of financing spreads cannot explain why policy is more powerful in a boom. The difference in the response - which would tend to generate an opposite result to the one we find for the impact of monetary policy on expenditure and prices - is not quite significant at standard levels.

2.4.2 The distribution of shocks in expansions and recessions

One possible explanation for these findings is that the response of the economy to monetary policy shocks is indeed nonlinear, but is not directly a function of the state of the economy. Rather, it is possible that policy shocks of different kinds are more common at certain times, and it is this that generates the apparent dependence of the IRF on the state of the business cycle. If, say, large or positive shocks are proportionally more powerful than small or negative shocks, and if they are more common in expansions than recessions, then an empirical model like ours that is linear in the shocks, conditional on the regime, would misleadingly uncover a larger IRF in expansions than in recessions.

Figure 2.6 and Table 2.2 shows IRFs for the state-independent model modified such that positive and negative shocks are allowed to have different effects. We plot $\left\{\beta_h^+, \beta_h^-\right\}, h \in \{0, H\}$ estimated from the following equation

$$y_{t+h} = \tau t + \alpha_h^b + \beta_h^+ \max\left[0, \varepsilon_t\right] + \beta_h^- \min\left[0, \varepsilon_t\right] + \gamma^{b\prime} x_t + u_t$$

and again scale β so that the shock raises the policy rate by 1 percentage point on impact. The figure shows that positive shocks (i.e. monetary tightenings) have a much larger impact on output than negative shocks. This finding is consistent with those in Cover (1992), Long, Summers, Mankiw, and Romer (1988) and Angrist et al. (2013). The effects of positive and negative shocks on inflation are much harder to distinguish, with the difference between them not significant at standard levels. The bottom row shows that contractionary shocks are substantially more persistent than expansionary shocks, hampering any reliable inferences about the effect of a given shift in the yield curve. However, the finding that contractionary shocks (monetary tightenings) appear to have a bigger impact on output, but not necessarily on inflation, than negative shocks is interesting in its own right. [8]

If positive shocks to the federal funds rate were more common in expansions than recessions, the results in Figure 2.6 might account for the finding

[8] We estimated another equation in which the impact of policy was allowed depend both on the sign of the shock and on the state of the economy when it hit - i.e. to take on four values at any given horizon. We did not find any consistent statistically significant evidence of non-linearities by the sign of the shock, but the precision of our estimates was low given the loss of degrees of freedom inherent in this procedure.

that policy tends to be more powerful in expansions than recessions. But no such regime-dependent pattern in the shocks exists. Figure 2.7 shows estimates of the pdf and the cdf of the shocks overall and depending on the state of the business cycle.[9] There is little difference between the central tendencies of the distributions of shocks in booms and recessions - positive shocks do not preponderate in booms.

The main difference between the two regimes, apparent in Figure 2.7, is that the distribution of shocks is more variable during recessions. If smaller shocks, which are more common in booms, are proportionally more powerful, this could also explain our finding of a larger average impact of shocks. To check this we estimated the following equation

$$y_{t+h} = \tau t + \alpha_h^b + \beta_h^s \varepsilon_t + \beta_h^l \varepsilon_t^3 + \gamma^{b\prime} x_t + u_t$$

i.e. adding the cubed value of the policy shock as an additional explanatory variable. If the coefficent β_h^l on this variable were significantly positive (negative), this would count as evidence that large shocks of either sign are more (less) powerful. The left-hand column of Figure 2.8 plots the functions $\beta_h^l, h \in \{1, H\}$ with associated 90% confidence intervals, while the right-hand column shows t-statistics associated with the null hypothesis that $\beta_h^l = 0$ for each of the variables. Table 2.3 shows estimates of the cumulative IRFs and signficance as above. The bottom row shows some evidence that larger shocks have tended to die out more quickly, which may explain why their negative effects on output and inflation are attenuated (top and middle rows). However, the third row of Figure 2.2 shows that a given impulse to interest rates is more persistent in a recession than in a boom. So it is unlikely that this is driving our main results.

In summary, positive shocks appear to be more powerful than negative shocks, but they are not more common in expansions than recessions. Larger shocks are more common in recessions than expansions, but the effect of a shock does not clearly depend more or less than proportionally on its size once the difference in its persistence has been accounted for. This suggests that differences across regimes in the distribution of the shocks, as opposed

[9]The linear estimate is the raw Romer shocks smoothed with a normally distributed kernel. The expansion and recession estimates are generated by weighting the kernel function with the $F(z_t)$ and $1 - F(z_t)$ respectively.

to differences across regimes in the response to a given shock, do not explain our key findings.

2.4.3 Non-linear Romer regression

The monetary policy shocks used in this paper are identified as the residuals from an estimated reaction function. If this reaction function is itself state-dependent, our results may arise from state-dependent measurement error in the estimated monetary policy shocks, rather than any state dependency in the true IRFs. To examine this possibility we estimate a smooth transition analogue of the original Romer regression and use the resulting shocks to examine the robustness of our results. To be precise, if we write the original regression as

$$\Delta FFR_t = \beta X_t + \varepsilon_t \tag{2.2}$$

where X are the control variables employed by Romer and Romer (2004) and the estimated residuals $\hat{\varepsilon}_t$ are the identified monetary policy shocks, then our state-dependent identification scheme is

$$FFR_t = F(z_t)\,\beta^b X_t + (1 - F(z_t))\,\beta^r X_t + \tilde{\varepsilon}_t \tag{2.3}$$

and $\hat{\tilde{\varepsilon}}_t$ are our non-linearly identified shocks. The resulting state-dependent policy shocks have a 0.902 correlation with the original series and 0.909 with our extension of it.

Figure 2.9 and Table 2.4 display the results of using the estimated $\tilde{\varepsilon}_t$ in equation (2.1). The second column in the figure and the solid blue line in the first column are state-independent IRFs identified using the original state-independent Romer regression and as such are identical to those in Figure 2.2. The third and fourth columns display the state-dependent IRFs to the shocks identified in equation (2.3). The results are qualitatively similar to the baseline approach, but somewhat less pronounced and less well-determined, such that the difference between the GDP IRFs is only marginally significant at standard levels, and the difference between the price level responses, while large, has a t-statistic around 1.

Figure 2.10 and Table 2.5 show the baseline IRFs calculated when ϵ_t are the structural shocks recovered from a VAR in the log-levels of GDP, the GDP deflator and the Federal Funds rate, with a Choleski identification scheme in which monetary policy is ordered last. The linear IRF is calculated using shocks from a linear VAR, whereas the non-linear IRFs employ shocks calculated with a non-linear VAR. The peak response of GDP is significantly larger in a boom, but the difference is generally nonmonotone. The difference in the inflation IRFs is similar in size and significance to the baseline case.

In summary, therefore, our baseline results do not seem to be driven by state-dependent measurement error in the estimated monetary policy shocks, although allowing for a state-dependent Romer reaction function widens the standard errors somewhat. Identifying our shocks with a state-dependent VAR produces similar results.

2.5 Sensitivity analysis

The following section examines the robustness of our findings to alternative choices of the state variable z_t (subsection 2.5.1), the phase shift of the state variable (subsection 2.5.2), θ (the intensity of regime-switching - subsection 2.5.3) and c (the proportion of the sample we call a recession - subsection 2.5.4). Using our preferred measure of the economic cycle, our results are qualitatively robust to reasonable alternatives along each of these margins.

2.5.1 The state variable

Our baseline results employ a measure of the economic cycle - a moving average of GDP growth - to which there are many reasonable alternatives. This subsection examines the sensitivity of our results to three of them.

Figure 2.11 and Table 2.6 show the response of our headline variables when Z_t is a moving average of a $[0, 1]$ indicator of recession defined as the proportion of the quarter in which the economy was in recession, as determined by the NBER. The differences across the cycle between the inflation responses are similar to our baseline results, and remain statistically significant at the 10 per cent level. The responses of output are somewhat less dissimilar across regimes than in the baseline setup.

Figure 2.12 and Table 2.7 contain the results of the same test but define Z_t as an HP-filtered output gap. An HP filter is already essentially a centered moving average of the level of GDP, so no further filtering or phase shifting is undertaken. The charts show that there is limited evidence of differences between the impulse responses in the different regimes for GDP, and no such evidence for inflation. Finally, figure 2.13 and Table 2.8 contain the results of an analogous exercise using a centered moving average of (the negative of) the unemployment rate. Once again, there is limited evidence of a bigger effect on output when the unemployment rate is low relative to when it is high, and no clear picture of any differences in the impact on inflation.

To sum up, the growth rate of the economy - measured as a moving average of either GDP growth or the NBER recession indicator - is the most reliable determinant of the effect of monetary policy shocks on output and inflation. Measures of the **level** of resource utilisation - the output gap and the unemployment rate - are much less informative about the impact of monetary policy.[10]

2.5.2 Phase shift of state variable

Figure 2.14 and Table 2.9 show the baseline IRFs calculated when, following Auerbach and Gorodnichenko (2011) rather than Ramey and Zubairy (2014), z_t is a centered rather than lagging moving average of output. The gap between booms and recessions shrinks somewhat and appears earlier in the case of GDP growth, but the broad picture remains for both output and inflation, and remains statistically signficant. So our results do not appear to be an artefact of using a centered moving average to calculate the state of the economy.

2.5.3 Intensity of regime switching (θ)

Tables 2.10 and 2.11 are analogues of table 2.1 but where we have set θ equal to 1 and 10 respectively. They show that the qualitative message of the earlier analysis is unchanged - our results are robust to reasonable changes in the intensity of regime switching.

[10]Using leads of the output gap and unemployment works better, given that they essentially cumulate past growth rates, but such leads do not describe the state of the economy as the shock hits

2.5.4 Proportion of sample in a recession (c)

Figure 2.15 and Table 2.12 show that the main qualitative conclusions are robust to increasing to 50 per cent the proportion of the sample judged to be more in a recession than in a boom. The response of inflation in a recession is now significantly negative, but still significantly smaller than that in a boom. Classifying a greater proportion of observations as a recession therefore does not overturn our main results.

2.6 Concluding remarks

We have found statistically strong evidence that standard measures of US monetary policy shocks have had more powerful effects on expenditure quantities and prices during economic expansions than during recessions. These findings are robust to several variations in the empirical model. They do not appear to be an artefact of different patterns in the shocks themselves, but rather the outcome of differences in the economic effects of a given shock at different points in the business cycle. We also find that monetary contractions are much more powerful than expansions. In other words, there is truth in the quote attributed to John Maynard Keynes that "you can't push on a string" - when the economy is weak, monetary policy can do little about it.

Standard estimates in the literature that do not allow for state-dependent impulse responses have masked these differential effects. The findings question the common wisdom that cuts in policy rates can stop or mitigate recessions, calling for the analysis of alternative policy measures during contractions. On the modelling side, the literature has hitherto focused on linear, state-independent models of monetary policy transmission. In contrast, these findings call for monetary models that generate a higher sensitivity in the response of durable goods during expansions, an asymmetry that has been largely glossed over in the theoretical literature.

Table 2.1: Cumulative impulse response of GDP and inflation: baseline specification

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0243	0.0050	0.0182	0.0344
	8	-0.0565	-0.0059	0.0652	0.0235
	12	-0.0939	-0.0110	0.0312	0.0038
	16	-0.0901	-0.0179	0.0894	0.0354
Inflation	4	0.0081	-0.0009	0.1958	0.6474
	8	-0.0100	0.0023	0.2480	0.3276
	12	-0.0601	0.0034	0.0058	0.0244
	16	-0.0973	-0.0070	0.0095	0.0112

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.2: Cumulative impulse response of GDP and inflation: positive and negative shocks

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0174	0.0101	0.0923	0.0175
	8	-0.0668	0.0153	0.0262	0.0020
	12	-0.1216	0.0269	0.0031	0.0014
	16	-0.1349	0.0288	0.0046	0.0068
Inflation	4	0.0030	0.0098	0.2871	0.2832
	8	-0.0157	0.0185	0.1721	0.1244
	12	-0.0409	0.0039	0.2368	0.1627
	16	-0.0677	-0.0141	0.2612	0.2522

This table shows the cumulative impulse response of a monetary policy shock on GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs for positive and negative shocks in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.3: Cumulative impulse response of GDP and inflation: small and large shocks

Cumulative impact on	Horizon	Coefficient	Significance level of difference	
			Driscoll-Kraay	Boostrap
GDP	4	0.0005	0.2793	0.4177
	8	0.0029	0.0649	0.2905
	12	0.0086	0.0005	0.2882
	16	0.0111	0.0000	0.2581
Inflation	4	-0.0011	0.1547	0.2880
	8	-0.0003	0.4463	0.3901
	12	0.0023	0.2496	0.2113
	16	0.0054	0.0518	0.0158

This table shows the cumulative impulse response to a cubed monetary policy shock ε_t^3 on GDP and inflation, i.e. β_h^l in the equation $y_{t+h} = \tau t + \alpha_h^b + \beta_h^s \varepsilon_t + \beta_h^l \varepsilon_t^3 + \gamma^{b\prime} x_t + u_t$ at horizons of 4, 8, 12 and 16 quarters, over and above the linear response to the shock, in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of cubed term is zero.

Table 2.4: Cumulative impulse response of GDP and inflation: non-linearly identified shocks

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0081	0.0017	0.1932	0.2427
	8	-0.0252	-0.0090	0.2310	0.2225
	12	-0.0550	-0.0128	0.0537	0.0759
	16	-0.0523	-0.0203	0.0952	0.2106
Inflation	4	0.0005	0.0014	0.4627	0.4650
	8	-0.0045	-0.0003	0.4211	0.4382
	12	-0.0315	-0.0024	0.1106	0.2024
	16	-0.0578	-0.0135	0.0668	0.1456

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.5: Cumulative impulse response of GDP and inflation: VAR shocks

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0055	-0.0063	0.4677	0.0431
	8	-0.0177	-0.0172	0.4907	0.0439
	12	-0.0391	-0.0170	0.2112	0.0104
	16	-0.0321	-0.0199	0.3333	0.1011
Inflation	4	0.0058	-0.0044	0.0804	0.4658
	8	-0.0003	-0.0075	0.2653	0.2912
	12	-0.0343	-0.0059	0.0707	0.0372
	16	-0.0631	-0.0034	0.0190	0.0050

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.6: Cumulative impulse response of GDP and inflation: NBER state variable

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0138	0.0053	0.1298	0.0846
	8	-0.0359	-0.0045	0.1287	0.0844
	12	-0.0575	-0.0122	0.1109	0.0683
	16	-0.0568	-0.0193	0.1717	0.1496
Inflation	4	0.0033	0.0051	0.4377	0.4740
	8	-0.0106	0.0114	0.1634	0.1845
	12	-0.0436	0.0118	0.0141	0.0184
	16	-0.0666	-0.0017	0.0186	0.0230

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.7: Cumulative impulse response of GDP and inflation: HP filtered output as state variable

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0085	0.0110	0.1433	0.1239
	8	-0.0237	0.0016	0.1923	0.1964
	12	-0.0363	0.0065	0.1052	0.1571
	16	-0.0354	-0.0152	0.2493	0.3636
Inflation	4	-0.0032	0.0127	0.0481	0.2387
	8	-0.0145	0.0200	0.0215	0.1334
	12	-0.0289	-0.0072	0.2263	0.2816
	16	-0.0386	-0.0439	0.4456	0.5155

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.8: Cumulative impulse response of GDP and inflation: unemployment state variable

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	0.0010	0.0287	0.3622	0.3254
	8	-0.0112	0.0748	0.2392	0.2725
	12	-0.0217	0.0393	0.3092	0.3657
	16	-0.0280	0.0673	0.1777	0.3009
Inflation	4	0.0034	0.0311	0.2610	0.3492
	8	0.0032	0.0216	0.4048	0.4148
	12	-0.0070	-0.0483	0.3207	0.5164
	16	-0.0193	-0.1529	0.0382	0.6407

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.9: Cumulative impulse response of GDP and inflation: phase shift in state variable

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0060	0.0068	0.1282	0.0988
	8	-0.0425	0.0061	0.0562	0.0274
	12	-0.0491	-0.0133	0.2610	0.1486
	16	-0.0369	-0.0324	0.4709	0.3810
Inflation	4	0.0066	0.0071	0.4901	0.4951
	8	0.0009	0.0133	0.3527	0.3611
	12	-0.0465	0.0200	0.0265	0.0333
	16	-0.0813	0.0134	0.0210	0.0062

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.10: Cumulative impulse response of GDP and inflation: $\theta = 1$

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0513	0.0098	0.0283	0.0664
	8	-0.1110	0.0054	0.0441	0.0524
	12	-0.1747	0.0055	0.0276	0.0410
	16	-0.1528	-0.0060	0.1016	0.0827
Inflation	4	0.0119	-0.0005	0.3058	0.5737
	8	-0.0239	0.0056	0.2385	0.3229
	12	-0.1265	0.0171	0.0079	0.0556
	16	-0.1870	0.0118	0.0130	0.0478

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.11: Cumulative impulse response of GDP and inflation: $\theta = 10$

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0149	0.0015	0.0522	0.0844
	8	-0.0376	-0.0117	0.1460	0.0655
	12	-0.0708	-0.0177	0.0503	0.0090
	16	-0.0731	-0.0219	0.0854	0.0387
Inflation	4	0.0108	-0.0012	0.0526	0.7608
	8	0.0013	-0.0011	0.4303	0.5136
	12	-0.0363	-0.0055	0.0433	0.1070
	16	-0.0704	-0.0171	0.0239	0.0431

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Table 2.12: Cumulative impulse response of GDP and inflation: c=0.5

Cumulative impact on	Horizon	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Boostrap
GDP	4	-0.0389	-0.0022	0.0823	0.0890
	8	-0.1018	-0.0159	0.0465	0.0654
	12	-0.1578	-0.0267	0.0374	0.0428
	16	-0.1402	-0.0327	0.1093	0.0969
Inflation	4	-0.0029	0.0041	0.3799	0.4235
	8	-0.0357	0.0030	0.1828	0.2473
	12	-0.1340	-0.0045	0.0207	0.0406
	16	-0.1958	-0.0177	0.0180	0.0363

This table shows the cumulative impulse response to a monetary policy shock of GDP and inflation at horizons of 4, 8, 12 and 16 quarters. The third and fourth columns show the values of the cumulative IRFs in the two regimes in units of log differences. The two right-hand columns show three measures of the p-value of the null hypothesis that the values of the two IRFs are the same.

Figure 2.1: Monetary policy shocks and the state of the economy

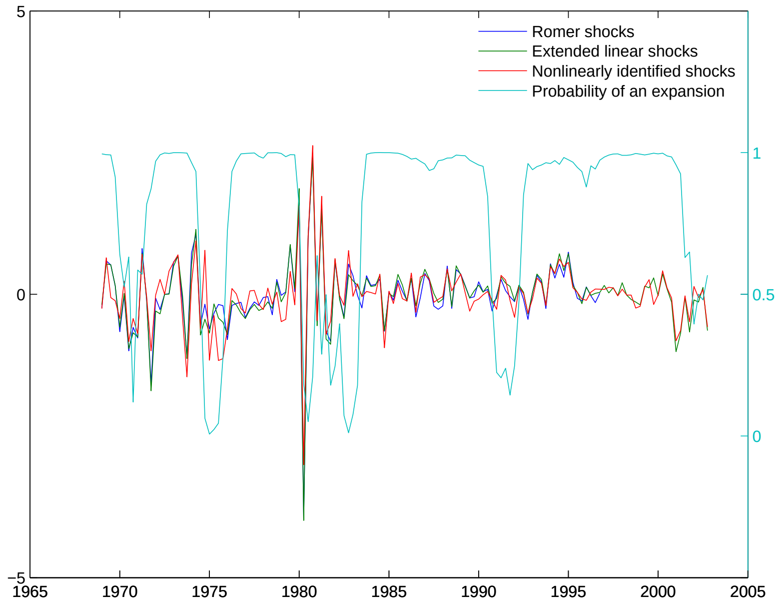

The blue line is the series of monetary policy in Romer and Romer (2004). The green line is constructed in an identical fashion but over a longer sample. The red line is constructed over the same longer sample but with a state-dependent regression model. The turquoise line is the value of the cdf of our state variable $F(z_t)$. See main text for details.

Figure 2.2: Impulse response of headline variables to a monetary policy shock

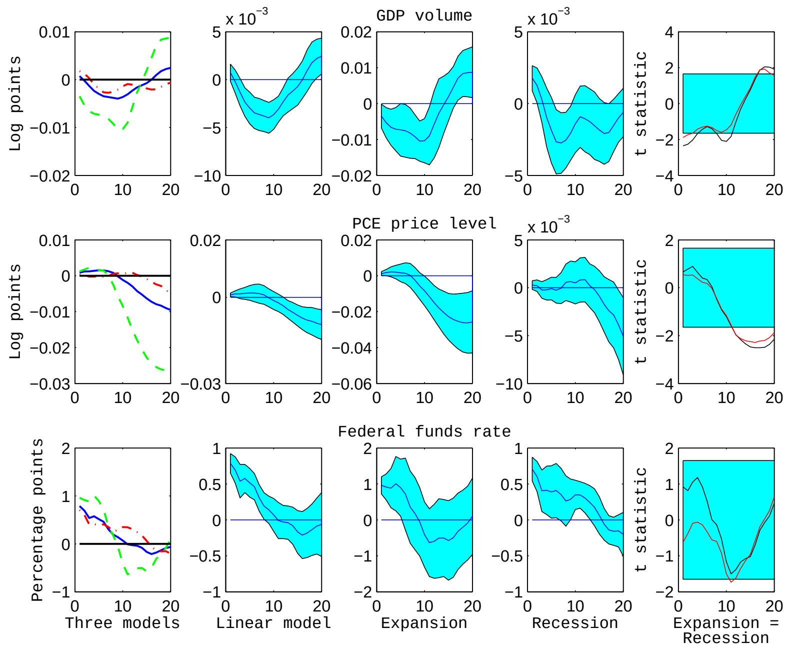

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.3: Impulse response of expenditure volumes

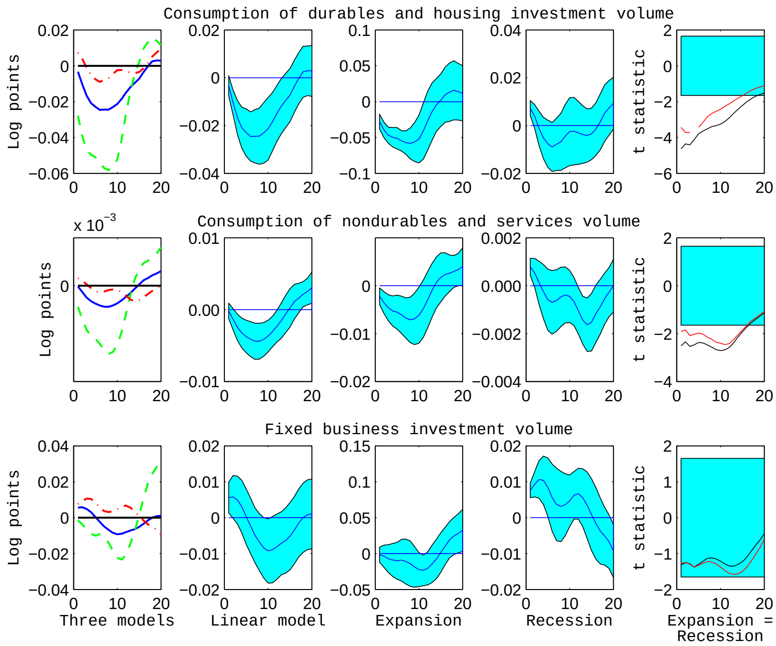

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ±1.65. The first row is the log-level of an index of real durables consumption and housing investment, the second row an index of real consumption of nondurable goods and services, and the third row an index of real fixed business investment.

Figure 2.4: Impulse response of expenditure prices

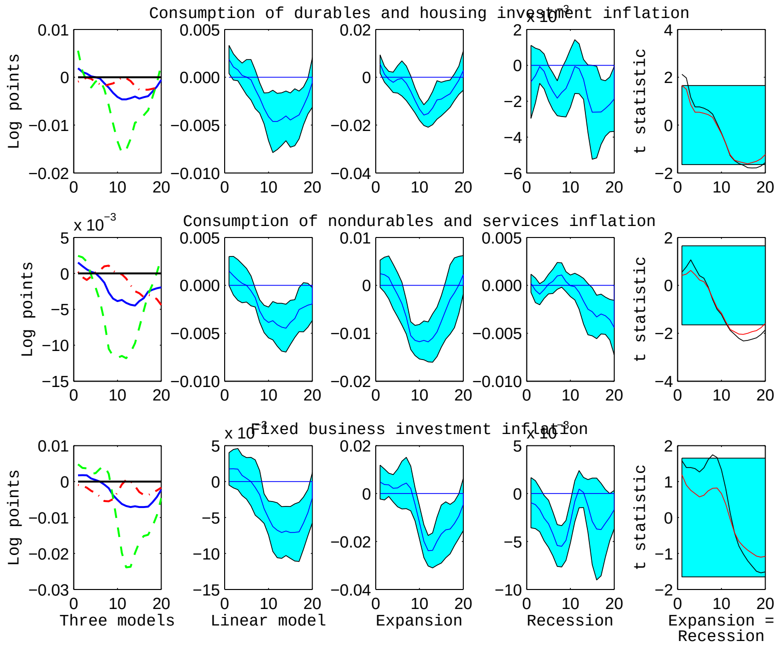

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the quarterly annualised inflation rate of an index of durables consumption and housing investment, the second row the inflation rate of real consumption of nondurable goods and services, and the third row the inflation rate of fixed business investment.

Figure 2.5: Impulse response functions of fiscal and credit variables

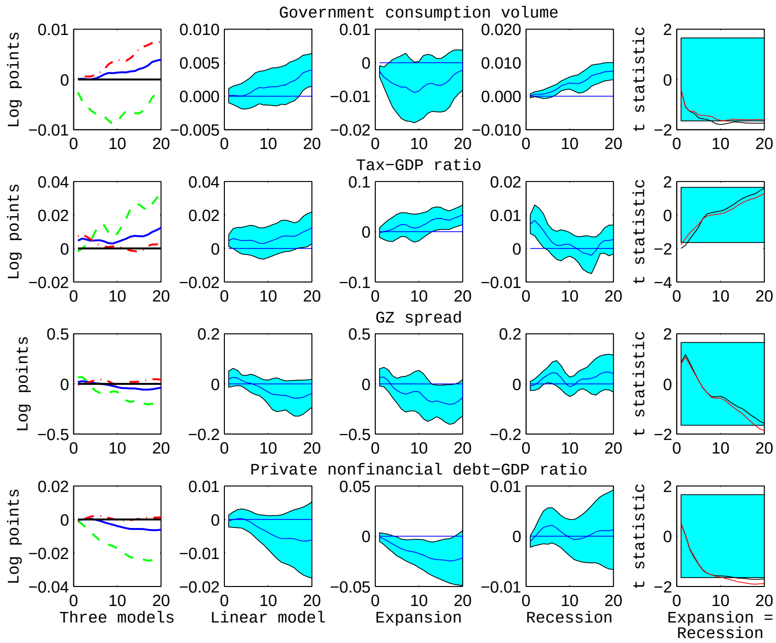

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of an index of real government consumption, the second row the level of the net tax-GDP ratio, the third row the Gilchrist-Zakrajsek bond spread index, and the fourth row the level of private nonfinancial debt to GDP.

Figure 2.6: Impulse response to positive and negative monetary policy shocks

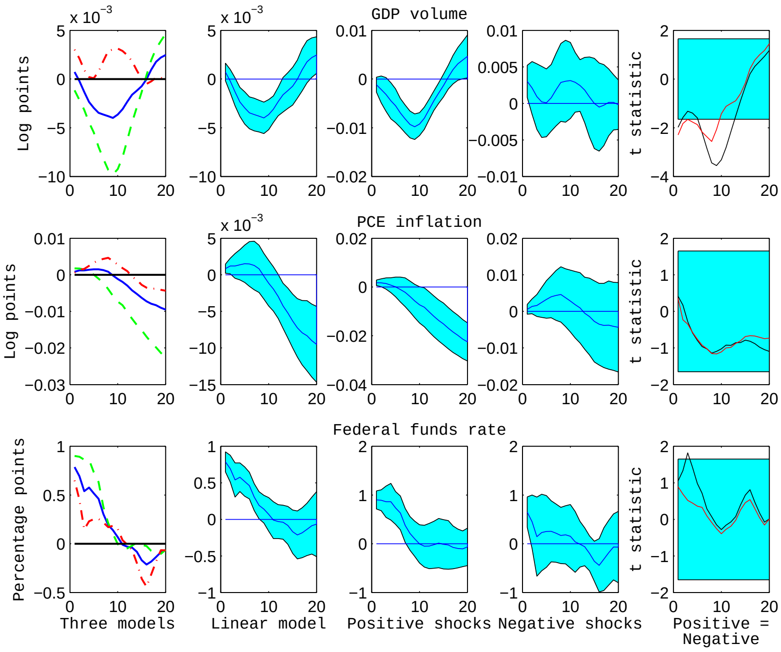

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response to a positive shock, and the red dotted line the response to a negative shock (normalised to have the same sign). The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response to a positive shock, and the fourth column the interval around the response to a negative shock. The fifth column shows t-statistics testing the hypothesis that the difference between the responses to positive and negative shocks is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.7: Pdfs and cdfs of the regime-specific shocks

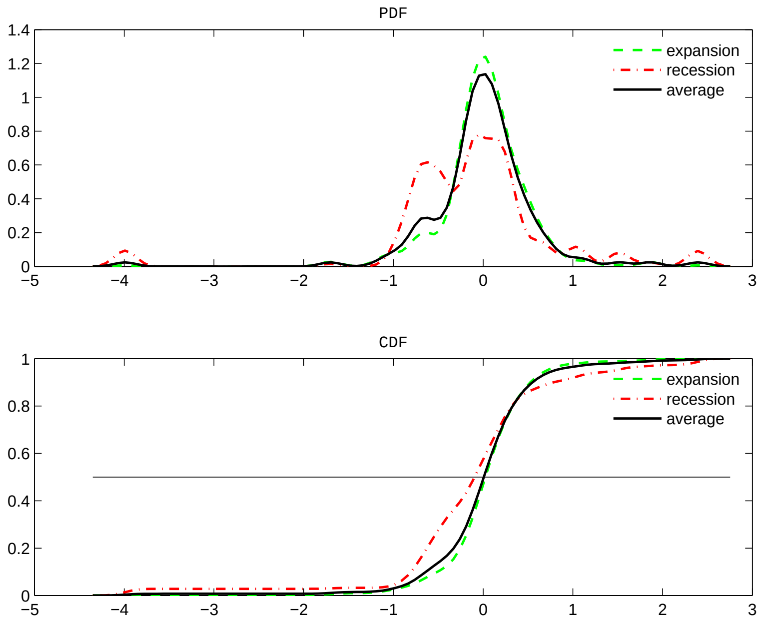

The top pane shows the pdf of the shocks in the different regimes. The bottom pane shows the cdf. The green lines show the distribution during an expansion, the red lines in a recession, and the black line the average of the two regimes.

Figure 2.8: Impulse response to cubed monetary policy shocks

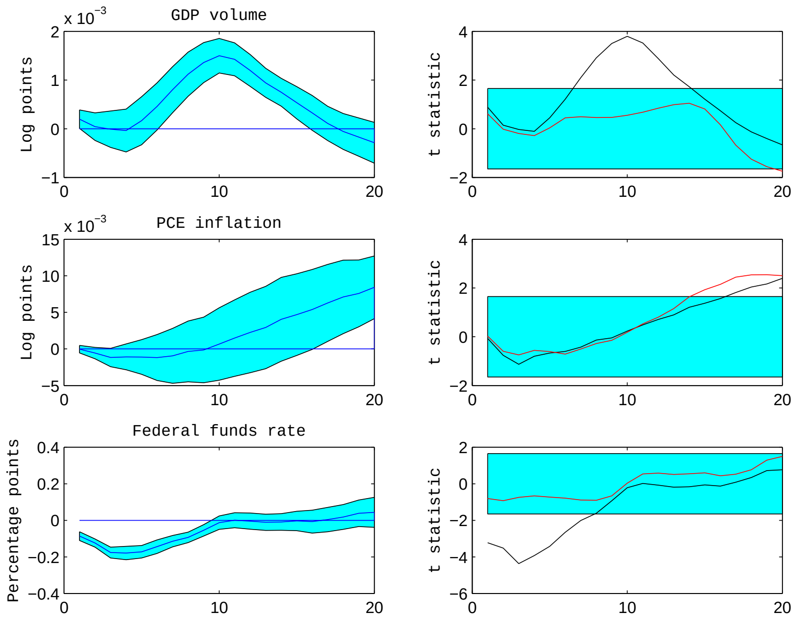

The left hand column shows point estimates and a 90 per cent confidence interval for the impulse response on cubed monetary policy shocks ε_t^3, i.e. β_h^l in the equation $y_{t+h} = \tau t + \alpha_h^b + \beta_h^s \varepsilon_t + \beta_h^l \varepsilon_t^3 + \gamma^{b\prime} x_t + u_t$. The right-hand column shows three estimates of the t-statistic testing the hypothesis that $\beta_h^l = 0$. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65.

Figure 2.9: Impulse response of headline variables to monetary policy shocks identified nonlinearly

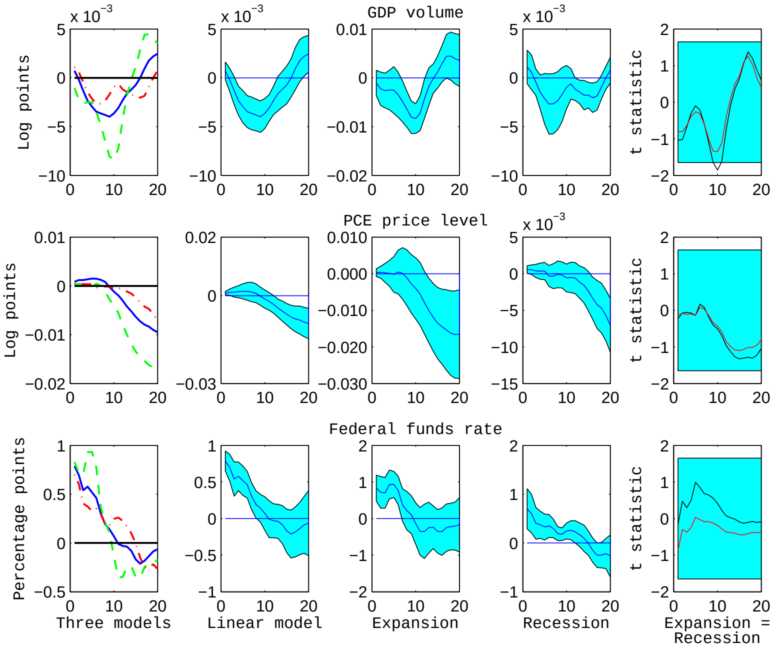

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.10: Impulse response of headline variables to monetary policy shocks identified with a VAR

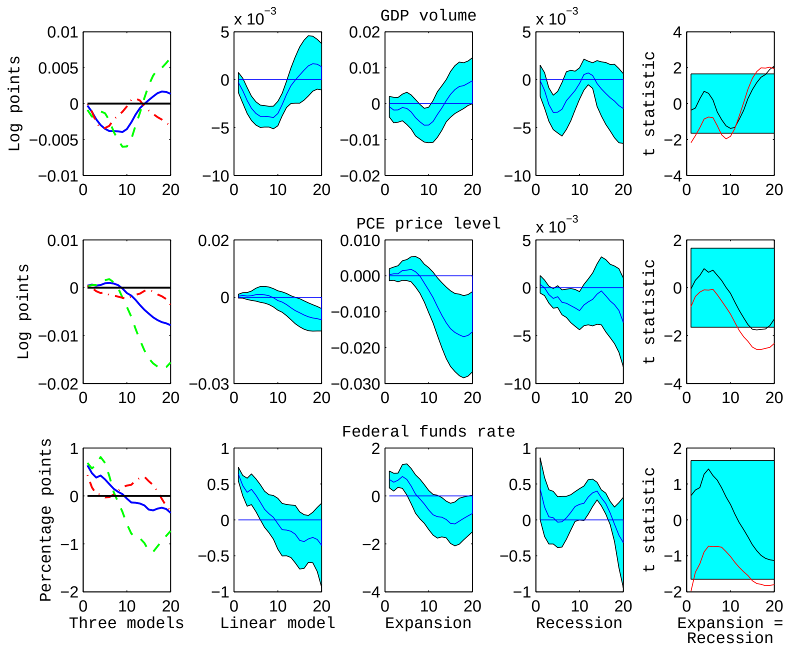

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ±1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.11: IRFs with NBER recession state variable

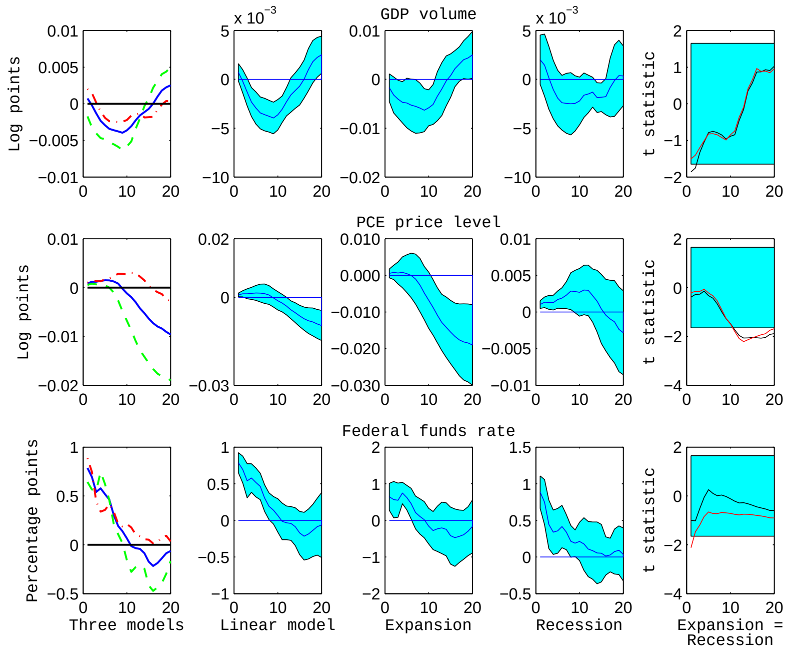

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.12: IRFs with HP-filtered output gap as state variable

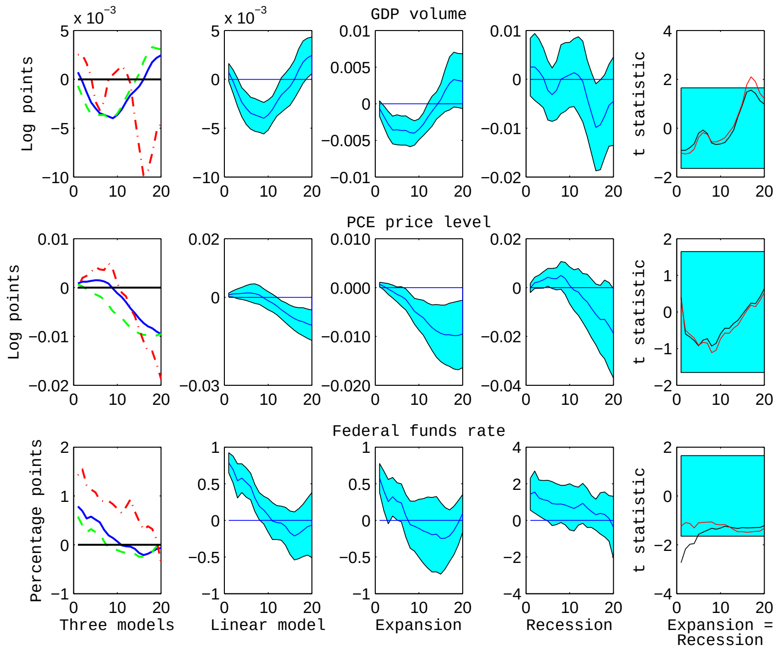

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.13: IRFs with unemployment rate as state variable

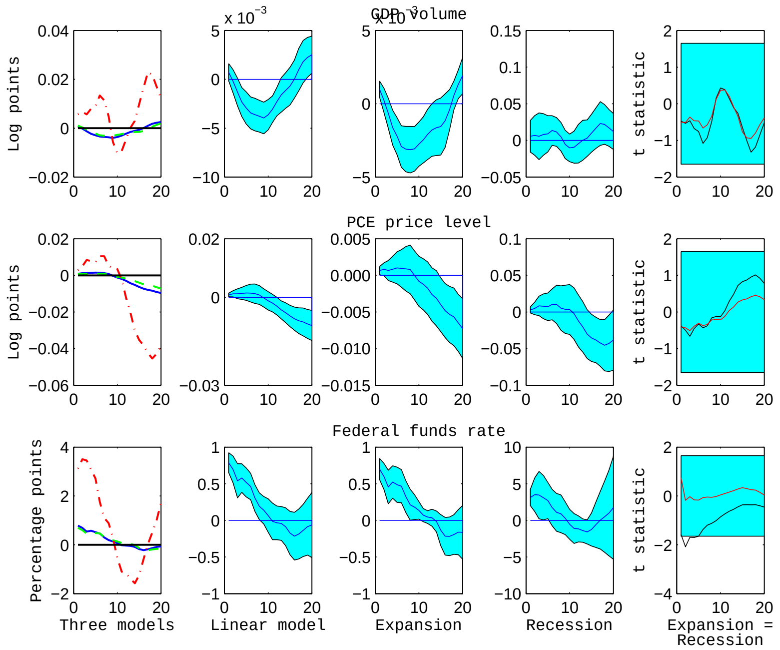

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.14: IRFs with centered state variable

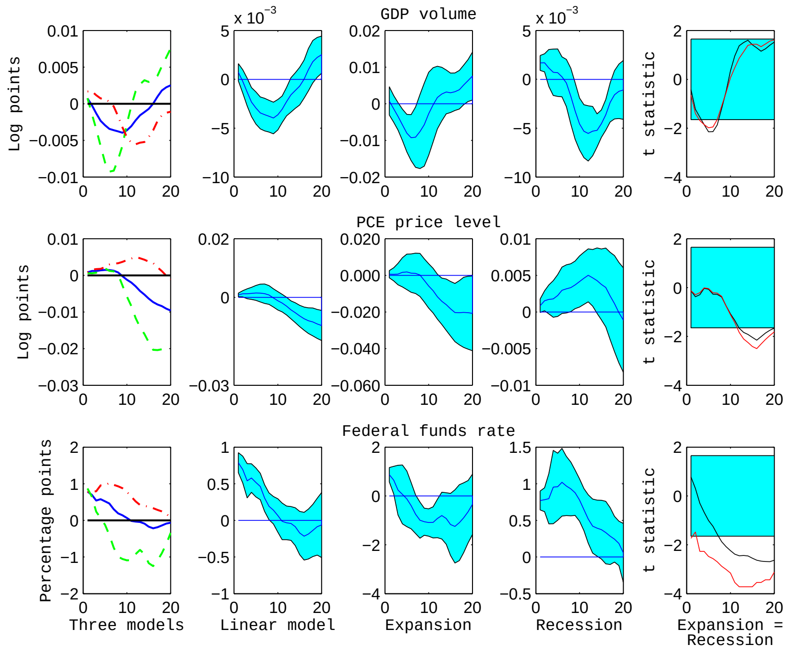

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure 2.15: Impulse response of headline variables to monetary policy shock, $c = 50$

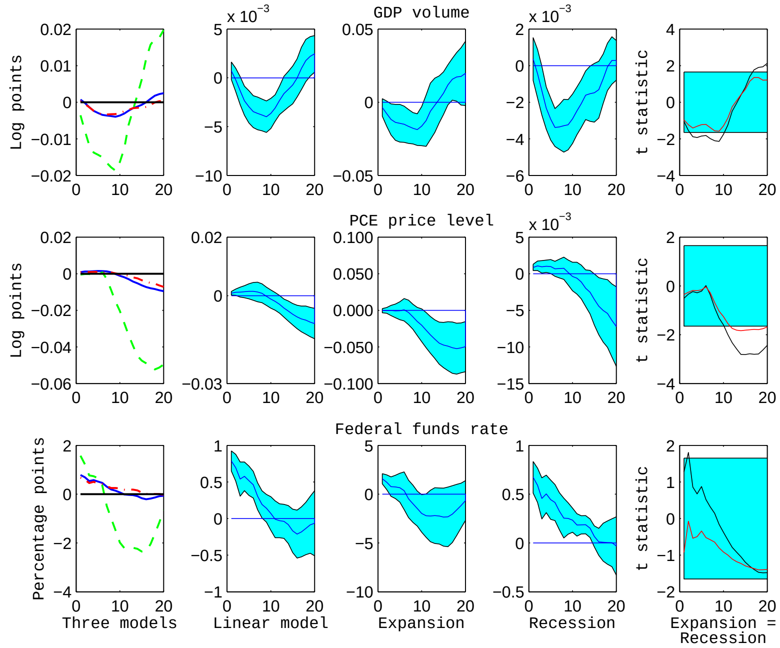

The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 per cent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated assuming spherical disturbances, the red line using a modified Newey-West method, and the blue line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65. The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Chapter 3

Why are real interest rates so low? Secular stagnation and the relative price of investment goods

3.1 Introduction

The financial crisis that began in 2007 pushed central banks in much of the industrialised world to the zero lower bound on nominal policy rates. Much ink has been spilled about how this happened, what central banks should have done when they got there, and how to avoid it happening again. But real interest rates had been trending down across the industrialised world for at least twenty years before this, and had already reached historic lows on the eve of the crisis (Summers (2013), King and Low (2014)). Alongside this fall in interest rates, much of the industrialised world saw house prices and household debt rise to historic highs before the crisis. While these series have subsequently fallen back somewhat, they appear at the time of writing to have stabilised at elevated levels in relation to GDP and real incomes in many countries.

There have been many explanations for this fall in industrialised-world interest rates, among which are three leading candidates. The first is demographics - in particular a rise in the weight of high-saving age-groups as baby-boomers enter late middle age. The second is inequality, whereby a rise in the share of income or wealth accruing to the high-saving rich has raised aggregate saving. And the third is emerging markets, whereby an excess of saving in the developing world has pushed down on rich-world interest rates.

Each of these explanations has merit. But what they all have in common is a rise in domestic or foreign saving as a cause of the fall in interest rates. They all predict, therefore, a rise in investment in the industrialised world.[1] But in contrast, nominal investment rates have fallen sharply across the industrialised world over the past thirty years, a fall which again long predates the recent financial crisis.

This paper fleshes out a new explanation for the falls in real interest rates and rises in household debt across the industrialised world, complementary to those which rely on higher saving, but which also explains the fall in investment rates. The story is based on the widespread fall in the price of investment goods relative to consumption over the past thirty or so years documented in Karabarbounis and Neiman (2014). I extend their data back in time for some countries, to show that this fall has not been a feature of the very long run, but rather began a few years either side of 1980.

In the model, households need to save to provide for retirement. The corporate sector invests the savings of the household sector in capital goods. If the price of capital goods falls, a given quantity of savings can buy more capital goods, raising the return on investment for a given marginal physical product of capital. But the increase in the volume of capital goods lowers the marginal product, thereby lowering the return on investment. The net impact of these two effects depends on the curvature of the production function.

I parameterise the model with a less-than-unit elasticity of substitution between labour and capital, in line with most estimates in the literature (see e.g. Chirinko (2008)). Consistent with the predictions of the model at these parameter values, I present cross-country evidence showing that nominal investment rates have fallen further in countries where the relative price of investment has fallen further.

Depending on parameterisation and the timing convention in the model,

[1] With the caveat that some demographic models that featuring slowing population growth may predict falling investment rates

the dynamics of the transition to this new steady state can involve a temporary rise in interest rates, as households attempt to bring forward the extra consumption afforded by the fall in the relative price of capital. This provides a new interpretation of the period of historically high world real interest rates experienced in the 1980s. More generally, the transitional dynamics operate for some decades both before and after the change in relative prices. For example, I find that the fall in the relative price of capital has been particularly good for the baby boomer generation whose housing wealth has been revalued by the shock.

But the new steady state is one of lower interest and investment rates and higher household debt ratios, even after investment goods prices have stopped falling. Lower interest rates reduce the user cost of housing, boosting housing demand. Housing supply is fixed, so house prices (or at least land prices) must rise. Houses are bought early in life and largely on credit, so household debt also increases. Acquiring these debt claims is an alternative form of retirement saving, so the capital investment rate falls in the steady state, as we see in the data. The model's implications for household debt and house prices receive qualified support in cross-country econometric analysis. I extend the model to allow for bequests, and for heterogeneity in the bequest motive. My core findings are robust to this modification. Furthermore I find that the real interest rate moves in the *opposite* direction to wealth inequality, in contrast to Piketty, Postel-Vinay, and Rosenthal (2014), but moves in the same direction as consumption inequality.

These findings cast recent debates on macroeconomic imbalances and household and government indebtedness in a new light, and have important policy implications. Some prominent policymakers (see, for example, Ingves (2014)) are seeking to prevent what they see as 'excessive' levels of household debt. But if low rates of interest and investment, accompanied by pressure for governments and households to become indebted, represent the transition to a new steady state in which the corporate sector's demand for household savings is weak, then attempts by macroprudential or monetary authorities to prevent this may be futile or counterproductive. I show that preventing the rise in household debt in response to a fall in capital goods prices makes interest rates fall further in response to the initial shock.

The mechanism in this paper builds on a long history of related ideas in the literature. Summers (2013) recently raised the issue of the pre-crisis falls in real interest rates and the possibility that they would stay low for an extended period in the future. But the idea that capitalist economies could be plagued by chronically low returns on capital, and that this could result from an overaccumulation, in some sense, of physical capital goes back at least to Marx (1867) and Hansen (1938). The fall in capital goods prices in the face of a need for retirement savings creates a form of asset shortage reminiscent of Caballero, Farhi, and Gourinchas (2008), which is satisfied by the endogenous creation of debt claims on the young. The focus on the fall in the relative price of investment goods builds on the important contribution of Karabarbounis and Neiman (2014), whose data I draw on for this study. Rey (2014) and Summers (2014) have, among others, linked secular stagnation to falls in the relative price of capital goods.

Two papers are particularly close, methodologically speaking, to the present study. Giglio and Severo (2012) examine the effect of a change in production technology in an OLG model. Like the present study, Eggertsson and Mehrotra (2014) address the issue of secular stagnation in an OLG model. They show that a tightening of the debt limits facing young households, reduced population growth and increased income inequality can reduce the equilibrium real interest rate in such a model, and explore the consequences for resource utilisation in a sticky price model. They also show that falling relative capital goods prices can lower the real interest rate. Relative to that study, this paper gives conditions under which interest rates can remain low even after capital goods prices have stopped falling; IMF (2014) finds that relative prices have been stable since 2002. Furthermore, this paper derives the implications of lower capital goods prices for house prices, household debt and wealth inequality. But unlike Eggertsson and Mehrotra (2014), this paper says nothing about resource utilisation or nominal variables.

The remainder of this paper is structured as follows. Section 3.2 sets out the key facts the model aims to explain. Section 3.3 describes the core economics of the paper in the simplest possible model. Section 3.4 describes the baseline model. Section 4.6 shows the results of model simulations in which I vary the relative price of investment and generate movements in interest rates, investment rates and household debt which are qualitatively similar to those presented in section 3.2. Section 3.6 examines the sensitivity of these findings to parameter values, and extends the model to allow for bequests. Section 3.7 extends the model to allow for bequests, heterogeneous agents, intangible capital and open-economy considerations. Section 3.8 adduces some econometric evidence in support of the model. Section 3.9 concludes.

3.2 Motivating facts

This section sets out the key stylised facts that the model aims to connect: falling real interest rates (subsection 3.2.1); rising household debt (subsection 3.2.2); and falling capital goods prices and nominal investment rates (subsection 3.2.3). This section also addresses cross-country movements in factor shares (subsection 3.2.4) I focus on the widest possible set of industrialised countries for each data series, but also, where possible, show data for a subset consisting of the 11 advanced countries[2] for which the EU-KLEMS database has sufficient data to calculate long time-series of nominal and real capital-GDP ratios.

3.2.1 Falling real interest rates

Ex-ante real interest rates can now readily be measured in many industrialised countries with reference to the yields on index-linked government liabilities. However, these securities were not issued before the 1980s, complicating the measurement of ex-ante real interest rates before then. IMF (2014) presents an attempt to solve this problem by constructing a parametric model of inflation expectations and subtracting the result from the yields on nominal government liabilities.

Figure 3.1 shows the result for the UK and the US. The figure shows that interest rates have been trending generally downwards for the 30 years since their recent peak in the early 1980s. The model-based series in IMF (2014) suggest that US ex ante real rates were close to current levels in the early 1970s, fell below zero in the middle part of that decade, before rising sharply in the late 1970s-early 1980s. King and Low (2014) and Laubach

[2]Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US

and Williams (2003) (updated to 2014) both also show declining real interest rates from the 1980s to the 2008 crisis.

3.2.2 Rising household debt ratios

Figure 3.2 shows an index of the ratio of household debt to GDP since 1970 for a broad sample of industrialised countries and our restricted sample of 11 countries. The figure shows a rise in the average ratio of around 50pp since 1970.

3.2.3 Price and quantity of capital investment and stock

Figure 3.3 shows the simple average across OECD countries and across our restricted sample of the ratio of nominal investment to nominal GDP. The nominal investment rate has been trending downwards since at least the mid-1970s. Figure 3.4 shows that the corresponding stock ratio (the current replacement cost of the capital stock as a proportion of GDP) had also fallen from nearly 4 times annual GDP around 1980 to nearly 3 times by 2007 for the 11 countries in the EU-KLEMS database for which data are available.

Figure 3.5 shows the real investment - GDP ratio across the same two sets of countries since 1970. The series show no strong trend over the whole sample, although there is weak evidence of an upward trend since the early 1980s. Figure 3.4 shows that the ratio of the real capital stock to real GDP (both at 1995 prices) has been trending upwards since the 1970s.

These divergent patterns in the nominal and real ratios are of course a manifestation of a trend fall in the price of investment goods relative to consumption or GDP, documented in IMF (2014) and Karabarbounis and Neiman (2014). Figure 3.6 shows four series of the ratio of the investment deflator to the consumption deflator. The red and blue lines are taken from the respective countries' national accounts data. The green line is the average change across all the countries in the dataset, and the purple line is the average among our restricted sample in this dataset. All three lines show that the relative price of investment goods has been falling in recent decades, with a fall of perhaps 30% since the mid-1970s. The longer series show that, prior to this fall, there has not been a secular trend in this relative price.

3.2.4 Factor shares

Karabarbounis and Neiman (2014) document a fall in the corporate and whole-economy labour share since 1975, within a large number of industries and countries. In the baseline model presented below, there are no pure profits in the economy and only two factors of production - capital and labour. In such a world, a falling labour share must imply a rising profit share, which in turn is equal to the product of the average return on capital and the capital-output ratio

$$\frac{\Pi}{Y} = \frac{\Pi}{Kp_K}\frac{Kp_K}{Y} \tag{3.1}$$

Figures 3.1 and 3.4 show that the real interest rate and the nominal capital-output ratio have typically fallen in industrialised countries over past thirty years which, according to the simple equation above, would generate a fall in the labour share. However, there are a number of explanations for the apparent discrepancy between the trends in the labour share, the real interest rate and the capital output ratio:

- The capital-output ratio could be somehow mismeasured, perhaps because of the omission of intangible or nonreproduced factors of production, and has in fact not fallen over time.

- The marginal and average returns on capital correspond to the real interest rate and the average profit rate respectively. There could be an increasing wedge between them, or an increasing wedge between the marginal product of capital and the real interest rate in financial markets, perhaps because of corporate taxes or physical depreciation.

- Corporate profits include a component of 'pure profit' as well as remuneration for capital investment, corresponding, for example, to producer markups over marginal cost. These markups could have risen over time.

- Relatedly, profits could be remunerating highly-skilled or managerial labour, e.g. through the granting of share options, such that the labour share, broadly conceived, has not fall as much as the wage share would suggest.

The baseline version of the model incorporates none of these features. Equation 3.1 holds in the model and therefore, in generating a falling real interest rate and a falling capital-output ratio, it also produces a rise in the labour share. Appendix 3.C outlines a variant of the baseline model with three factors of production which can generate both a rising profit share and a falling investment rate.

3.3 Real interest rates, capital goods prices and the curvature of the production function

Other things equal, lower capital goods prices p raise the return on capital when denominated in consumption goods: a foregone consumption good buys more capital goods, so for a given marginal product of capital, the return on investment

$$r = \frac{1}{p}\frac{\partial Y}{\partial K} - \delta \tag{3.2}$$

is higher.

But other things will not be equal - the fall in capital goods prices will mean that a given volume of savings will finance more of them, pushing down on the marginal product of capital to an extent that depends on the curvature of the production function. Whether the volume effect outweighs the price effect depends on the curvature of the production function. And savings may respond to the resulting change in interest rates in either direction, depending on the properties of the utility function. To crystallise these issues before I present the baseline model, this section of the paper analyses the role of the curvature of the production and utility functions in the simplest possible model with variable capital goods prices.

3.3.1 Simplest possible model

Consider a world populated by an identical series of overlapping generations, each of which lives for two periods.[3] Each generation of households has a standard isoelastic utility function defined over consumption in each

[3]Overlapping generations are necessary because the interest rate in an infinite horizon model would be pinned down by the household discount factor

generation of life

$$U\left(c_1, c_2\right) = \frac{c_1^{1-\theta}}{1-\theta} + \beta \frac{c_2^{1-\theta}}{1-\theta}$$

Households supply one unit of labour at wage rate W in the first period of life, and can lend money to firms at net interest rate r to provide for their retirement. So their intertemporal budget constraints are as follows

$$c_2 \leq \left(W - c_1\right)\left(1 + r\right)$$

Young households' saving in the first period of life as a fraction of their wage income can be shown to be given by [4]

$$s = \frac{W - c_1}{W} = \frac{\beta^{\frac{1}{\theta}}\left(1+r\right)^{\frac{1}{\theta}-1}}{1 + \beta^{\frac{1}{\theta}}\left(1+r\right)^{\frac{1}{\theta}-1}} \tag{3.3}$$

This familiar expression shows that the sign of the slope of the savings schedule in $\{s, r\}$ space depends on the intertemporal elasticity of substitution $\frac{1}{\theta}$. When this substitution elasticity is high (i.e. above unity), a fall in interest rates causes a fall in savings, as agents substitute away from relatively expensive retirement consumption. Infinite-horizon households pin the interest rate down at $r = \frac{1}{\beta} - 1$, and are thus equivalent to OLG households with linear period utility functions. When the elasticity is below unity, retirement saving is akin to a Giffen good: lower interest rates *raise* the savings rate out of wages, as the desire to offset the negative effect of lower interest rates on retirement consumption outweighs the higher price of it. When the elasticity is exactly one, these two effects cancel and the savings schedule is vertical.

Turning to the determination of factor prices, firms hire labour and borrow funds from young households, buy capital goods (which depreciatate at rate δ) at relative price p and maximise profits with them.[5] Factor prices $\{W, r\}$ will therefore be set equal to marginal product in the standard fash-

[4] See Appendix for derivation.

[5] We can for now think of a class of final goods firms turning intermediate goods into consumption goods one-for-one or into capital goods at rate p^{-1}. This will be made more explicit when describing the full model in section 3.4

ion

$$W = \frac{\partial Y}{\partial L}$$
$$r = \frac{1}{p}\frac{\partial Y}{\partial K} - \delta$$

In aggregate, the gross savings of the young will equal the replacement cost of the capital stock

$$pK = sW$$

If we assume a CES production function with elasticity of substitution σ and capital share parameter α we can derive an 'investment schedule' that implicitly maps s into $\{r, p\}$[6]

$$s = \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{1 - \alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}} \tag{3.4}$$

Which way does the interest rate schedule slope in $\{s, r\}$ space? There are two effects. The effect in the numerator is negative for the standard reasons: for given capital goods prices, more savings reduces the marginal product of capital and hence the interest rate. The effect in the denomnator is of ambigious sign, and comes through the labour share (for a Cobb-Douglas function $\sigma - 1 = 0$ it is absent). For low σ, an increase in r reduces the denominator, raising the quotient. This is because the saving rate is expressed here as a fraction of wages and when $\sigma < 1$, higher interest rates are associated with a lower labour share. To save enough for a given volume of capital goods, a lower labour share must mean a higher saving rate. For reasonable parameter values, the effect on the numerator will dominate, such that the investment schedule slopes down in $\{s, r\}$ space.

The derivative of the saving rate with respect to the price of capital goods p is the same sign as $1 - \sigma$. Consider a fall in the relative price of capital goods of x percent. Holding the marginal product of capital constant, the return on investment increases by x percent as each consumption unit of investment buys x percent more capital goods. But because of the price fall a given volume of savings can finance x per cent more capital goods, and

[6]Derived in Appendix 3.A.2

the marginal product of each will fall by $\frac{x}{\sigma}$ per cent, such that the sign of the effect is equal to the sign of $1-\sigma$.

Figure 3.7 depicts graphically how the effect of a fall in p is governed by the effects of the curvature parameters $\{\sigma, \theta\}$. The top left panel shows the effect of a fall in the relative price of investment goods on the investment schedule for values of σ either side of unity. The top right panel adds an upward-sloping savings schedule, corresponding to a relatively elastic utility function. In this case, the rates of interest and of investment/saving will covary positively, with the sign of the change once again depending on the sign of $dp\,(1-\sigma)$. The cases of a small open economy or of infinite-horizon households correspond to a horizontal saving schedule - no change in interest rates and a change in investment rates of the same sign as $dp\,(1-\sigma)$. The bottom-left panel depicts a highly inelastic utility function. In this case, the changes in the rates of interest and saving are of opposite sign, but the former is still the same sign as $dp\,(1-\sigma)$. Finally, the bottom-right panel shows an extreme case in which the savings schedule slopes downward but is *shallower* than the investment schedule. In this case a fall in the relative price of capital would lead to a fall in the investment rate and a rise in interest rates if $dp\,(1-\sigma)<1$.

3.4 The baseline model

In this section we augment the heuristic model above with an intermediate period of working life, and with the requirement for households to buy a house when young. This enables us to analyse the effect of capital goods prices on house prices and household debt, and how the existence of both alters the determination of interest rates.

3.4.1 Households

The economy is closed and comprises three overlapping generations of constant and equal size. Each generation has a standard separable CES utility function over consumption and 'housing'[7]

$$U(c_1, c_2^{'}, c_3^{''}, h) = \frac{1}{1-\theta}\left(c_1^{1-\theta} + \beta_2 c_2^{'1-\theta} + \beta_3 c_3^{''1-\theta}\right) + \phi\frac{h^{1-\gamma}}{1-\gamma} \tag{3.5}$$

where I denote leads one and two periods hence with primes and double-primes respectively. In period 1 (young adulthood), the household supplies η units of labour inelastically (remunerated at wage W), consumes goods and buys a house. She can borrow or save a net amount S_1 at rate r. In period 2 (middle age), the household remains in said house, supplies $(1-\eta)$ units of labour, and can again borrow or save $S_2^{'}$. In period 3 (retirement), she sells her house and consumes the proceeds plus her accumulated savings. So each of the three periods is associated with a budget constraint as follows

$$c_1 + hp_h + S_1 = \eta W \tag{3.6}$$

$$c_2^{'} + S_2^{'} = (1-\eta)W + (1+r)S_1 \tag{3.7}$$

$$c_3^{''} = (1+r^{''})S_2^{'} + hp_h \tag{3.8}$$

Forming and solving the Lagrangean yields standard consumption Euler equations thus

$$\frac{c_1^{-\theta}}{(1+r^{'})(1+r^{''})} = \frac{\beta_2 c_2^{'-\theta}}{(1+r^{''})} = \beta_3 c_3^{''-\theta} \tag{3.9}$$

We also get a housing demand equation that depends on future house prices and consumption as you would expect

$$\phi h^{-\gamma} + \beta_3 c_3^{''-\theta} p_h^{''} = c_1^{-\theta} p_h \tag{3.10}$$

This is intuitive. The LHS is the marginal utility of housing plus the discounted marginal utility of the retirement consumption paid for by the sale of the house. The RHS is the consumption utility cost of buying a unit of housing.

[7] Housing is in fixed supply so might be more usefully thought of as land, or more generally any non-produced asset that yields utility

3.4.2 Firms

A measure of perfectly competitive firms produce intermediate goods, combining capital and labour with a CES production technology

$$Y = A[(1-\alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \tag{3.11}$$

These intermediates can then either be consumed directly, or transformed into capital goods at rate p units of intermediate for every one unit of capital. The relative price of investment goods - the key exogenous parameter in our model - is therefore p. This means of introducing investment-specific technological change is isomorphic to that in Greenwood, Hercowitz, and Krusell (1997).

Wages are set equal to the marginal product of labour

$$W = \frac{\partial Y}{\partial L} \tag{3.12}$$

Firms equate the user cost of capital to its marginal product, both denominated in consumption goods

$$1 + r^{'} = \frac{1}{p_K}\frac{\partial Y'}{\partial K'} + \frac{p_K^{'}}{p_K}(1-\delta) \tag{3.13}$$

3.4.3 Market clearing

At the end of each period, the net savings of households of young and middle age are transformed into next period's capital stock (at this period's relative prices), such that the following capital-market clearing condition holds in stock terms

$$S_1 + S_2 = K^{'} p_K \tag{3.14}$$

There is a fixed measure $\bar{H}$ of housing or land for each of the first two generations to live in, so that in equilibrium

$$h = \bar{H} \tag{3.15}$$

3.5 Results

In this section of the paper we set out the results of the baseline model. Subsection 3.5.1 explains how it is parameterised, subsection 3.5.2 how it is solved, and subsections 3.5.3 and 3.5.4 discuss comparative static and dynamic results respectively.

3.5.1 Parameterisation

Each of the three periods of adult life lasts twenty years. The discount factors $\{\beta_1, \beta_2\}$ and capital share parameter α are set to hit an annualised steady-state interest rate of 3% and a capital share of one-third respectively. The depreciation rate δ is set at set at the standard value of .05 in annualised terms. I set ϕ to hit a sensible value for the ratios of housing wealth to GDP. The elasticity in the utility function $\{\theta\}$ is set to unity (log utility).

A fall in p_K amounts to an improvement in the overall level of technology, in the sense that the lower is p_K, the larger is the total volume of consumption and investment goods a given factor endowment can produce. However, the overall growth rate of TFP has not notably accelerated over the past several decades. So when considering changes in p, I change A so as to keep potential GDP unchanged given existing factor endowments. The fall in capital goods prices in the simulations prompts an accumulation of capital goods, so potential GDP rises nonetheless.

The production elasticity σ is set to 0.7. As demonstrated below in the sensitivity analysis of the full model and above in the exposition of its toy analogue, this parameter is crucial for the behaviour of the model. Chirinko (2008) discusses a number of approaches for estimating this parameter and the resulting range of estimates. These are typically based either explicitly on a firm's optimisation problem, choosing capital subject to adjustment costs, or on estimating a relation between investment or capital intensity and the level or change in the user cost of capital. Chirinko (2008) cites over thirty estimates, typically based on firm-level panel or aggregate time-series data. The median of these estimates is 0.6 and the mean 0.5, with about 15% of the estimates above the critical value of unity. The author quotes a preferred range of $\sigma = 0.4 - 0.6$, i.e. much further below the critical value of unity than the baseline assumption used in this paper. Section 3.8

adds to this evidence of a below-unit elasticity using cross-country panel data on nominal investment rates and prices.

A noteworthy recent addition to this set of estimates is provided by Karabarbounis and Neiman (2014), who compile a large cross-country panel dataset on the relative price of capital goods and the labour share in the corporate sector. They write down a standard model in which the production side of the economy is very similar to that in the present study, but in which consumers have infinite horizons and therefore the interest rate is pinned down by the discount factor. They identify country-specific time trends in the data with transitions between steady states in their model, and find using a 'robust regression' algorithm that the labour share has fallen faster on average in countries in which the relative price of capital has fallen faster.[8] This points to a value of σ exceeding unity; the authors central estimate is around 1.25. The contrasting conclusions in Section 3.8, based on regressions of the nominal investment rate on the relative price of capital, employ the same dataset.

3.5.2 Solution method

We first solve for the steady state of the model for a given value of capital goods prices p. An initial assumption is made about house prices and the savings of each generation $X^0 = \left\{S_1^0, S_2^0, p_h^0\right\}$, which implies a certain constellation of factor prices $\left\{W^0, r^0\right\}$. Household behaviour is then optimised taking these prices as given, the resulting optimal values of $X^* = \{S_1^*, S_2^*, p_h^*\}$ are computed, and the initial guess is updated toward them - i.e. $X^1 = \lambda X^* + (1-\lambda) X^0$, where $\lambda \in (0,1)$ is a gain parameter. This process is repeated until the solution converges to a fixed point, i.e. until $X^n \approx X^{n-1}$.

To assess the dynamic effects of a change in p, I consider a simulation path of sufficient length T that the economy will be at the steady state at the beginning and end of the simulation, with the exogenous changes to p occurring in the middle. I first calculate the steady state in each period t of the simulation $\{X_t^{ss}\}_{t=1}^{T}$, given the extant values of the exogenous parameters. I then optimise the behaviour of each generation t, taking the behaviour

[8]The authors use the *rreg* command in STATA as a means of downweighting outliers

of the other generations as given, obtaining $X_t^*\left(\{X_s^{ss}\}_{s\neq t}\right)\forall t$. As above, the initial guess is updated towards this solution until it converges. I verify that the model converges to the steady state well inside the endpoints of the simulation.

3.5.3 Comparative statics

The blue lines in figure 3.8 show the the effect of varying the relative price of capital goods p on the steady state of the model at the baseline parameter values. In the baseline model, the annualised real interest rate falls by 20 basis points. The nominal investment rate falls about 1 percentage point in response to the lower relative price of capital (bottom left panel). This implies a somewhat upward-sloping savings schedule in the model, notwithstanding the assumption of log utility which, in the simple two-period model of section 3.3. This is because the fall in interest rates lowers the user cost of housing for a given house price. House prices must rise to choke off the resulting increase in demand - by about 10 per cent in the baseline, relative to GDP. To fund the purchase of more expensive houses, the ratio of the net debt of young households to GDP increases by about 20 percentage points (top right panel). Housing is a store of value as well as a consumption good, so the purchase of a house is an alternative to the purchase of capital goods as a means to fund retirement consumption; in general equilibrium, part of the money that would have gone to fund the purchase of capital goods is instead lent to the young to fund their house purchase, who live off the sale proceeds in old age.

3.5.4 Dynamic results

What are the dynamic consequences of the experiment considered above? We analyse the dynamic impact of a 30% fall in the relative price of capital goods over one model period (20 years). The exercise of mapping into the data 20-year model time periods, each of which contains a series of supposedly discrete events, is somewhat nuanced. According to the model's timing conventions, savings accumulated at the end of period $t-1$ become productive in period t. An important question is which period's capital goods prices are used to convert savings into capital goods, and back into consumption goods. In the baseline simulation shown here, consumption foregone in period $t-1$ becomes productive capital goods in period t, with the conversion happening at period t prices. For this reason, the first period of low capital goods prices - period 10 in the charts - corresponds roughly to the 1980s and 1990s in the data. Interest rates are measured ex ante - so the period 10 interest rate is the return on savings made in the 1990s, paying a return in the 2010s. The interest rates observed at the time of writing (the mid-2010s) correspond to period 11 of the model. Alternative timing conventions are possible - for example turning period $t-1$ savings into period t productive capital goods at period $t-1$ prices - and are explored in the sensitivity analysis below. Timing conventions would matter less in a model with shorter time periods or in periods with more stable capital goods prices, and of course do not matter at all when analysing the steady state.

Figure 3.9 shows the results of the baseline dynamic simulation. In each panel, the blue line shows the relative price of capital goods produced in the period in question. The top left panel shows the path of the ex-ante real interest rate. The ex-ante interest rate earned on savings made at the end of period 9 (before the fall in capital goods prices) rises. This corresponds to the late 1970s, a period of rising world interest rates. The middle left panel shows that the saving (or investment) rate falls before the shock hits, recovers partially, and then resumes its fall. A fall in the saving rate combined with a rise in the interest rate is indicative of a shift inwards in the saving schedule. Consistent with this, the top right panel shows the path of household debt, which begins to rise in advance of the fall in capital goods prices. Younger generations can look forward to funding their retirement in part by selling more expensive houses, and thus begin consuming and dissaving more. The middle-right panel shows that the rise in housing wealth in relation to GDP takes several generations to be completed.

The bottom right panel shows that the profit share initially rises and then falls when the shock hits, as the fall in the interest rate outweighs the rise in the capital-output ratio at the assumed parameter values. How this feature of the model relates to the evidence is discussed in section 3.3.

Finally, the bottom-right panel shows the response of output and the consumption of each age group. Output initially falls very slightly as capital is decumulated, but eventually rises as the fall in capital goods prices affords a larger real capital stock. GDP rises despite the assumed fall in Hicks-neutral productivity in the intermediate goods sector, which is calibrated to be sufficient to offset the improvement in technology that the fall in p represents without any increase in factor endowments. The consumption of the young generations rises sooner, and by more, than that of older generations, such that in the steady state the age-consumption profile is flatter. It rises more because the steady-state interest rate is lower, encouraging households to consume earlier, and sooner because households who are young on the eve of the shock anticipate capital gains on their house purchases. This pattern of capital gains can also be observed in the consumption of the old - the generation that is old in period 11 (i.e. the baby boomers) consumes more in retirement than any other retired generation, because it enjoyed the biggest capital gains on housing, buying them relatively cheaply in the 1960-1970s and then trading down in the early 21st century.

Overall, the simulation results generate a qualitatively similar pattern in the real interest rate, housing wealth, the real and nominal capital-output ratios and the household debt-GDP ratio to those which we have observed over the past four decades. The shock is particularly beneficial for the baby-boomer generation. Furthermore, the simulations provide forecasts of what may happen in years to come. In particular, even if the relative price of capital has stopped falling, the interest rate may continue to fall somewhat, as the capital deepening process brought on by the fall in the relative price of capital runs its course. And future generations of retirees will consume less than the current one.

3.6 Sensitivity analysis

In this section paper we conduct sensitivity analysis on the main model. Subsection 3.6.1 varies the availability of household debt. Subsection 3.6.2 varies the key parameters of the model and subsection examines the timing conventions.

3.6.1 Housing and debt

The availability of debt and housing as alternative savings vehicles attenuates the fall in interest rates in the model. This is illustrated in figure 3.10, which considers two alternative regimes for household debt alongside the baseline model. The green line represents a regime in which household debt is forbidden (the net savings of the young must be nonnegative, i.e. $S_1 > 0$). The red represents a simulation in which there is an upper bound on debt that binds at an intermediate level of p. These constraints attenuate the rise in house prices (bottom right panel), as young consumers cannot spread the extra cost of housing over their lives. The fall in the aggregate savings rate (bottom left panel) is also attenuated, as the debt of the young and more expensive houses are less readily available as savings vehicles. Higher savings means more capital and thus lower real interest rates - the top left panel shows that, without household debt, a fall in capital goods prices gives rise to a fall in real interest rates of about 60 basis points, i.e. about three times larger than in the baseline simulation.

Figure 3.11 shows the dynamic effects of a shock to p when household debt is prohibited. The key difference is that the path of interest rates is now monotone. Interest rates do not rise ahead of the shock because young households are not able to borrow to bring forward consumption. The investment rate follows the same falling-rising pattern but now settles at a higher rate than before the shock, as household debt is no longer available as an alternative destination for retirement savings. The bottom-right panel of the figure shows that it is now the middle-aged rather than the young whose consumption rises the most. As before, lower interest rates dissuade retirement saving, but the young cannot respond by dissaving more; only the middle-aged can respond, by reducing retirement consumption at the expense of higher consumption in middle age.

3.6.2 Parameterisation

Curvature of the production function

The key parameter in this model is the elasticity of substitution in the production function σ between capital and labour. Figure 3.12 shows how the impact of p on the steady state of the model depends on the elasticity of substitution between capital and labour σ. When the production function is Cobb-Douglas, the relative price of capital has no effect on the interest rate, house prices, household debt or the investment rate in the steady state. The heuristic model presented in section 3.3 explains why. The volume of capital goods bought with a given quantity of consumption goods is inversely proportional to the relative price. With a Cobb-Douglas production function, the marginal product of capital is inversely proportional to the real capital-output ratio, so these two effects exactly offset. There are nonetheless some dynamic effects during the transition to lower relative capital prices (figure 3.13). In particular, the interest rate rises and then falls, as consumers attempt to bring forward some of the higher consumption afforded by the lower capital goods prices.

Figures 3.12 also the effect of the fall in relative capital goods prices on the steady state of the model when the $\sigma = 1.3$, in line with the estimates in Karabarbounis and Neiman (2014), and symmetric around the Cobb-Douglas case with the baseline value of $\sigma = 0.7$. Everything now goes in the opposite direction to the baseline: interest rates and the investment rate rise, while house prices and the household debt ratio fall. Figure 3.14 shows the dynamics of the transition. The interest rate and investment rate overshoot their long run value during the transition. The profit share rises and the household debt ratio falls monotonically.

For all the values of σ considered here, the interest rate falls in the period after the shock hits, and in this sense the model can account qualitatively for interest rates being lower now than during the early 1980s. However, the amount further that the interest rate is expected to fall, and where it will settle relative to its previous value, depend crucially on σ, and in particular whether it is bigger or smaller than one. Furthermore, the model can only account for rising house prices and debt when σ is below one. This result can be viewed in one of two ways. Either the model is 'wrong', or at least insufficiently general to account for the facts it is trying to explain without particular values for key parameters. Or it helps to identify a value for σ that is in line with the range of estimates reported in Chirinko (2008) and with the econometric evidence presented in section 3.8, but not with the findings in Karabarbounis and Neiman (2014).

Parameters of the utility function

The curvature parameters in the utility function are important because, as discussed in section 3.3, there could be no effect on interest rates in a model with overlapping generations and infinitely elastic utility (and hence savings) functions. Figure 3.15 shows how the effect of p on the steady state of the model depends on the the curvature parameter $\{\theta\}$ in the utility function. In these experiments we recalibrate β to hit the same initial interest rate but leave fixed the other parameters, in particular the utility function parameters $\{\gamma, \phi\}$. The figure shows that, if the utility function is more inelastic (setting $\theta = 1.5$), the interest rate falls by less, while house prices and household debt fall by somewhat more. This is in sharp contrast to the simplified model presented in section 3.3, which predicts that the interest rate varies by less when the utility function is more elastic. The reason for the discrepancy is the addition of housing to the model. If we omit housing from the model (see figure 3.16), the effect of the curvature of the utility function is in line with section 3.3: more elastic utility means that the interest varies by less.

3.6.3 Timing conventions in the model

Time periods in this model are 20 years long. Within-period timing conventions may accordingly have an important effect on the dynamics of the model. In particular, in the baseline simulations above, savings accumulated in period t are assumed to be turned into capital goods at the end of period t, at period t prices, and yield a physical return in period $t+1$ before implicitly being transformed back into intermediate goods at period $t+1$ prices. A reasonable alternative would be to assume that savings made in period t are transformed into capital goods at the beginning of the following period $t+1$, at the prices extant in that period, and then transformed back into intermediate goods in the same period at the same prices. There would then be no price depreciation component in interest rates.

Figure 3.17 shows the dynamic effects of a change in capital goods prices when the timing convention is altered in this manner. In a steady state with constant capital goods prices, these timing conventions clearly would not affect the steady state, so only the dynamic solution is shown. The most striking difference is in the path of interest rates, which fall in the period before the shock, rise above their pre-shock value and then fall again to their new, lower steady state value. Depending on when one dates the fall in capital goods prices that took place over this period, this pattern may rationalise the relatively low world interest rates observed in the early to mid-1970s. Figure 3.6 shows that the fall in capital goods prices was relatively steady over the period between the mid-1970s to mid-1990s, such that either timing assumption is reasonable in a model with 20-year periods.

3.7 Extensions

This section of the paper extends the model in various directions. Subsection 3.7.1 adds a bequest motives to households' objective functions. Subsection 3.7.2 looks at the inequality that arises between agents when only a subset have a bequest motive, and how this changes when the relative price of capital moves. Subsection 3.7.3 extends the model to study the behaviour of asset prices and the external balance in a small open economy.

3.7.1 Bequests

In the baseline model, households spend all their wealth by the end of their lives, including their housing wealth. In practice, bequests form a large part of households' total resources and a large fraction of GDP is bequeathed in any one year (Piketty (2011)). Retirees often live in owner-occupied housing until the end of life (Yang (2009). These features can be introduced into our framework by adding bequests $\{b^{'}, b\}$ respectively given and received to the utility function and budget constraints as follows[9]

$$U(c_1, c_2^{'}, c_3^{''}, h, b) = \frac{1}{1-\theta}\left(c_1^{1-\theta} + \beta_2 c_2^{'1-\theta} + \beta_3 c_3^{''1-\theta}\right) + \phi\frac{h^{1-\gamma}}{1-\gamma} + \xi\frac{b^{'1-\zeta}}{1-\zeta} \tag{3.16}$$

$$c_1 + hp_h + S_1 = \eta W \tag{3.17}$$

$$c_2^{'} + S_2^{'} = (1-\eta)W + (1+r)S_1 + b \tag{3.18}$$

$$c_3^{''} + b^{'} = (1+r^{''})S_2^{'} + hp_h \tag{3.19}$$

[9]These are 'warm glow' preferences over bequests. Households still care about the consumption value of their assets in retirement, because they evaluate their bequests in consumption rather than utility terms. Adding later generations' utility directly to the utility function would collapse the model into an infinite horizon setup.

Again forming and solving the Lagrangean we get Euler equations thus

$$\frac{c_1^{-\theta}}{(1+r')(1+r'')} = \frac{\beta_2 c_2^{'-\theta}}{(1+r'')} = \beta_3 c_3^{''-\theta} = \xi b^{'-\zeta} \tag{3.20}$$

Piketty (2011) finds that bequests in France are around 15% of GDP. With log utility bequests will be a constant fraction of old-age consumption. Consistent with this, we set $\xi = 0.75\beta_3$. Figures 3.18 and 3.19 show that, qualitatively speaking, the steady state and dynamic solutions of the model are unaffected by this change, although the change in interest rates is attenuated somewhat. The levels of debt and house prices are nonetheless higher at any given level of p. Households save for bequests much like they save for retirement consumption. By the time the middle-aged receive their requests, they themselves are on the cusp of retirement and are therefore no long debtors. Anticipating bequests in middle age, the young accumulate more debt and push up house prices.

3.7.2 Heterogeneous agents

In the versions of the model presented above, the only dimension along which agents are heterogeneous is age. The change in p affects different generations differently, but there can be no intra-generational inequality. In reality, inherited wealth is distributed highly unequally across the members of any given generation. The asset-price consequences of a change in p are accordingly likely to have consequences for intra-generational inequality. To study this, we simulate a version of the model in which the population is divided into two kinds of dynasty - life-cycle households without a bequest motive, as in the baseline model, and households with a bequest motive. For illustrative purposes, we set the proportions of each kind to one-half, and apportion to them equal labour endowments (and thus labour income).

Figure 3.20 sets out the dynamic consequences of a change in p on the ratio of the consumption of households with a bequest motive to life-cycle households. In the steady state before the fall in p, households with the bequest motive accumulate more wealth (about one-and-a-half times as much) and consume about 13% more than those without it, even though they have the same labour income. In the long run, as in the baseline and dynastic models, the fall in p lowers the real interest rate and raises house prices.

Consumption inequality falls because lower interest rates reduce the returns to inherited wealth, and thus the extra consumption that can sustainably be financed from bequests. But the fall in interest rates revalues non-produced assets (land) such that the wealth-income ratio rises. Households that do not receive bequests have only life-cycle saving, which falls somewhat, as a source of lifetime wealth, so wealth inequality increases. In this model, therefore, and in contrast to Piketty et al. (2014), r moves in the *opposite* direction to wealth inequality, but moves in the same direction as consumption inequality.

The dynamic consequences of the shock to p are non-monotone and vary a great deal according to date of birth. Households in bequest-giving dynasties who are young on the eve of the shock do especially well, because they receive a disproportionate slice of the one-off capital gains that accrue to asset holders.

3.7.3 Open economy

The baseline model in this paper treats the industrialised world as a large, closed economy, with a view to explaining a global trend. The world is of course composed of many economies which are open to trade in goods and financial assets with each other as well as with emerging markets. For any one of these countries, foreign assets are an important store of value, such that we might expect the external position of any given economy to depend on the domestic relative price of its capital goods. The real interest rate in any given country may accordingly not depend to a great extent on the relative price of capital goods in that country, even if the interest rate and the relative price of capital goods are linked at a global level. Furthermore, testing the implications of the model presented above is hampered by the fact that, at a global level, we only have one very short time series when time is denominated in model units, whereas an open-economy version will lend itself to testing along the cross-country dimension. Last but not least, the low-frequency behaviour of the current account dynamics is of independent interest.

To study the open-economy implications of capital goods prices, and to take our model more readily to the data, we therefore consider a simple open-economy version of our baseline model. We assume that intermediate goods, and financial claims denominated in them, are perfectly tradable across borders. They are transformed into consumption and investment goods at home using the domestic technology. This technology can vary across countries, and hence so can the relative price of capital goods. Each country takes the world real interest rate as given. Relative to the closed-economy baseline model set out in section 3.4, all prices and quantities except the interest rate r acquire country subscripts i, and the only equation to materially change is the asset-market clearing condition (3.14)

$$S_{1i} + S_{2i} - K_i^{'} p_{Ki} = NFA_i \tag{3.21}$$

where NFA_i denotes the net foreign assets of country i.

Figure 3.21 shows the steady-state effect of changing the relative price p_{Ki} of capital goods in country i holding the world interest rate fixed and starting from a position in which, at $p_{Ki} = 1$, the net foreign asset position is zero. The experiment can be thought of as describing the behaviour of a small open economy in which the relative price of capital goods falls by more than the world average. The blue lines show the impact in a closed economy (i.e. the baseline simulation), and the green lines show the small open economy. The top left panel reminds us that, by construction, interest rates do not change. The fall in capital goods prices leads the corporate sector to demand less in the way of investable funds, as we see by the fall in the nominal investment rate (top right panel). Middle-aged households' savings go overseas rather than to young households when the economy is open, such that net foreign assets rise from zero to about 60% of GDP (bottom right), while the household debt ratio is essentially unchanged. House prices relative to incomes (bottom left panel) rise nonetheless, albeit by about 10% rather than the 30% we see in the closed-economy case. This is because the fall in investment and the rise in net foreign assets afford a rise in the consumption-GDP ratio, and consumption and housing demand are positively related.

Figure 3.22 shows the dynamic behaviour of a small open economy when faced with a 30% fall in the relative price of capital over one period as above, but where world (and hence domestic) real interest rates are fixed. The blue line (read against the left-hand axes) show the relative price of capital, and the green lines show the behaviour of the closed-economy baseline model for ease of comparison. The top left panel shows that, by assumption, the interest rate in the small open economy is unchanged. The two most striking results are in the top and bottom right panels. House prices rise ahead of the fall in capital goods prices (middle panel), in anticipation of higher housing demand after the shock. This raises household debt as young households seek to smooth consumption in the face of higher house prices. However, once capital goods prices fall, output and wages rise sharply, such that the ratio of house prices to wages falls and young households need less debt. Net foreign assets rise sharply around the shock as the interest rate that would prevail in a closed economy falls below the world interest rate: the fall in domestic capital goods prices makes available savings that can fund foreign investments.

3.8 Econometric evidence

This section of the paper confronts the predictions of the model with econometric evidence. Subsection 3.8.1 details evidence on the elasticity of substitution between capital and labour σ. Subsection 3.8.2 looks at the model's predictions for house prices, debt and the external balance.

3.8.1 Evidence on the elasticity of substitution between capital and labour

Karabarbounis and Neiman (2014) present a model in which, like the model above, the relationship between the labour share and relative price of investment goods depends on the elasticity of substitution between capital and labour

$$\frac{s_{Lj}}{1-s_{Lj}}\hat{s_{Lj}} = \gamma + (\sigma - 1)\,\hat{\xi}_j + u_j$$

where s_{Lj} is the labour share in country j, ξ_j is the relative price of investment, and hats denote low-frequency, country-specific time trends. They regress the time trend in the labour share on the time trend in the relative price of investment goods, obtain a coefficient averaging 0.28 across datasets and and infer that the elasticity of substitution between capital and labour is 1.28.

Their model also implies an analogous relationship between the nominal investment rate $\frac{I_N}{Y}$ and the relative price of capital [10]

$$\widehat{\frac{I_N}{Y}}_j = \tilde{\gamma} + (1 - \sigma)\,\hat{\xi}_j + \tilde{u_j}$$

Similarly, the baseline model presented in this paper predicts that, when the supply of funds is perfectly interest-elastic, as in the small open economy case in section 3.7.3 or a model with infinite-horizon consumers, the investment rate and the price of investment are related as follows in the steady state

$$\widehat{\frac{pI}{Y}} = c + (1 - \sigma)\,\hat{p}$$

These equations motivate regressing the nominal investment rate on the relative price of investment goods across countries as a way of quantifying the crucial parameter σ: if the elasticity is greater than one, we would expect a negative relationship between the relative price of capital and the nominal investment rate as the quantity falls (rises) proportionally more than the price falls (rises). We employ the dataset in Karabarbounis and Neiman (2014) to this end.

Table 3.1 sets out the results, showing the central estimate for the coefficient on the relative price of capital, along with its standard deviation and implied confidence intervals for σ. Results are shown for three different estimators - robust regression (the estimator used by Karabarbounis and Neiman) and OLS, both on country-specific time trends, and panel fixed effects, on country-year observations - and for two different sources for the relative price of investment (Penn World Tables and the World Bank). In all cases, following Karabarbounis and Neiman, the sample is restricted to contain only countries with 15 or more years of observations, and contains the corporate investment rate where it is available, and its whole-economy analogue where it isn't. The results are clearly sensitive to the choice of estimator and sample. Using Karabarbounis and Neiman's preferred robust regression methodology, the central estimates of σ are 0.2 and 0.3, depending on the source data for relative prices. None of the confidence intervals for σ contain 1. So on the face of it, these results do not corroborate the

[10] See Appendix 3.B for derivation

greater-than-unit elasticity presented in Karabarbounis and Neiman (2014), instead leaning toward a less-than-unit elasticity in line with most estimates in the literature (Chirinko (2008)). Nominal investment rates are typically increasing in the relative price of capital, suggesting as a model with a less-than-unit elasticity of substitution.[11]

3.8.2 Testing the model's predictions on house prices, household debt and net foreign assets

The baseline closed-economy model predicts that countries with relatively low capital-goods prices will have low steady-state interest rates. In the open-economy version, these lower shadow real rates translate into positive net foreign asset positions. And in a state state with growth in nominal GDP, these more positive external positions would necessitate more positive current account balances. Furthermore, around the time of the transition, the bigger the fall in capital goods prices, the more positive the current account balance. The model therefore predicts a negative relation between capital goods prices and the current account, both in the steady state and during the transition.

The closed-economy model also predicts a negative steady-state relationship between capital goods prices and household debt. There is no such long-run relationship between debt and relative prices - taking world interest rates as given - in the open economy model. Finally, both models predict a negative relationship between relative prices and real house prices.

Given the rising, but on average intermediate, degree of de facto capital-account openness in the world economy over the past forty years, it is not clear a priori whether the open- or closed-economy versions of the model will turn out to be better approximations to the real world.

With these caveats in mind, we can take the model to the data in a manner analogous to the previous subsection, regressing the household debt-GDP ratio, real house prices and the current account-GDP ratio on the level of the relative price of capital. We use the same three estimators used in the previous subsection: panel fixed effects on annual country-year observations; OLS on country-specific time trends in relative capital goods prices and the

[11] Appendix 3.C sets out a model to recognise the seemingly conflicting results obtained from estimating σ from the labour share and the nominal investment rate.

current account, the household debt-GDP ratio, and real house prices; and robust regression on the same.[12]

Tables 3.2, 3.3 and 3.4 set out the results. Table 3.2 shows that, for panel fixed effects and robust regression on time trends, we find a large and significant negative relationship between the relative price of capital goods and household debt. This is in line with the prediction of the closed-economy version of the model: lower capital goods prices reduce interest rates, stimulating household borrowing. For OLS on time trends, we find a positive but insignificant effect. Table 3.3 displays a similar pattern - large negative and significant coefficients when using panel fixed effects; negative coefficients when using robust regression, but which are significant for only one of the measures of relative prices; and inconclusive results from OLS. Finally, table 3.4 shows the results for the current account. Here the results for cross-country trends are inconclusive, but panel fixed-effects deliver a significant negative coefficient, in line with the predictions of the open-economy model.

Overall, cross-country econometric analysis provides qualified support for the assumptions in and predictions of the model. There is strong evidence that nominal investment rates are increasing in the relative price of capital p, and thus that the key elasticity parameter σ is below 1. There is some evidence that household debt and house prices are both negatively related to p, consistent with the predictions of the closed-economy model. And there is weak evidence that the current account is negatively related to p, consistent with the open economy model. But taken together, the economies in our sample appear to have behaved more like financially closed economies on average over the period in question.

3.9 Conclusion

This paper presents a model of 'secular stagnation' - persistently low real interest rates - driven by the interaction of life-cycle savings motives and an improvement in the technology for producing investment goods. The model is complementary to other explanations for low real interest rates that rely on

[12]At any point in time, the behaviour of the current account in country i will depend on the path of capital goods prices in country i relative to other countries. So in the panel regression of current account, we first condition p on time and country fixed effects

demographics, emerging markets and inequality. Using standard parameter values and the observed path for capital goods prices over the past few decades, it is able to reproduce part of the rising-falling pattern in real interest rates, the falling ratios of nominal investment and capital to GDP, and the rise in household debt observed across the industrialised world.

The dynamic simulations predict that the real interest rate will stay low, even if the relative price of capital goods has stopped falling. The model suggests that limiting the accumulation of household debt would have made the fall in interest rates larger. And it suggests that the rise in the wealth-income ratio the shock has produced may have increased inherited wealth inequality, even though interest rates have fallen.

The model has important normative and positive implications. First, it and the accompanying econometric results provide additional evidence on a below-unit elasticity of substitution between capital and labour. Secondly, it stresses the point that capital goods are simultaneously a social savings technology, the means of production, and a produced asset in themselves. So changes in the way that capital goods are produced will have implications for other stores of value, such as housing, land, public debt, and any 'bubbly' asset. And thirdly, it suggests that real interest rates may remain low, or have further to fall, even if the relative price of capital goods has stopped falling. Fiscal and monetary rules that are calibrated implicitly on the real interest rate may need to be revised as a result.

3.10 Tables and charts

Table 3.1: Estimates of the elasticity of substitution σ

Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
Log(p)	0.491***	1.121***	0.776***	0.290***	0.999***	0.695***
	[0.04]	[0.21]	[0.17]	[0.04]	[0.25]	[0.16]
$\hat{\sigma}$	0.509	-0.121	0.224	0.71	0.001	0.305
$\hat{\sigma_H}$	0.589	0.299	0.564	0.79	0.501	0.625
$\hat{\sigma_L}$	0.429	-0.541	-0.116	0.63	-0.499	-0.015
N	1632	54	54	1643	52	52
no. of countries	99			100		

Table 3.2: Regression of household debt on relative price of capital

Left-hand side variable	Household debt/GDP					
Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
log(p)	-0.993***	0.702	-0.779***	-1.179***	0.571	-0.888***
	[0.05]	[0.65]	[0.25]	[0.07]	[0.72]	[0.30]
N	535	18	18	551	18	18
no. of countries	21			21		

Table 3.3: Regression of real house prices on relative price of capital

Left-hand side variable	Real house prices					
Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
log(p)	-1.082***	0.121	-0.672	-0.976***	-0.277	-1.520**
	[0.10]	[0.89]	[0.79]	[0.12]	[0.91]	[0.65]
N	535	18	18	551	18	18
no. of countries	21			21		

Table 3.4: Regression of current account on relative price of capital

Left-hand side variable	Current account/GDP					
Dataset	Panel	Time trends		Panel	Time trends	
Estimator	FE	OLS	Robust	FE	OLS	Robust
RHS source		PWT			WDI	
log(p)	-0.055***	0.006	0.020	-0.025**	0.025	0.028
	[0.01]	[0.05]	[0.05]	[0.01]	[0.05]	[0.05]
N	1004	35	35	992	34	34
no. of countries	50			51		

Figure 3.1: 10-year real interest rates

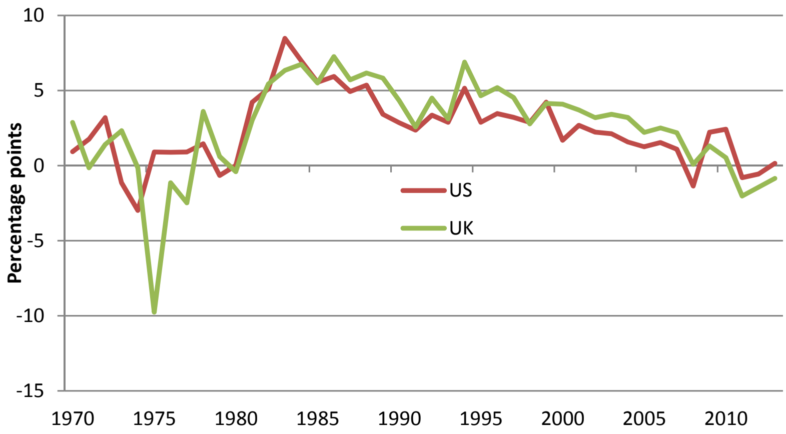

This figure shows estimates of ex-ante 10-year real interest rates for the US and UK, calculated as the difference between nominal government bond yields and model-based estimates of inflation expectations taken from Figure 3.2 of IMF (2014).

Figure 3.2: HH debt-GDP ratio, % of GDP

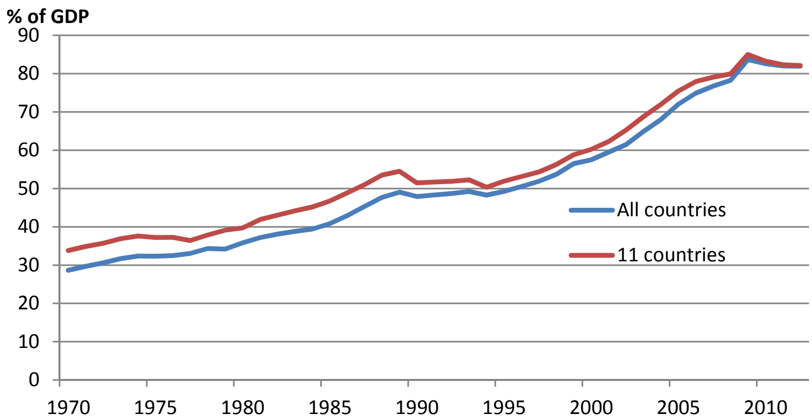

This figure shows the change in the ratio of household debt to GDP since 1970 for a broad sample of industrialised countries and a restricted sample of 11 countries (Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US). The source for household debt is the BIS, and the source for GDP is OECD StatBase. The chart is constructed from an unbalanced panel of data by running a fixed-effects panel regression of the household debt ratio on year dummies, then adding the dummy for each year to the intercept of the equation. This allows other countries to affect the change in the ratio in years after they have been added to the sample.

Figure 3.3: Nominal investment-GDP ratios

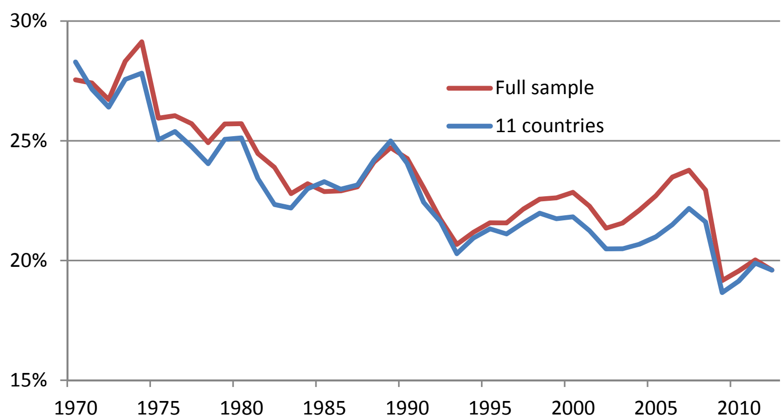

This figure shows a simple average across 24 OECD countries and a restricted sample of 11 countries (Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US) of the ratio of nominal gross capital formation to nominal GDP. The source is OECD Statbase.

Figure 3.4: Nominal and real capital stock-GDP ratio

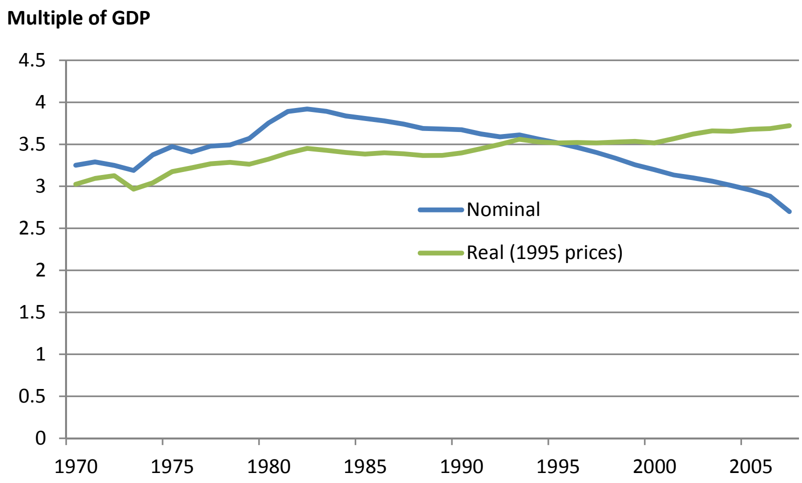

This figure shows the average nominal and real capital-output ratios for Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US. The source is the 'All capital input files' file from the November 2009 release of the EU-KLEMS database. The nominal capital stock was computed for each country as the product of the real capital stock at 1995 prices and the gross fixed capital formation price index rebased to 1995, and then divided by nominal GVA taken from EU-KLEMS to give the nominal capital output ratio. Country-year observations were regressed on country and year dummies, and average index in year t was obtained as the sum of the intercept and the dummy for year t. The real capital-output ratio was constructed analogously, then by rebasing the average ratio of the real capital stock and real GDP to the 2005 nominal ratio.

Figure 3.5: Real investment-GVA ratio, 11 industrialised countries, 2007=1

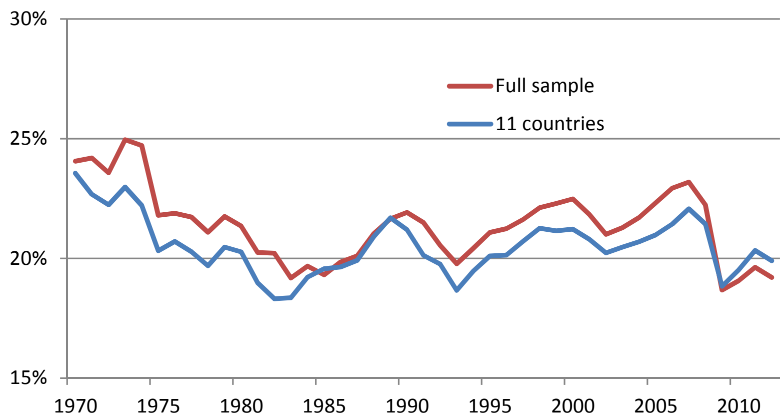

This figure shows a simple average across 24 OECD countries and a restricted sample of 11 countries (Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US) of the ratio of gross capital formation to GDP, both at constant prices. The source is OECD Statbase.

Figure 3.6: Price of investment relative to consumption

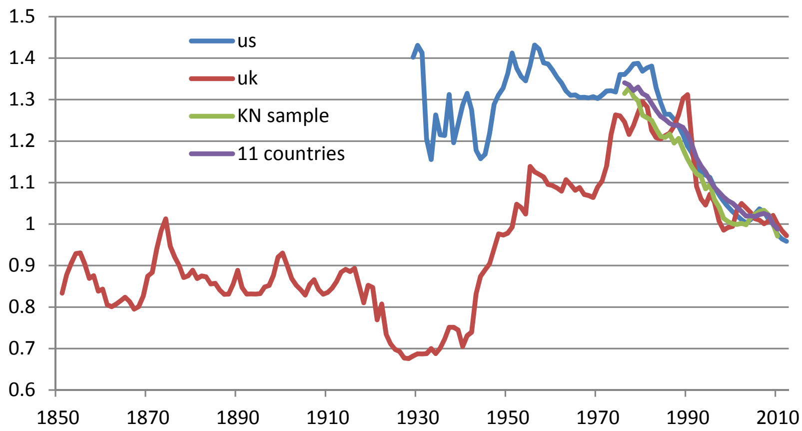

This figure shows four series of the relative price of investment to consumption. The red line is the ratio of the deflators of gross capital formation and private consumption in the UK, taken from the Bank of England's internal long-run database. The blue line is the analogous ratio for the US, taken from the FRED database. The green line is constructed from an unbalanced panel of data by running a fixed-effects panel regression on the data for all countries in Karabarbounis and Neiman (2014), then adding the dummy for each year to the intercept of the equation. The purple line is constructed in the same way for Australia, Austria, Denmark, Finland, Germany, Italy, Japan, Netherlands, Sweden, the UK and the US.

Figure 3.7: Simple savings-investment diagram

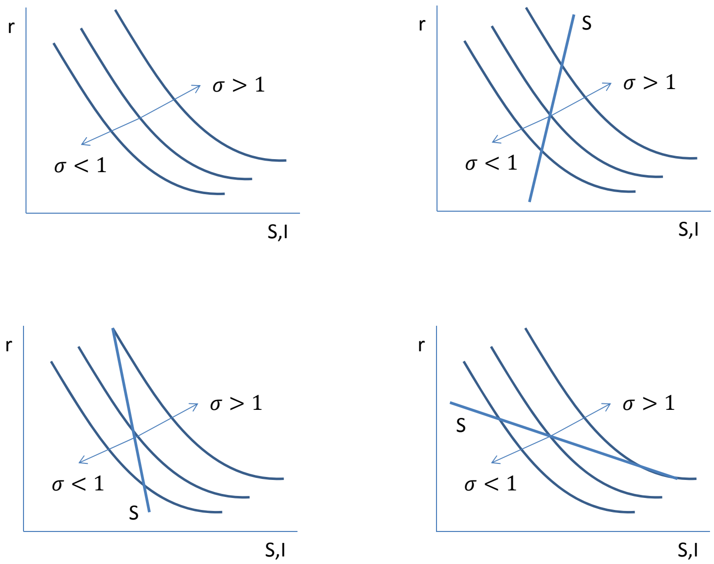

Figure 3.8: Steady state, baseline setup

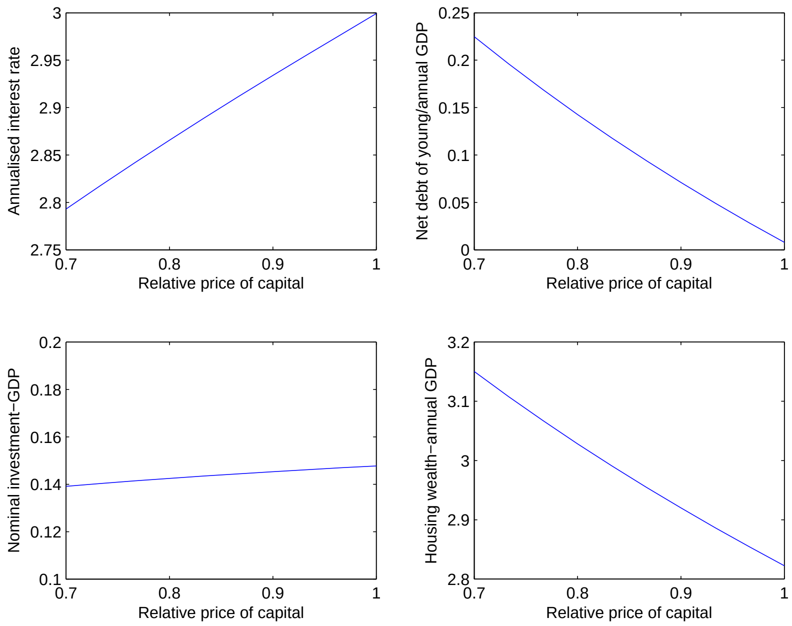

Figure 3.9: Dynamic solution, baseline setup

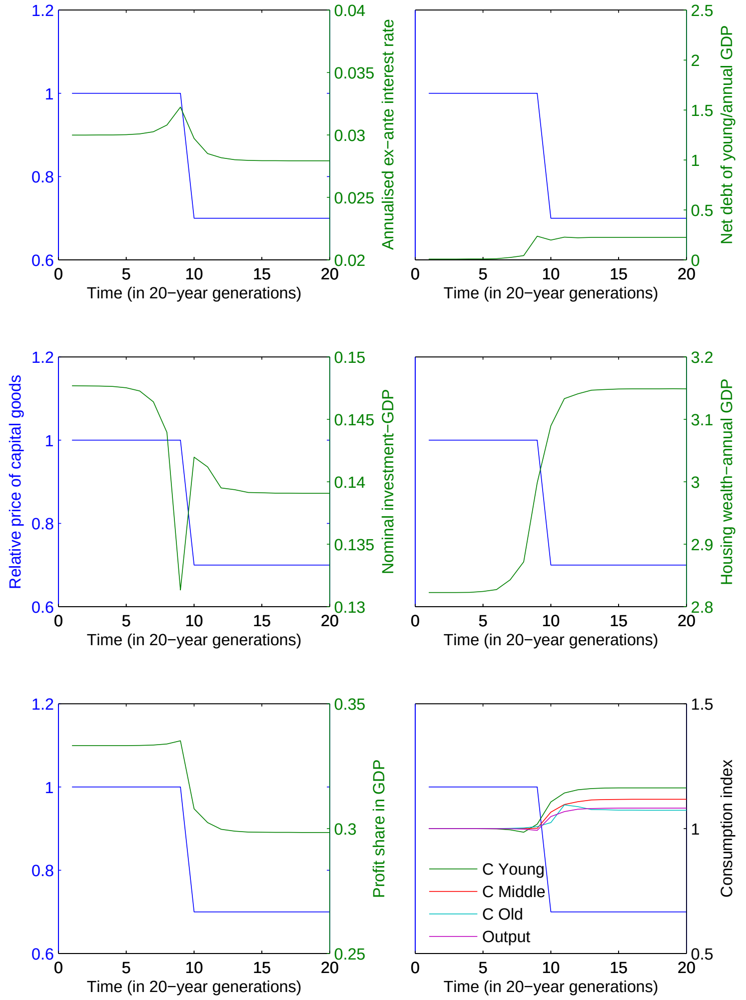

Figure 3.10: Steady state, varying availability of household debt

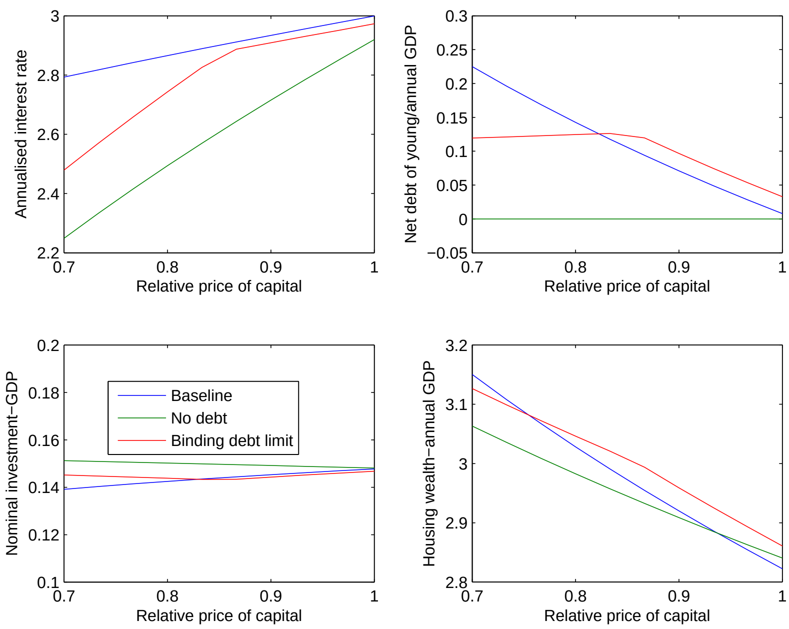

Figure 3.11: Dynamic solution, no household debt

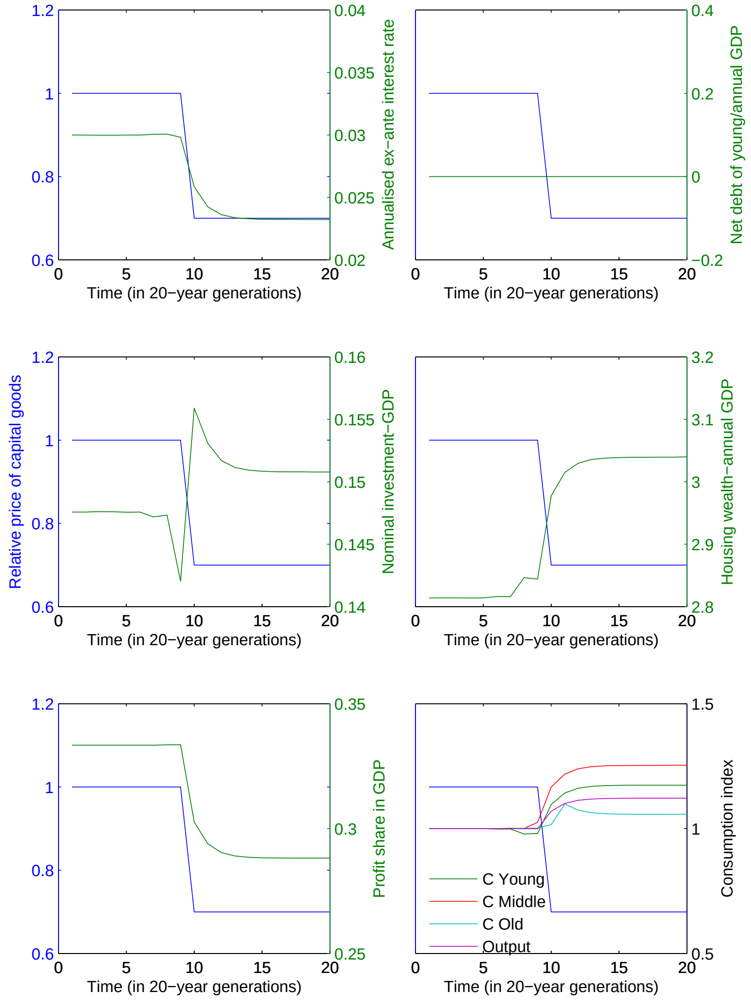

Figure 3.12: Steady state, varying curvature of the production function

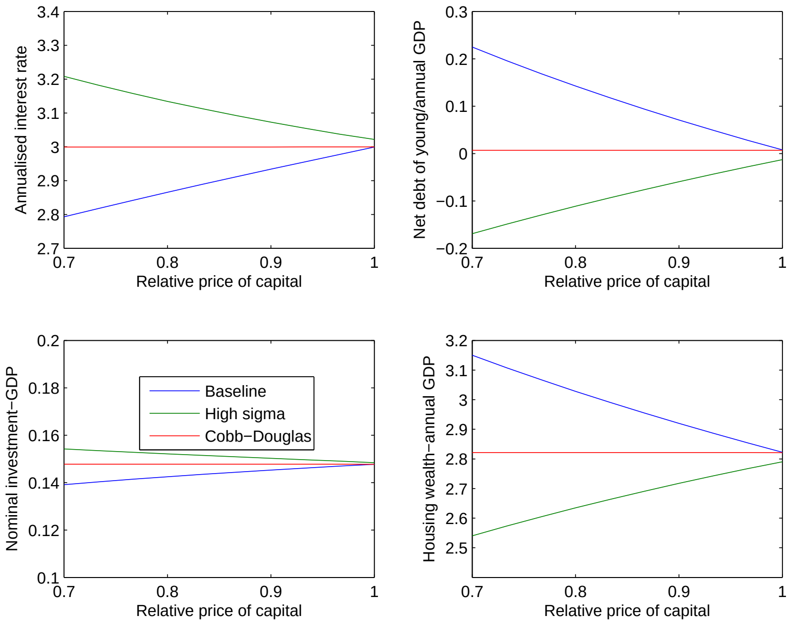

Figure 3.13: Dynamic solution, Cobb-Douglas production function

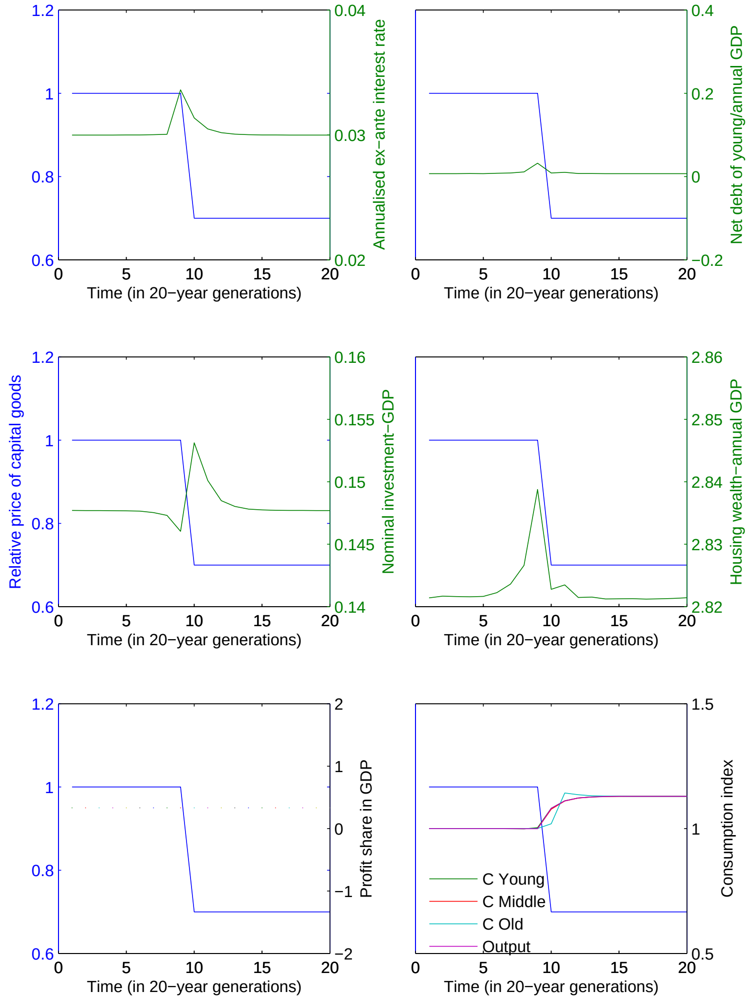

Figure 3.14: Dynamic solution, highly elastic production function

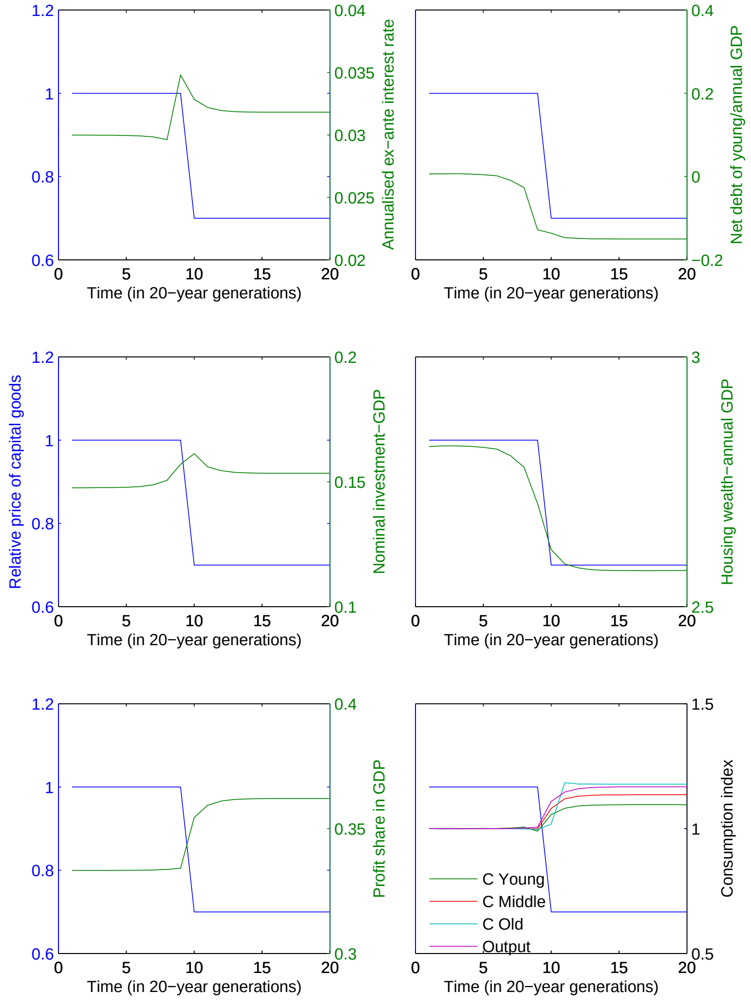

Figure 3.15: Steady state, varying curvature of the utility function

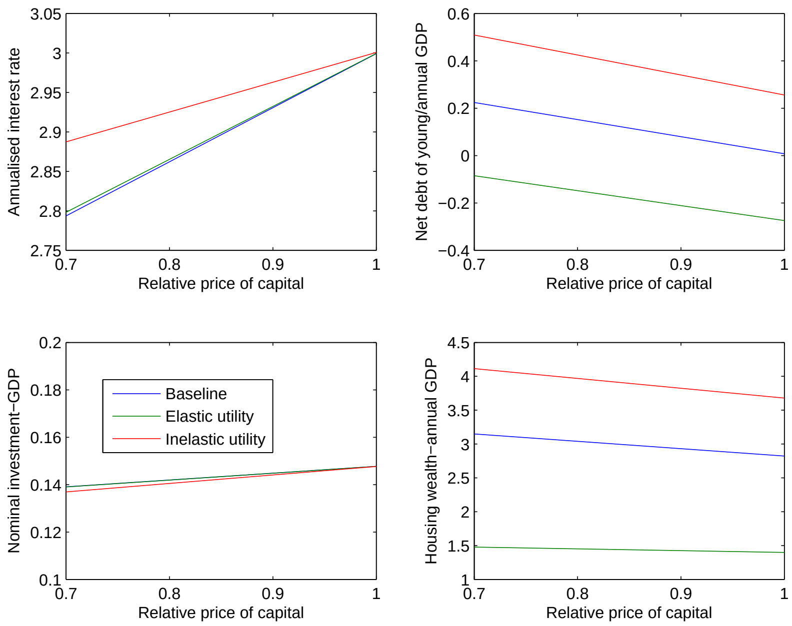

Figure 3.16: Steady state, no housing, curvature of the utility function

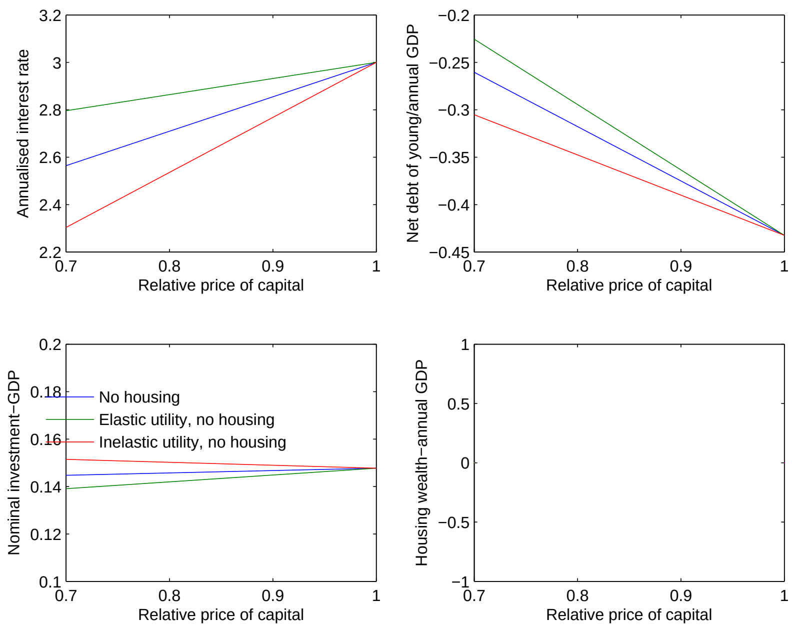

Figure 3.17: Dynamic solution, alternative timing convention

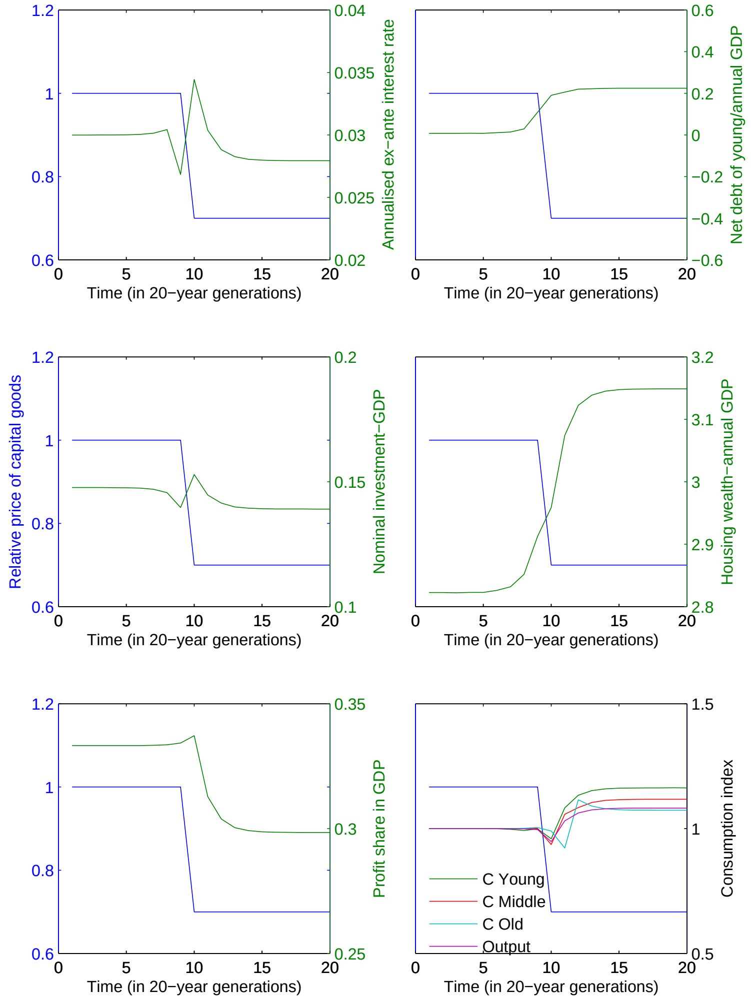

Figure 3.18: Steady state, bequests and heterogeneous agents

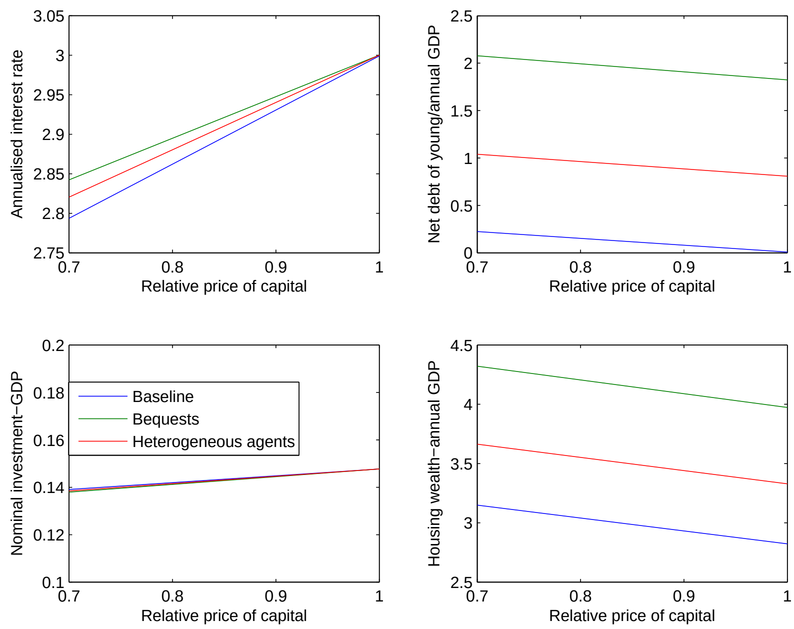

Figure 3.19: Dynamic solution with bequests

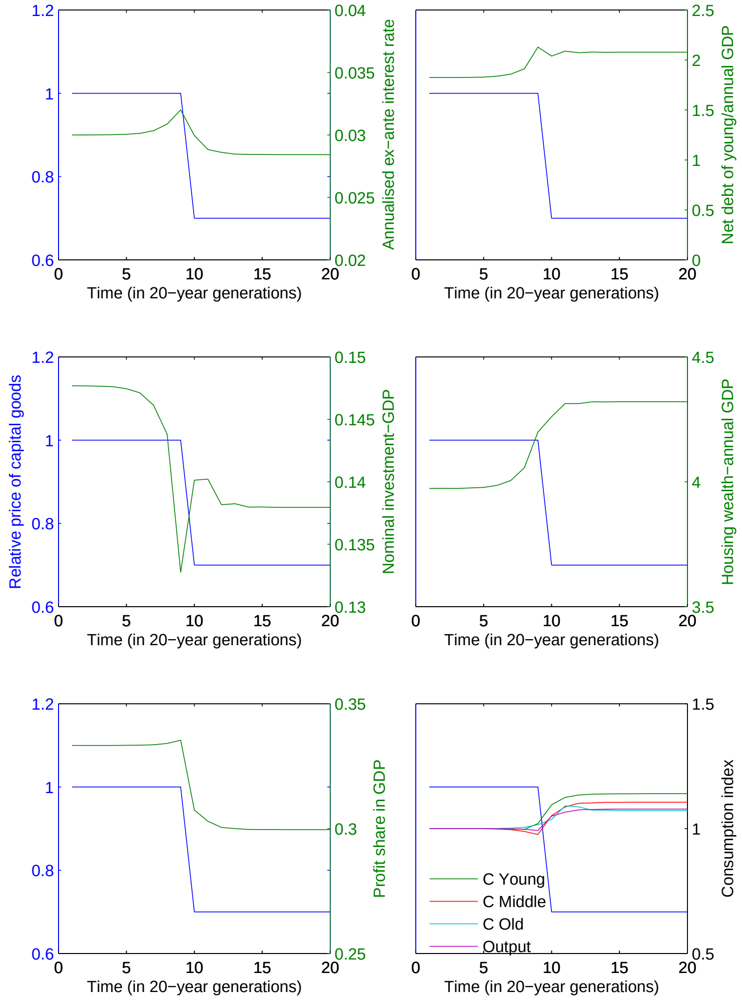

Figure 3.20: Inequality within generations

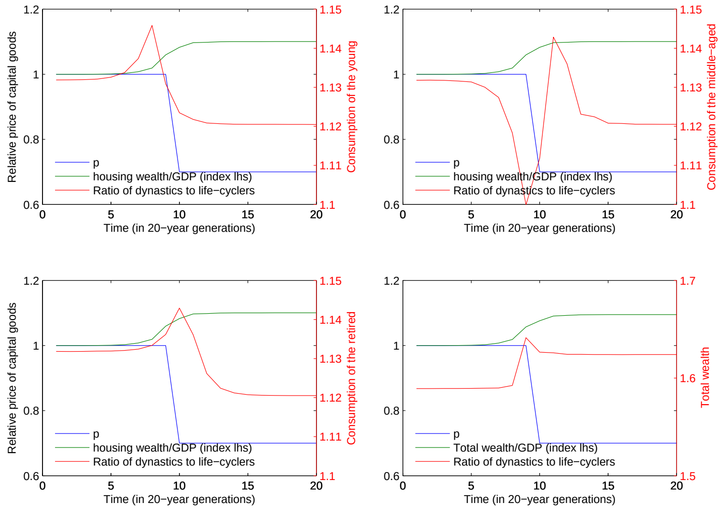

Figure 3.21: Steady state, small open economy

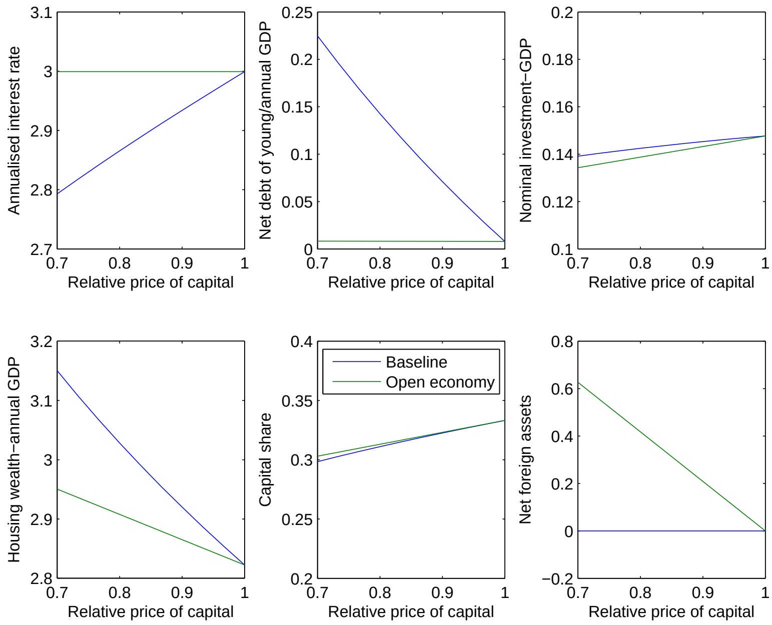

Figure 3.22: Dynamic solution, small open economy

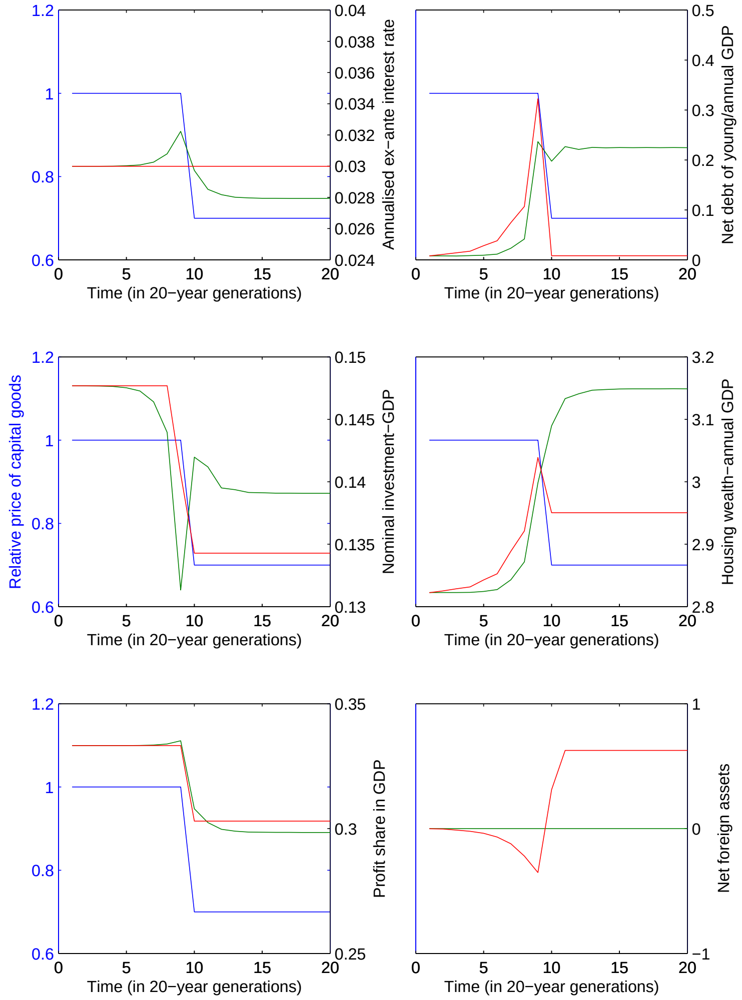

Figure 3.23: Region of parameter space in which labour share and investment are both increasing in p

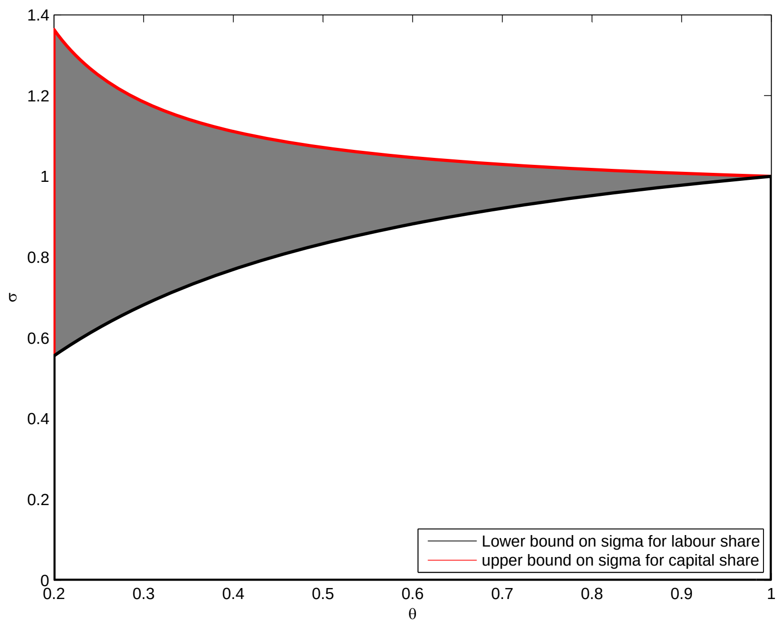

The shaded area shows the region of the parameter space of the three-factor production function set out in Appendix 3.C in which the labour share and the nominal investment rate are both increasing in the relative price of capital

3.A Derivations

3.A.1 Saving

Here we derive the savings schedule for the simple two period model in Section 3.3

$$\begin{aligned} c_1^{-\theta} =& \beta(1+r)c_2^{-\theta} \\ c_2 =& [\beta(1+r)]^{\frac{1}{\theta}} c_1 \\ S = W - c_1 =& \frac{c_2}{(1+r)} \\ =& W\frac{\beta^{\frac{1}{\theta}}(1+r)^{\frac{1}{\theta}-1}}{1+\beta^{\frac{1}{\theta}}(1+r)^{\frac{1}{\theta}-1}} \end{aligned}$$

3.A.2 Investment schedule in $\{s, r\}$ space

Here we derive the investment schedule for the simple two-period model in Section 3.3. From the CES production function we have

$$\begin{aligned} r+\delta =& \frac{1}{p}\frac{\partial Y}{\partial K} \\ =& \frac{1}{p}\alpha\left(\frac{Y}{K}\right)^{\frac{1}{\sigma}} \\ \frac{(r+\delta)p}{\alpha} =& \left(\frac{K}{Y}\right)^{-\frac{1}{\sigma}} \\ \frac{K}{Y} =& \left[\frac{\alpha}{p(r+\delta)}\right]^{\sigma} \end{aligned}$$

From CRS and Euler's theorem we have

$$\begin{aligned} W + K\frac{\partial Y}{\partial K} =& Y \\ \frac{W}{Y} =& 1 - \frac{K}{Y}\frac{\partial Y}{\partial K} \\ =& 1 - \alpha\left(\frac{K}{Y}\right)^{\frac{\sigma-1}{\sigma}} \end{aligned}$$

Now we can rewrite the saving rate as

$$
\begin{aligned}
s =& \frac{S}{W} \\
=& \frac{\left(\frac{pK}{Y}\right)}{\left(\frac{W}{Y}\right)} \\
=& \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}}
\end{aligned}
$$

So the derivative of the saving rate with respect to the price of capital goods p is the same sign as $1 - \sigma$.

Which way does the interest rate schedule slope? There are two effects. The effect in the numerator is negative for the standard reasons: for given capital goods prices, more savings reduces the marginal product of capital and hence the interest rate. The effect in the denomnator is of ambigious sign, and comes through the labour share (for a Cobb-Douglas function $\sigma - 1 = 0$ it is absent). For low σ, an increase in r reduces the denominator, raising the quotient. This is because we are expressing the saving rate as a fraction of wages and when $\sigma < 1$, higher interest rates are associated with a lower labour share. To save enough for a given volume of capital goods, a lower labour share must mean a higher saving rate.

For reasonable parameter values, the effect on the numerator will dominate, such that the investment schedule slopes down in $\{s, r\}$ space. To see this, differentiate the schedule with respect to r

$$\begin{aligned}\frac{ds}{dr} =& \frac{-\sigma p^{1-\sigma}\alpha^{\sigma}(r+\delta)^{-\sigma-1}}{1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}} \\ &- \frac{p^{1-\sigma}\left[\frac{\alpha}{(r+\delta)}\right]^{\sigma}}{\left(1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}\right)^2}\left(-\alpha p^{1-\sigma}\alpha^{\sigma-1}(1-\sigma)(r+\delta)^{-\sigma}\right) \\ =& \frac{-\sigma}{r+\delta}s + s\frac{\alpha p^{1-\sigma}\alpha^{\sigma-1}(1-\sigma)(r+\delta)^{-\sigma}}{1-\alpha p^{1-\sigma}\left[\frac{(r+\delta)}{\alpha}\right]^{1-\sigma}} \\ =& \frac{-\sigma}{r+\delta}s + (1-\sigma)s^2 \\ =& s\left((1-\sigma)s - \frac{\sigma}{r+\delta}\right)\end{aligned}$$

3.B Estimating σ from the nominal investment share in the Karabarbounis and Neiman (2014) model

The Karabarbounis and Neiman (2014) model decomposes income into the capital share s_K, the labour share s_L and markups μ

$$\mu(s_K + s_L) = 1$$

Taking logs and then the derivative with respect to time gives

$$\begin{aligned}0 = \frac{d}{dt}\log(\mu(s_K+s_L)) =& \hat{\mu} + \frac{1}{s_K+s_L}\left[\frac{ds_K}{dt} + \frac{ds_L}{dt}\right] \\ =& \hat{\mu} + \mu s_K \hat{s_K} + (1-\mu s_K)\hat{s_L}\end{aligned}$$

Following Karabarbounis and Neiman (2014), we set $\mu = 1, \hat{\mu} = 0$ and get

$$s_K\hat{s_K} + s_L\hat{s_L} = 0$$

and therefore we can rewrite the left-hand side of their equation (19) as follows

$$\frac{s_L}{1-s_L}\hat{s_L} = \frac{s_L}{1-s_L}\frac{-s_K\hat{s_K}}{s_L} = -\hat{s_K}\frac{s_K}{1-s_L} = -\hat{s_K}$$

Karabarbounis and Neiman (2014) write the capital share as

$$\begin{aligned} s_K &= \frac{RK}{Y} \\ &= \frac{1}{\mu} F_K \frac{K}{Y} \\ &= \frac{\alpha_K A_K^{\frac{\sigma-1}{\sigma}}}{\mu} \frac{Y}{K}^{\frac{1}{\sigma}} \frac{K}{Y} \\ &= \frac{\alpha_K A_K^{\frac{\sigma-1}{\sigma}}}{\mu} \frac{K}{Y}^{\frac{\sigma-1}{\sigma}} \end{aligned}$$

If we assume away changes in technology, capital shares and markups we have

$$\hat{s_K} = \frac{\sigma-1}{\sigma} \widehat{\left(\frac{K}{Y}\right)}$$

In the steady state, the nominal investment rate is proportional to the nominal capital-output ratio

$$\frac{I_N}{Y} = \frac{p\delta K}{Y}$$

so that

$$\frac{K}{Y} = \frac{I_N}{Y} \frac{1}{p\delta}$$

If we assume away changes in depreciation rates we have

$$\hat{s_K} = \frac{\sigma-1}{\sigma} \widehat{\frac{K}{Y}} = \frac{\sigma-1}{\sigma} \left(\widehat{\frac{I_N}{Y}} - \hat{p} \right)$$

Combining these results with the estimating equation (19)

$$\frac{s_{Lj}}{1-s_{Lj}} \hat{s_{Lj}} = \gamma + (\sigma - 1) \hat{p} + u_j$$

we have

$$\begin{aligned} \frac{s_{Lj}}{1-s_{Lj}} \hat{s_{Lj}} &= -\hat{s_K} = \frac{1-\sigma}{\sigma} \left(\widehat{\frac{I_N}{Y}} - \hat{p} \right) \\ \widehat{\frac{I_N}{Y}} - \hat{p} &= \tilde{\gamma} + \frac{\sigma}{1-\sigma} (\sigma - 1) \hat{p} + \tilde{u_j} \\ \widehat{\frac{I_N}{Y}} &= \tilde{\gamma} + (1-\sigma) \hat{p} + \tilde{u_j} \end{aligned}$$

In other words, if the elasticity of substitution between capital and labour σ is greater than unity, then a fall in the relative price of capital should lead to a rise in the ratio of nominal investment to GDP, as the volume of investment rises by a greater proportion than the fall in its price.

3.C Three-factor model

This appendix sketches out an alternative production function for intermediate goods which breaks the negative relation between the labour share and the investment rate and can, at certain parameter values, deliver a labour share and a nominal investment rate which are both increasing in the relative price of capital. It then illustrates the region of the parameter space at which this is so.

3.C.1 Modified production function

Consider a production function defined over three factors - capital K, unskilled labour L and entrepreneurial labour M, which is paid in profits. K and L are aggregated first into an intermediate X, then combined with M as follows

$$\begin{aligned}
X =& G\left(K, L\right) \\
=& [(1-\alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \\
Y =& F\left(X, M\right) \\
=& [\mu M^{\frac{\theta-1}{\theta}} + (1-\mu)X^{\frac{\theta-1}{\theta}}]^{\frac{\theta}{\theta-1}} \\
Y =& F\left(G\left(K, L\right), M\right) \\
=& \tilde{F}\left(K, L, M\right) \\
=& \left[\mu M^{\frac{\theta-1}{\theta}} + (1-\mu)\left[\left[(1-\alpha)\, L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}\right]^{\frac{\theta-1}{\theta}}\right]^{\frac{\theta}{\theta-1}}
\end{aligned}$$

3.C.2 Labour share and investment rate

The nominal investment rate is proportional to the capital share at fixed interest rates, because the marginal product of capital is inversely proportional to the price of capital. So a shift inward in the investment schedule is equivalent to a fall in the capital share

$$\frac{I_N}{Y} = \frac{\delta PK}{Y} = \frac{\delta}{r+\delta} MPK \frac{K}{Y}$$

The factor share for factor Z is

$$\Omega_Z = \frac{\partial Y}{\partial Z}\frac{Z}{Y}$$

Denoting logs with lower case we have

$$\begin{aligned}
\omega_K =& A + log\left(\left(\frac{Y}{X}\right)^{\frac{1}{\theta}}\left(\frac{X}{K}\right)^{\frac{1}{\sigma}}\frac{K}{Y}\right) \\
=& A + log\left(Y^{\frac{1-\theta}{\theta}} K^{\frac{\sigma-1}{\sigma}} X^{\frac{1}{\sigma}-\frac{1}{\theta}}\right) \\
=& A + \left(\frac{1-\theta}{\theta}\right) y + \left(\frac{\sigma-1}{\sigma}\right) k + \left(\frac{1}{\sigma}-\frac{1}{\theta}\right) x
\end{aligned}$$

and (because $L = 1$)

$$\begin{aligned}
\omega_L =& A + log\left(\left(\frac{Y}{X}\right)^{\frac{1}{\theta}}\left(\frac{X}{L}\right)^{\frac{1}{\sigma}}\frac{L}{Y}\right) \\
=& A + log\left(Y^{\frac{1-\theta}{\theta}} X^{\frac{1}{\sigma}-\frac{1}{\theta}}\right) \\
=& A + \left(\frac{1-\theta}{\theta}\right) y + \left(\frac{1}{\sigma}-\frac{1}{\theta}\right) x
\end{aligned}$$

3.C.3 Parameter space

To calculate the derivative of each factor share with respect to the log relative price of capital p note that

$$\frac{d\omega_Z}{dp} = \frac{d\omega_Z}{dk}\frac{dk}{d(mpk)}\frac{d(mpk)}{dp}$$

For a fixed interest rate, the user cost condition implies that $\frac{d(mpk)}{dp} = 1$. The middle term in the chain is negative and depends on the curvature of the production function. So for both the labour share and capital share (and hence investment rate) to be increasing in the relative price of capital, as suggested by the econometric evidence in Karabarbounis and Neiman (2014) and this paper respectively, we require

$$\frac{d\omega_L}{dp} < 0$$
$$\frac{d\omega_K}{dp} < 0$$

The elasticities of output and the subfunction X with respect to capital are equal to the factor shares of capital in each, so the parameter condition for the labour share is

$$\left(\frac{1-\theta}{\theta}\right)\omega_K + \left(\frac{1}{\sigma} - \frac{1}{\theta}\right)\frac{\omega_K}{\omega_K + \omega_L} < 0$$

and for capital is

$$\left(\frac{1-\theta}{\theta}\right)\frac{dy}{dk} + \left(\frac{1}{\sigma} - \frac{1}{\theta}\right)\frac{dx}{dk} + \frac{\sigma - 1}{\sigma} < 0$$

Setting the labour and capital shares $\{\omega_L, \omega_K\}$ to reasonable values of 0.6 and 0.25 respectively (such that the share going to entrepreneurial labour M is 0.15), figure 3.23 displays the region of the parameter space in which both conditions are met. The elasticity of substitution between capital and unskilled labour σ must lie in a region around unity that is increasing in size the further below unity is the elasticity between entrepreneurial labour M and the other factors.

The intuition behind this result is as follows. A fall in the relative price of capital increases the quantity of it employed in production, for a given interest rate. For fixed supply of unskilled labour L and entrepreneurial labour M, an increase in K means an increase in the intermediate aggregate X and, for $\theta < 1$, a fall in the income share going to X. If the substitutability between K and L in X (i.e. σ) is low, then the rise in K pushes up the share of L in X quickly enough to offset the fall of the share of X in overall output, and the labour share rises. Conversely, if σ is high enough, then in response to the increase in K the capital share in X will rise quickly enough to allow it (and therefore the nominal investment rate) to rise as a share of Y. The further θ is below unity, the more that the overall share of X falls, and so the bigger is the region of σ in which these offsetting effects do not dominate.

Chapter 4

The banks that said no: banking relationships, credit supply and corporate outcomes in the United Kingdom

4.1 Introduction

The financial crisis of 2008 was associated with falls in corporate lending, business investment, labour productivity and real wages in the United Kingdom. What were the causal links between these events? Did firms retrench because they could not get financing? Or did they become pessimistic about demand for their products, and demand both less financing and fewer factors of production as a result? Empirical research in this area is only just beginning to address adequately the issue of causation.

This paper provides new evidence on the impact of the credit supply shock on corporate outcomes in the UK. Our aim is to identify what the impact of the reduction in credit supply following the 2007/8 financial crisis has been on wages, investment, employment and labour productivity. We employ a new identification strategy that exploits information on pre-crisis credit relationships within a large firm-level dataset of UK companies. We exploit the stickiness of these relationships, together with the fact that different banks tightened credit conditions to different degrees, to generate exogenous variation in credit supply at the firm level. This instrument for credit supply can be used to quantify the impact of a given change in credit volumes on firm outcomes.

Controlling for demand conditions in the product market and conditional upon firm survival, we find that firms facing a reduction in credit supply experienced greater falls in capital investment and labour productivity. Capital per worker fell significantly more in firms facing tighter financing conditions. This may be due to an increase in the shadow price of capital causing firms to substitute towards more labour-intensive technologies in production. Wages fell further in firms more exposed to the credit shock, and in similar proportion to labour productivity, even though these firms were hiring labour in the same markets as less exposed firms. This observation lends support to rent- or risk-sharing theories of wage determination (see e.g. Van Reenen (1996)) - firms were able to share some of their idiosyncratic productivity shocks with their workers. We also find that firms facing adverse credit supply shocks were more likely to fail

In the UK, firms are required to register the identity of any party (a 'chargeholder') that has a claim on the firm's assets as collateral for a loan. We construct a proxy for pre-crisis banking relationships by identifying UK banks among these chargeholders. We show that these relationships are persistent, and that they help to predict the amount firms borrow after the crisis.

Our identifying assumption is that banking relationships are only correlated with firm outcomes through their effect on credit supply, conditional on the control variables in our model. This puts our paper in a similar vein to Chodorow-Reich (2014), who looks at employment in US firms, Edgerton (2012) (equipment spending in California), Paravisini, Rappoport, Schnabl, and Wolfenzon (2011) (exports in Peru), Amiti and Weinstein (2013) (investment in Japan) and Bentolila, Jansen, Jimnez, and Ruano (2013) (employment in Spain). It will be violated if the banks which cut lending most during the crisis lent to firms which performed systematically worse, controlling for observables, whether this is due to reverse causation (bad firms harming their banks) or common causation (bad or risky decisions in several parts of the bank). We provide narrative evidence that the main cause of variability in banks' performance after the crisis was not their corporate lending decisions. Unlike a number of papers in this field, our identification strategy allows us to report standard tests of overidentifying restrictions, which are typically not rejected. We discuss how they are affected by the endogenous attrition in our sample. Our instruments are shown to be strong, and our parameter estimates are statistically significant and economically large, but only when we address the endogeneity of credit volumes with two-stage least squares - OLS estimates are typically much smaller in absolute value, and statistically less significant. We show that this is consistent with highly variable credit demand shocks at the firm level.

Relative to the existing literature, our study makes two principal contributions. First, our paper is the first to our knowledge to look at the effect of the credit shock on labour productivity and wages in the cross section, both of which have been puzzlingly weak in the United Kingdom and a number of other economies after the financial crisis. Second, our paper is the first to use bank relationships to study the credit shock in the United Kingdom, an economy which is both heavily dependent on banks and which suffered a relatively large credit shock.

The remainder of this paper is structured as follows. In section 4.2 we discuss how our work relates to existing studies of the impact of credit on corporate outcomes. Section 4.3 provides a brief overview of the behaviour of key macroeconomic variables in the UK since the 2008 crisis, and the structure of the UK banking system and corporate sector. Section 4.4 presents the dataset used in our analysis. Section 4.5 sets out our empirical methodology. Section 4.6 presents our results and compare them to existing estimates from the literature. Section 4.7 gauges the economic magnitude of our results. Lastly, we conclude in section 4.8 with an interpretation of our key results and suggestions for further work in this area.

4.2 Existing literature

The existing literature on the impact of credit on the macroeconomy can broadly be split into two strands.

The first strand uses aggregate macroeconomic data and time series econometric methods to identify the impact of a credit shock on the real economy. For example, Gilchrist and Zakrajsek (2012b) examine the relationship between credit spreads and economic activity in the time series. Reinhart and Rogoff (2012) find that output growth following financial crises remains persistently weak for longer than recessions that are not associated with financial crises. Oulton and Sebastia-Barriel (2013) use a similar dataset and find that financial crises are associated with prolonged periods of weak TFP growth.

The second strand, which relates more directly to our work, uses microeconometric techniques to identify the impact of a credit shock in the cross section. Several papers use pre-crisis variation in exposure to different lenders as a means of generating cross-sectional variation in credit supply during the crisis. (Peek and Rosengren (1997)) find that Japanese banks cut lending in the US following deterioration in their parent banks' capital positions, and these authors go on to show with regional data that this in turn affected US construction activity (Peek and Rosengren (2000)). Amiti and Weinstein (2013) use matched Japanese bank-firm data over the period 1990-2010 to decompose loan movements into bank, firm, industry and common shocks. They find that idiosyncratic bank shocks have a large effect on investment. Greenstone and Mas (2012) use geographic variation in the pre-crisis market share of different banks across the US, along with variation in the credit crunch across banks, and finds that US counties in which poorly performing banks had bigger market shares saw fewer new loans, less employment and fewer business start-ups during the crisis. Their empirical model assumes that each bank's loan growth in each county is the sum of county fixed effects, bank effects (varying over time for each bank) and an idiosyncratic shock. The predicted credit supply shock for each county is the average of the bank-time shocks, weighted by each bank's pre-crisis market share in that county. The measure is correlated with county-level loan and employment growth in the crisis, but not before.

Chodorow-Reich (2014) measures banking relationships by identifying the lead arrangers for a given firm's syndicated loans and showing that bank-firm pairs are sticky. He finds that employment fell more sharply during the crisis among the clients of less healthy lenders, particularly when those clients were small firms. The withdrawal of credit can explain roughly one-third of the employment decline in the sample in the year following the Lehman collapse. Edgerton (2012) uses data on lending relationships for a sample of equipment finance loans to identify the impact of restricting supply of credit to firms. Again he uses 'distressed' lenders as a proxy for restricted credit supply. He finds that variation across lenders account for a 17% decline in aggregate equipment financing, or about one-third of the total decline in financing in the sample of small businesses used in the paper. Paravisini et al. (2011) estimate the elasticity of exports to credit using matched Peruvian customs and firm-level bank credit data. To account for non-credit determinants of exports, they compare changes in exports of the same product and to the same destination by firms borrowing from banks differentially affected by the 2008 financial crisis. The results suggest that the credit shortage reduces exports through raising the variable cost of production, rather than the cost of financing sunk entry investments. Flannery, Giacomini, and Wang (2013) find that US firms which had relationships with banks with higher non-performing real estate loans borrowed less and invested less during the crisis.

Ongena, Peydro, and Horen (2013) examine how corporate outcomes of firms that are dependent on credit differed from those that are not credit constrained. They focus on firms located in Eastern Europe and Asia as that region was not initially affected by the global financial crisis. Their identification strategy relies on distinguishing between 3 types of banks according to whether they are domestic or foreign-owned and whether they can borrow on the international wholesale market. They find that banks with access to international wholesale funding cut back their lending by more than domestic banks who didn't; and that firms dependent on credit from those banks had lower returns on asset growth and revenue growth. Bentolila et al. (2013) merge the Spanish credit register with balance sheet data and find that Spanish firms who entered the crisis with relationships to weak banks experienced larger falls in employment.

In summary, there is a small but growing literature using bank relationships to study the effect of credit supply on corporate outcomes, principally borrowing, investment and employment. Our paper is the first to do so in the UK, and the first to look at capital intensity, labour productivity and wages.

4.3 Macroeconomic context

This section discusses the macroeconomic context for the research question, namely the weakness in productivity in the UK since the crisis (subsection 4.3.1) and the relationship between the UK banking and non-financial corporate sectors (subsection 4.3.2).

4.3.1 Productivity, investment and employment in the United Kingdom after the 2008 crisis

In common with many other industrial countries, the UK experienced a deep recession during the 2008 financial crisis. Since the crisis, labour productivity has been exceptionally weak in relation to its previous trend and past financial crises. It has also been weaker in the UK than in most other industrial countries (Barnett, Batten, Chiu, Franklin, and Sebastia-Barriel (2014)). This weakness in productivity has been puzzling, and is explained by surprisingly weak unemployment rather than weak output. It has coincided with a sharp fall in UK corporate borrowing (Figure 4.1) and real wages. A key question is the extent to which the disruption in credit supply witnessed during the 2008 crisis has been a cause of the weakness in capital formation, productivity and wages.

4.3.2 Banks and corporate borrowing in the United Kingdom

Firms in the United Kingdom are highly dependent on banks as a source of debt finance: the approximately 250 UK firms with access to public bond markets account for only about 12% of UK private employment. In their reliance on banks instead of bond markets, UK firms are much closer to continental European firms than US firms; bank loans account for about three-quarters of euro area corporate debt, about two-thirds in the UK, and about one-quarter in the US (Pattani, Vera, and Wackett (2011)). So variation in the credit supplied by banks is, for most UK firms, coterminous with variation in overall credit supply.

The banking sector is highly concentrated in the UK - the top six banks account for about 70% of the stock of lending to UK firms (Bank of England (2013)). The major UK banks - Barclays, HSBC, Lloyds Banking Group and

RBS (the Royal Bank of Scotland) - experienced widely differing fortunes during the financial crisis that began in 2008. In particular, Lloyds Bank and HBoS (Halifax Bank of Scotland) merged to form the Lloyds Banking Group and were subsequently part-nationalised by the UK government, as was RBS. Chart 4.2 shows the premia on credit default swaps on the senior unsecured bonds of the big four UK banks, a measure that is highly correlated with their funding costs. These were low and similar before the crisis, and high and dispersed afterwards.

A necessary condition for our identification strategy is that this dispersion was not in large part caused by systematic differences in the health of UK banks' corporate loan books. Official narrative accounts of the failures of HBoS and RBS (FSA (2011) and PCBS (2013)) support this idea, laying the blame instead on trading book losses and reliance on wholesale funding. The key exception to this is the particularly large losses made in RBS' and HBoS' commercial real estate (CRE) portfolios. For this reason, and in line with Bentolila et al. (2013), we exclude CRE firms from our sample.

4.4 Data description

Our dataset is compiled from information taken from the Bureau Van Dijk FAME database. This extracts information from UK companies annual accounts that are submitted to Companies House. This database contains information on around 1.3m registered UK companies. These data provide a range of balance sheet as well as profit and loss accounting information about individual companies.

4.4.1 Sample selection

A key variable of interest for our study is the level of total debt held by individual companies, which we define as the sum of overdrafts, short term loans and long term debt. The sample used in our analysis is based on all companies that reported a nonzero level of total debt in 2007 - giving us a total sample size of around 90,000 firms. We then track what happened to these companies in the two subsequent years. However, reporting restrictions and thresholds mean that only a subset of these companies will report data on all the variables of interest.

There are three main issues with the way we have selected our sample. First, the panel is unbalanced, since not all firms will submit data for every variable in each period (depending, inter alia, on whether they meet the appropriate reporting thresholds), particularly but not exclusively if they are liquidated. Second, since there are various accounting exemptions available for SMEs, they will be underrepresented in our sample relative to the population of firms. In subsection 4.4.3 we describe how the firm size distribution in our sample compares to the population. Third, we do not observe the effect of credit supply shocks on entry and firm creation, nor on firms which existed but had been denied credit in prior to the crisis.

4.4.2 Chargeholder information

Our identification strategy relies on information regarding pre-crisis relationships between companies and individual banks. To get this we extract information on registered charges from the BvD FAME database. A charge is the security a company gives for a loan and must be registered at Companies House (the UK business registry) within 21 days. There are two types of charge: a fixed charge is a charge or mortgage secured on particular property; and a floating charge can be against all the company's assets, such as stock in trade, plant and machinery and vehicles. The BvD FAME database captures information on persons entitled to an outstanding charge raised at Companies House. It also includes information on when the charge was created and when it ended (when the loan matured). The way in which we encode this information for our econometric analysis is set out in section 4.5.

Figure 4.6 shows the proportion of companies that started with an outstanding charge in 2007 with one of the four major UK banks and tracks whether or not they still had an outstanding charge to that particular institution in subsequent years. By the end of the sample period around 90% of companies still had an outstanding charge registered to the same institution as they did in 2007. Figure 4.6 also shows the proportion of companies that started off with an outstanding charge to one of the four major UK banks in 2007 and tracks whether or not a new charge was registered with a different institution. The chart shows that by 2011 only around 5% of companies had raised a charge with a different institution to the one they had an outstanding charge with in 2007. This suggests banking relationships - or at least the proxy in our dataset - appear to be very sticky.

4.4.3 Summary statistics

This subsection provides an overview of the key variables we use in our estimation strategy.

Summary statistics of key variables

Table 4.1 provides summary statistics for the left- and right-hand variables in our regressions, discussed in more detail in the next section. As is typically the case in firm-level data, the distributions of all our variables are heavily skewed to the left. For the analysis in the remainder of this paper, we remove firms in the 1st and 99th percentiles from each variable, to ensure that our results are not distorted by any outliers or extreme observations.

Size and industry distributions

Figure 4.3 shows the distribution of companies across three firm size bands: small, medium and large. Since information on the number of employees is less regularly reported than turnover, we use turnover our measure of firm size. Small firms, defined as companies with a turnover of less than 6.5 million, make up the vast majority of our sample. According to the 2012 Business Population Estimates (ONS (2013)), small firms (in those statistics defined as firms with fewer than 50 employees) represent around 99% of UK companies. As described earlier, the reason for this difference is that smaller firms are not required to submit full annual accounts and will in many cases, therefore, be excluded from our sample. Figure 4.3 shows a reasonable correspondence between the proportion of companies in different industries in our sample and in the population.

Bank market shares

Figure 4.5 below shows the proportion of companies with an outstanding charge registered at Companies House in 2007 within our sample. RBS and Lloyds had the largest market shares at 17% and 12% respectively, consistent with their dominance of the UK corporate lending market during that period.

4.5 Empirical approach and identification strategy

The availability of credit enables firms to invest, pay workers and buy intermediate goods before the revenues they create are generated. A disruption in credit supply in the form of a change in the terms on which a firm can access debt financing will therefore affect the ability of a firm to conduct its business. Our aim is to quantify the effect of a credit supply shock on various measures of firm $i's$ behaviour at time t - principally investment, productivity and turnover - represented by the variable y_{it}, conditional on other time-invariant covariates x_i. This is the parameter β in equation (4.1)

$$y_{it} - y_{i,2007} = \beta_t(d_{it} - d_{i,2007}) + \gamma x_i + \epsilon_{it} \tag{4.1}$$

where we measure a credit supply shock in terms of its effect on the volume of credit d_i the firm owes, holding demand fixed, and where x_i is a vector of time-invariant firm characteristics. One reasonable alternative to this specification would be to group banks according to whether they are strong or weak, and therefore more or less likely to provide credit, and then perform a diff-in-diff analysis comparing firms who have relationships with these two groups. The problem with this approach is that it is not obvious how to group banks. On one hand, banks like RBS and Lloyds became so weak that they were nationalised. On the other, nationalisation itself may have prompted a change in lending policy and actually boosted credit supply from the affected banks (see Rose and Wieladek (2014) for evidence that nationalisation affected the lending of UK banks).

At the level of a firm, the amount a firm borrows will be driven by both the supply of and demand for credit. For example, a firm might reduce borrowing because of a reduction in credit supply, but also because it might want to dispose of physical capital or otherwise alter its capital structure. In each case, the correlation between credit and investment will be different. So a simple OLS regression of, say, investment on the change in debt will typically deliver biased estimates of the effect of a credit supply shock.

To see the determinants of the OLS bias, consider a very simple model of borrowing (d), output (y) and credit supply (perfectly elastic at interest rate r) at the firm level

$$
\begin{aligned}
d =& \alpha_1 y + \alpha_2 r + \epsilon_d && \text{(Credit demand)} \\
r =& \epsilon_s && \text{(Credit supply)} \\
y =& \gamma_1 \epsilon_s + \epsilon_y && \text{(Output)}
\end{aligned}
$$

In terms of the exogenous shocks of the model, realised output and borrowing are given by

$$
\begin{aligned}
d =& \alpha_1 \left(\gamma_1 \epsilon_s + \epsilon_y\right) + \alpha_2 \epsilon_s + \epsilon_d \\
=& \left(\alpha_1 \gamma_1 + \alpha_2\right) \epsilon_s + \alpha_1 \epsilon_y + \epsilon_d \\
y =& \gamma_1 \epsilon_s + \epsilon_y
\end{aligned}
$$

If we were simply to regress output on credit volumes, the expected value of our OLS parameter estimate would be

$$
\begin{aligned}
E\left[\hat{\beta}_1^{OLS}\right] =& \frac{cov(y,d)}{var(d)} \\
=& \gamma_1 \frac{\sigma_s^2 \left(\alpha_1 \gamma_1 + \alpha_2\right)}{\left(\alpha_1 \gamma_1 + \alpha_2\right)^2 \sigma_s^2 + \alpha_1^2 \sigma_y^2 + \sigma_d^2} + \frac{\alpha_1 \sigma_y^2}{\left(\alpha_1 \gamma_1 + \alpha_2\right)^2 \sigma_s^2 + \alpha_1^2 \sigma_y^2 + \sigma_d^2}
\end{aligned}
$$

Now suppose we have an instrument b which is correlated with credit supply but uncorrelated with any of the other shocks, such that $b = \mu \epsilon_s + \epsilon_b$. In expectation, our IV estimator is then

$$
\begin{aligned}
E\left[\hat{\beta}_1^{IV}\right] =& \frac{cov(y,b)}{cov(d,b)} \\
=& \frac{E\left[\left(\gamma_1 \epsilon_s + \epsilon_y\right)\left(\mu \epsilon_s + \epsilon_b\right)\right]}{E\left[\left(\left(\alpha_1 \gamma_1 + \alpha_2\right) \epsilon_s + \alpha_1 \epsilon_y + \epsilon_d\right)\left(\mu \epsilon_s + \epsilon_b\right)\right]} \\
=& \frac{\gamma_1 \mu \sigma_s^2}{\mu \left(\alpha_1 \gamma_1 + \alpha_2\right) \sigma_s^2} \\
=& \frac{\gamma_1}{\left(\alpha_1 \gamma_1 + \alpha_2\right)}
\end{aligned}
$$

This is the effect of credit supply shocks on output normalised by their effect on borrowing. Relative to this parameter, the OLS estimator is biased for two reasons. First, the bias is an increasing function of $\alpha_1\sigma_y^2$ - i.e. output shocks will bias the OLS parameter upwards to the extent that borrowing is an increasing function of output and that there are output shocks in the sample. Intuitively, if credit volumes are strongly increasing in output, and output varies autonomously a great deal, the parameter estimate in an OLS regression of output on credit volumes will be revised upwards. Secondly, credit demand shocks will bias the parameter towards zero, as they will raise the variance of the right-hand side variable in the regression.

For this reason, we adopt an instrumental variables approach. For each of the firms in our sample, we have information about the identity of any legal person with a charge on the assets of the firm. When the chargeholder is a bank, we take this as evidence of a possible banking relationship between the firm and the bank. These relationships are in turn an indicator of firm-specific credit supply on account of two features of the UK banking system in the recent crisis: the stickiness of banking relationships (subsection 4.4.2), and the exogenous differences across banks in the severity with which they were hit by the credit shock (section 4.3.2).

Our dataset is an unbalanced panel of firms, so we could in principle estimate a variant of equations 4.2-4.3 below in levels terms rather than first differences, controlling for time-invariant unobserved heterogeneity using standard methods. This approach would, however, suffer from a number of important problems. Most obviously, the credit shock itself varied over time and affected each bank differently over our sample. Furthermore, our identifying strategy relies on using pre-crisis banking relationships as an instrument for firm-level credit supply. Over the passage of time, for a variety of reasons, banking relationships will change and end. This means that the coefficients of equation 4.2 are highly likely to be unstable over time, and we verify that this is the case in the next section.

We therefore estimate our model in terms of changes between the year 2007, before the most serious phase of the credit crisis, and each of the post-crisis years 2008-2009 in our sample. Beyond 2009, our identification strategy begins to fail as pre-crisis banking relationships have longer to decay, and do so non-randomly, so we do not present results for later years.

Our model - estimated by two-stage least squares - is thus

$$(d_{it} - d_{i,2007}) = \theta b_i + \xi x_i + \mu_i \tag{4.2}$$

$$y_{it} - y_{i,2007} = \beta(\widehat{d_{it} - d_{i,2007}}) + \gamma x_i + \epsilon_{it} \tag{4.3}$$

where b_i is an indicator of the identity of the bank with which firm i had relationships before the crisis, explained in more detail above. This is similar to a time-differenced version of the specification in Paravisini et al. (2011).

We encode the bank relationships as a set of seven binary dummy variables. Each of four dummies represents a firm having a relationship with exactly one of the big four banks - Barclays, HSBC, Lloyds (including HBoS) and RBS (including NatWest) - given the relative infrequency with which other banks appear in our sample. The remaining three respectively code for more than one relationship with one of the big four, a relationship with a bank outside the big four, and a relationship with a non-bank. The vector of control variables consists of a constant, dummy variables for the industrial sector of the firm, and dummy variables for whether the firm is a subsidiary, a parent company or a standalone firm. Our IV results are qualitatively robust to reasonable alternatives to this scheme. For example, we get similar results (not reported here) when we code a relationship with each of the big banks in a separate dummy variable, irrespective of whether the firm has relationships with other banks.

Our identifying assumption is that the identity of a firm's bank is only correlated with its performance, conditional on observables, through the effect that bank identity has on credit supply. In terms of equations 4.2-4.3, validity requires that

$$E[\epsilon_i b_i] = 0. \tag{4.4}$$

Put another way, banking relationships are only valid instruments if, as well as being correlated with credit supply, they are also uncorrelated with any other determinants of our dependent variable we do not control for.

There are two principal reasons why this might not be the case. First, there could be reverse causation, i.e. from firm performance to credit supply rather than vice versa. If the firms who had relationships with a given bank performed systematically worse (conditional on observables) than others, say because that bank had selected riskier or less promising borrowers than others, causation could run from corporate performance to bank relationships - violating our identifying assumption.

Second, there could be selection on unobservables, whereby a firm's performance and the lending behaviour of its bank are influenced by a common, unobserved factor. For example, suppose a given bank took above-average risks on both sides of its balance sheet in the lead-up to the crisis. In the event of a system-wide financial shock, its lending would have contracted more than average on account of funding difficulties, and its borrowers may have performed less well than average because they were more exposed to the economic cycle. We would then observe a conditional correlation between performance and borrowing at the firm level, but it would not be causal - another violation of our identifying assumption.

We can provide two kinds of evidence in support for our identifying assumption. First, we report in section 4.6 standard tests of overidentifying restrictions, which are typically not rejected at standard significance levels. Where they are rejected, this may be due to the role of credit supply in firm attrition, as set out in subsection 4.5.1 below. Second, as referred to in section 4.3, there is narrative evidence that the main cause of variability in banks' performance after the crisis was not their corporate lending decisions.

4.5.1 Credit supply and firm survival

When firms cease to operate, the left-hand side variable is typically not recorded or recorded at zero (such that the log difference is undefined), so the firm in question is dropped from the sample. So the preceding analysis can necessarily only be done on the intensive margin - i.e. on those firms which do not cease to function for reasons of credit supply or otherwise. An interesting question is whether this sample attrition is random or not with respect to the other variables in our model. It seems likely that changes in credit supply will influence firm survival. But with our dataset we can also quantify the impact of credit supply shocks on firm survival. We attempt to quantify this below.

Our identification strategy involves predicting the cumulative change in borrowing between 2007 and another year on the basis of *time-invariant* variables observed in 2007. This means that we can evaluate the predicted change in a firm's borrowing among dead firms, i.e. how much a firm with similar pre-crisis characteristics would have been expected to borrow had it survived. This is a natural metric with which to assess the impact of predicted borrowing, and in particular the contribution of bank identity to it, as a determinant of firm survival. We construct a binary cumulative failure indicator f_{it} taking the value of zero if firm i is alive in year t and 1 if the firm failed in or before year t.[1] We then run logit regressions of firm failure or survival on the predicted value of credit supply, plus the non-bank controls in equation (4.2)

$$f_{it}^{*} = \alpha_t \widehat{\Delta d_i} + \beta_t x_i + u_{it} \tag{4.5}$$

$$Pr(f_{it}^{*} > 0) = F(\alpha \widehat{\Delta d_i} + \beta x_i) \tag{4.6}$$

The estimated coefficient on credit supply will therefore capture the effect that credit supply, as identified with bank ID and measured in units of credit volumes, has on bank survival. A negative coefficient would mean that firms which would have been able to borrow more, had they survived, would have been less likely to fail.

To the extent that credit supply affects firm survival, it may also affect the results of the standard tests of overidentification restrictions used to assess instrument validity. Standard tests of overidentifying restrictions work by looking how far from zero our sample analogues $b'\hat{\epsilon}$ of the population moment conditions $E\left[\epsilon_i b_i\right] = 0$ are. However, if there is nonrandom attrition in our sample, tests based on such overidentifying restrictions are unlikely to work. In particular, if the disturbances in the observation equation ϵ are correlated with those in the selection equation u, and our instrument d is a determinant of the latent selection variable f^{*} (this is a necessary condition for identification), then in general our instrument will be correlated with the disturbances in the observation equation, conditional on them being observed, i.e.

$$E\left[b'\epsilon | f^{*} > 0\right] \neq 0$$

To see this, assume that the moment condition $E[b'\epsilon] = 0$ holds in the

[1] We create a proxy for firm death by looking at whether a firm's status (when the data were collected) is not 'active' and looking for the first year in which balance-sheet data such as assets are either zero or missing: we assume the firm failed in that year.

population. However, what we observe is instead the sample analogue of $E\left[b'\epsilon|f^*>0\right]$

$$\begin{aligned}
E\left[b'\epsilon|f^*>0\right] =& E_b\left[E\left[b'\epsilon|f^*>0,b\right]\right] \\
=& E_b\left[b'E\left[\epsilon|f^*>0,b\right]\right] \\
=& E_b\left[b'E\left[\epsilon|\alpha_t\widehat{\Delta d_i}+\beta_t x_i+u_{it}>0,b\right]\right] \\
=& E_b\left[b'E\left[\epsilon|u_{it}>-\alpha_t\widehat{\Delta d_i}-\beta_t x_i,b\right]\right]
\end{aligned}$$

Suppose for the sake of illustration that u and ϵ are both mean zero, and jointly normally distributed with covariance $\sigma_{u\epsilon}$. Then

$$\begin{aligned}
E_b\left[b'E\left[\epsilon|u_{it}>-\alpha_t\widehat{\Delta d_i}-\beta_t x_i,b\right]\right] =& E_b\left[b'\frac{\sigma_{u\epsilon}}{\sigma_\epsilon}\frac{\phi\left(\frac{-\alpha_t\widehat{\Delta d_i}-\beta_t x_i}{\sigma_\epsilon}\right)}{1-\Phi\left(\frac{-\alpha_t\widehat{\Delta d_i}-\beta_t x_i}{\sigma_\epsilon}\right)}\right] \\
\neq& 0
\end{aligned}$$

So in general, even if our identifying restrictions hold, our instruments will be correlated with the residuals in the second-stage equations for the continuous variables. The intuition is that our instruments determine selection, so they will be systematically related to the unobserved variables u in the selection equation *among surviving firms.* The latter will be related systematically to the unobserved variables ϵ to the extent that firms which are more likely to survive are also more likely to invest, hire labour, and so forth. This means that our instruments can be correlated with the residuals of the second-stage equation, invalidating tests of validity based on overidentifying restrictions. We may therefore reject the null of validity more frequently than indicated by the significance level of the test.

4.6 Results

This section of the paper present our results. Subsection 4.6.1 sets out our baseline results. Subsection 4.6.2 presents tests of the robustness of these results. And subsection 4.6.3 provides results disaggregated along various dimensions.

4.6.1 Baseline results

First stage regression

Table 4.3 sets out the results of the first stage regression (equation 4.2) of the change in credit volumes on our vector of identifiers for bank relationships b_i and (not shown) our additional observable controls x_i, a vector of dummies of the industrial sector of the firm and whether it is a subsidiary of other firms. The four columns show the results in four different cases, depending on whether we are measuring the change from 2007 to 2008 or 2009, and whether we are using all available observations in a given year, or restricting ourselves to firms for which data is available in both years.

The last two rows of the table show that the regression as a whole and the bank dummies are highly significant, with the p-values of the null hypotheses that all variables or the bank dummies can be omitted very close to zero, even though the R^2 of the regression is small. The dummies themselves are precisely estimated and in many cases bilaterally significantly different from each other.

Second stage regression

Tables 4.4 - 4.8 set out the results of the second stage regressions (equation 4.3). The eight columns in each table show all possible permutations of the time period in question (whether we are measuring changes from 2007 to 2008 or 2009), whether the sample is balanced or unbalanced (i.e. whether we restrict attention to observations available in both years), and whether we use OLS or 2SLS as our estimator. The rows of the table show the estimated parameter β and its (robustly) estimated standard error in square brackets, the number of observations used, the R^2 of the regression, the Sargan statistic of overidentifying restrictions and its p-value, and the F-statistic of the signficance of the first stage of the regression, again with its p-value.

Table 4.4 shows the results with turnover per head - our measure of labour productivity - as the dependent variable. In both years, and for both samples, the OLS estimates of β are economically very small and statistically insignificant. In contrast, our IV estimates are economically large and relatively precisely estimated. Based on the argumentation in section

4.5, the apparent bias towards zero in the OLS numbers is evidence that credit demand shocks are an important source of variation at the firm level, relative to shocks to credit supply and output.

Interpreting the numbers, column (2) shows that a credit shock that reduces a firm's level of borrowing by 1% in 2008 reduced labour productivity by around 0.8%. The diagnostic tests reported with the regression suggest that any bias in this estimate is unlikely to be very large due to any problems with instrument validity and/or relevance: the Sargan test of overidentifying restrictions (where the null is that our instruments are valid) is not rejected at standard levels for either year, while the F-statistic for the joint significance of the instruments suggests that chance of a large weak instrument bias is low for 2008 (Stock and Yogo (2002)).

Comparing the unbalanced sample estimates for 2008 and 2009, we see that the estimated parameter falls approximately 2.5 standard deviations from 0.78 to 0.45. There are at least three reasons, two econometric and one economic, why this might be the case. First, the instruments are weaker for 2009, which may be one reason why the estimated coefficient is smaller in absolute value (i.e. it may be biased towards the OLS parameter). Second, the sample shrinks between the two years, in part because some of the firms in the 2008 sample had gone bankrupt by 2009. In the (likely) event that some of this attrition is non-random with respect to credit supply, the effect we observe among surviving firms will be different; it is noteworthy in this regard that the difference between the estimates for 2008 and 2009 is smaller on the balanced sample. Third, the effect of a persistent credit shock on productivity may fade over time, as firms have more time to reorganise production or finance expenditures with internal funds.

Table 4.5 shows analogous results for capital per head.[2]. Once again the IV estimates of β are reasonably well-determined and much larger in absolute value than their OLS analogues although, in contrast to the productivity regressions, the latter are now statistically significant. The 2SLS coefficients now get somewhat bigger over time, perhaps as firms have more time to adjust their factors of production.

Tables 4.6 and 4.7 separately show how credit shocks affect the numerator and denominator of capital per worker. As with previous results, we

[2]To proxy productive capital we use tangible assets

find that our IV estimates are generally well determined and larger in absolute value than the OLS numbers. The results show that the adjustment in capital per worker was driven more by capital in 2008 than 2009. The results on employment are surprising at first sight - they show no effect on employment in 2008 but by 2009 a credit shock that reduced borrowing by 1% by 2009 relative to 2007 was associated with an approximately 0.5% *rise* in employment among surviving firms.

One interpretation of this result is that the credit shock raised the shadow interest rate. The up-front costs and duration of returns to capital are longer than those to labour, such that a rise will cause firms to substitute towards labour and away from capital. It is important to emphasise that this result does not predict that aggregate employment will rise in response to an adverse credit shock, for at least two reasons. First, our method implicitly partials out the demand conditions that firms face by comparing employment growth across firms operating in the same markets but with different banking relationships. Second, we only quantify these effects among surviving firms: to the extent that firm-level credit supply shocks cause firms to go bust (and hence to cut employment to zero), their overall effect is of ambiguous sign. An estimate of the effects of firm-level credit shocks on firm survival is presented below.

Finally among our benchmark results, table 4.8 shows our estimate of the effect on wages (per head) of a firm-level credit shock that reduces borrowing. As above, the IV estimates are much larger in absolute value than their OLS analogues, and suggest that a firm-level credit shock reduces wages in approximately the same proportion as it reduces debt. This suggests that firms were able to share some of the costs by the credit shock with their employees, and relatedly that these firms were operating in a labour market that was non-Walrasian to an important extent, such that wages were not equalised across similar employees but differed ex post according to the credit supply experienced by their employers. Comparing them to table 4.4, it seems that firms in our sample were able to cut wages per head by somewhat more than the fall in productivity per head.

Credit supply and firm survival

Table 4.9 sets out our estimates of the parameter α in equation 4.5, the effect of a reduction in credit supply on the probability of firm death. The results are negative and significant for all years under consideration. A negative parameter means that higher predicted credit supply (i.e. the amount a firm's borrowing would have changed conditional on survival based on its banking relationships and other observables) increases the probability that the firm survived up until the period in question. This suggests that a widespread contraction in credit supply will tend to increase corporate insolvencies across the economy, an effect we quantify in section 4.7

4.6.2 Robustness tests

Our identifying strategy works by exploiting variation across banks in the changes in the terms on which they supplied credit during the financial crisis. This variation was likely to have been relatively large during the period of funding and credit market turbulence, asset-price volatility and bank nationalisation of 2008-9. Conversely, if our instruments are valid and banking relationships are randomly assigned with respect to corporate outcomes, our instruments are more likely to be only weakly relevant at a time of tranquil market conditions. But if there is also endogenous variation in bank relationships that is relevant for credit supply and corporate performance, violating our identifying assumptions, then our instruments could turn out to be relevant when they should not be.

With this in mind, we re-run the regressions above but substitute 2005 in place of 2007 for our sample selection rule[3] and as the base year against which changes in debt, productivity and so on are measured. Tables 4.10-4.14 present the results and are analogous to tables 4.4-4.8 discussed above. In contrast to our baseline results, the F-statistics for the first-stage regressions and the associated IV estimates of β are typically insignificant at standard levels. The only exception to this is when capital is the dependent variable and the change is measured between 2005 and 2007, in which case the F-statistic indicates that our instruments are significant. This could be

[3]that is, we select firms in the Bureau Van Dijk database who had a positive level of debt in 2005 rather than 2005

a false positive, or could suggest that the market turbulence experienced in the second half of 2007 - which culminated in the UK with the nationalisation of Northern Rock - may already have been having an effect on the real economy. However, in no case are our instruments relevant in 2006.

In summary, the failure of our identification strategy when used to identify credit shocks before the credit crisis (when the variation across banks in credit supply was more limited) provides greater confidence that we are indeed picking up credit supply shocks - and nothing else - in our baseline results for 2008-9.

4.6.3 Disaggregated results

The following subsection disaggregates the firms in our sample by size (subsubsection 4.6.3) and then by age (subsubsection 4.6.3). For brevity, we do not report OLS results. We disaggregated the sample along various other margins but did not obtain meaningful results, and therefore do not report these exercises.

Disaggregation by size

Tables 4.15 to 4.19 set out estimates of the parameter β in equation 4.3 where the sample has been split into firms with turnover above and below 50m (the standard criterion for 'small business' in UK statistics). This corresponds to around the 90th percentile of the turnover distribution in our dataset. For brevity, we drop the OLS results from the table. The F statistics from the first stage regression are typically highly significant but in single figures, suggesting that weak instruments bias may be more of a problem in our disaggregated samples. Table 4.15 does not present a consistent picture of the relative importance of credit shocks for the productivity of small and large firms. Table 4.16 suggests that capital per head adjusted proportionally more in large firms, with coefficients in the range of 0.7-1 rather than 0.5, although this evidence is fairly weak. Table 4.17 shows that capital adjusted much more among larger firms, although our identification seems to be weak among the smaller firms, such that this result may be a statistical artefact from weak identification among the latter. Table 4.18 shows the same for employment - an apparently larger (negative) response among large firms, and weak identification among small firms. Finally, table 4.19 shows that the coefficient for wages per head is apparently similar between small and large firms, with some identification questions among small firms in 2009. To summarise, there is some weak evidence that large firms adjusted factors of production proportionally more than small firms in response to the credit shock, but none that productivity or wages behaved differently.

Disaggregation by age

Tables 4.20 to 4.24 repeat the exercise in the previous subsubsection but disaggregate firms by (above and below median) age rather than by size. And as with that exercise, the F statistics, while highly significant, again suggest that the issue of weak instruments bias may affect the results somewhat. With this caveat in mind, however, table 4.21 suggests that older firms adjusted capital per head more than younger firms, with more adjustment in both the numerator and the denominator. This may be because they are themselves larger.

4.7 Economic magnitude

It is not straightforward to translate our microeconometric findings into macroeconomic numbers. But the simple calculations in this section show that the impact of the credit shock on the macroeconomy may have been substantial.

In a logistic regression such as that defined by equations 4.5-4.6, a slope parameter α implies that a one-unit increase in the explanatory variable increases the odds of the event by a factor e^{α}. So our estimate of the parameter α in the survival equation 4.6 of around -7.5 suggest that a credit supply shock that would reduce borrowing by 10% in the event the firm survives would raise the odds of bankruptcy by about $e^{.75} \approx 50\%$. Corporate liquidations in England and Wales, a close proxy in the aggregate data (ref to QB) rose by around two-thirds in the course of the crisis that began in 2008, while the stock of corporate debt fell by about one-third relative to its pre-crisis trend path (figure 4.1). For the sake of illustration, if even one-third of this fall in corporate debt represented a credit supply shock, then the shock could in turn account for most of the rise in insolvencies at this time.

Turning to the impact on surviving firms, our baseline estimates suggest that a contraction in credit supply sufficient to reduce borrowing by 10% over 2008 would have reduced productivity by about 7% among our sample of firms. Productivity (GDP per head) fell by approximately 3% in the UK over the course of 2008, and had fallen by about one-sixth relative to trend by 2013 (Barnett et al. (2014)). Even after allowance is made for the fact that firms without debt may have been less affected, it is still possible to account for a substantial part of the weakness in UK productivity on the basis of these estimates.

4.8 Conclusion

This paper provides new evidence on the impact of the credit supply shock caused by the financial crisis of 2008 on corporate outcomes in the UK. Controlling for demand conditions in the product market and conditional upon firm survival, we find that firms facing a reduction in credit supply experienced greater falls in capital investment and labour productivity. Capital per worker fell significantly more in firms facing tighter financing conditions. This may be due to an increase in the shadow price of capital causing firms to substitute towards more labour-intensive technologies in production. Wages fell further in firms more exposed to the credit shock, and in similar proportion to labour productivity, even though these firms were hiring labour in the same markets as less exposed firms. We also find that firms facing adverse credit supply shocks were more likely to fail.

Our identifying assumption is that banking relationships are only correlated with firm outcomes through their effect on credit supply, conditional on the control variables in our model. We provide narrative evidence that the main cause of variability in banks' performance after the crisis was not their corporate lending decisions, and report standard tests of overidentifying restrictions, which are typically not rejected. Our instruments are shown to be strong, and our parameter estimates are statistically significant and economically large, but only when we address the endogeneity of credit volumes with two-stage least squares - OLS estimates are typically much smaller in absolute value, and statistically less significant. We show that this is consistent with highly variable credit demand shocks at the firm level.

The limitations of our identification strategy, which relies on pre-crisis banking relationships that decay non-randomly over time, mean that we are unable to say how persistent the corporate sector was damaged beyond the first two years after the financial crisis. On the intensive margin, it could be that firms are somehow permanently scarred by temporary credit shocks, or it could be that they are able to catch up to the counterfactual no-crisis path once the credit shock abates. And on the extensive margin, this study does not tell us how firm entry is affected by financial crises, so we cannot say what happens to the factors that become unemployed when firms fail, either during or after the period of crisis. The durability of the productivity slowdown following the crisis make this an important avenue for future research.

Figure 4.1: UK macroeconomic data, log-linearly detrended

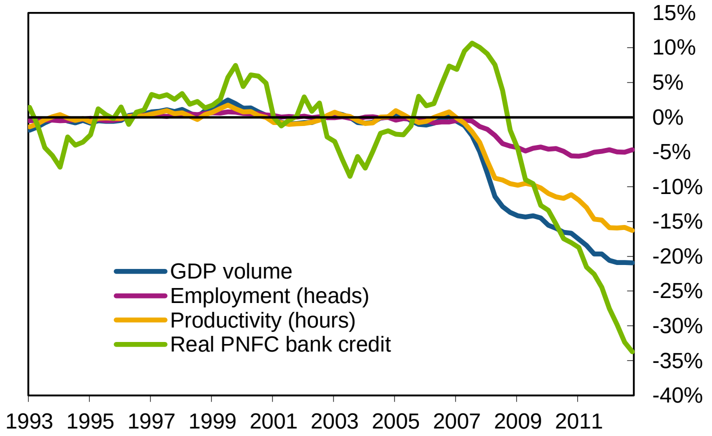

Figure 4.2: UK bank CDS spreads

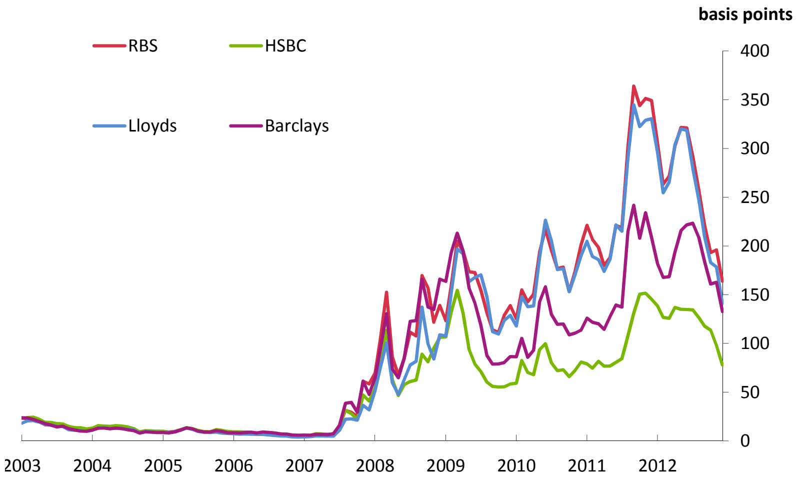

Figure 4.3: Distribution of firm sizes

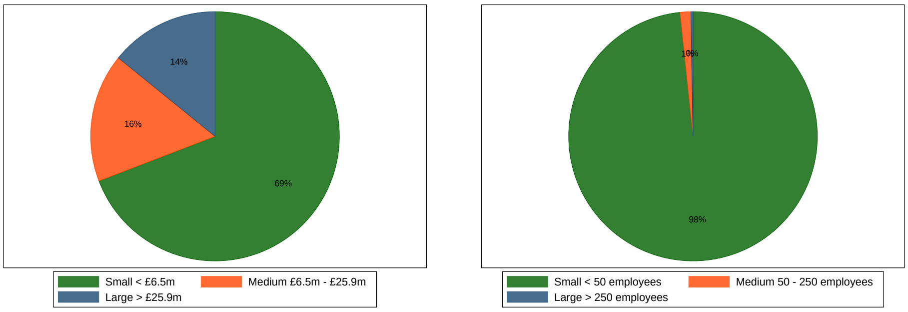

(a) Sample (b) BPE(2013)

Figure 4.4: Distribution of sectors

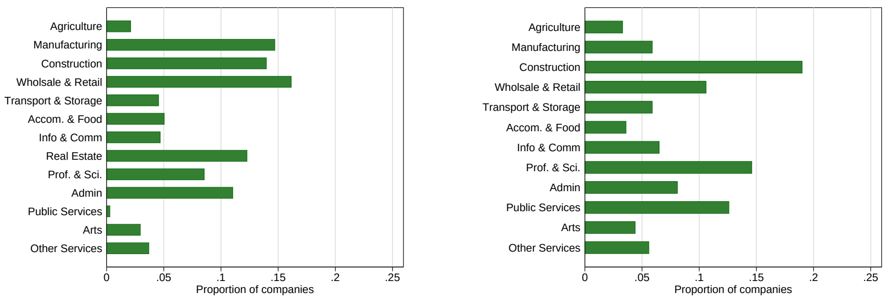

(a) Sample (b) BPE(2013)

Figure 4.5: Market shares of the major banks

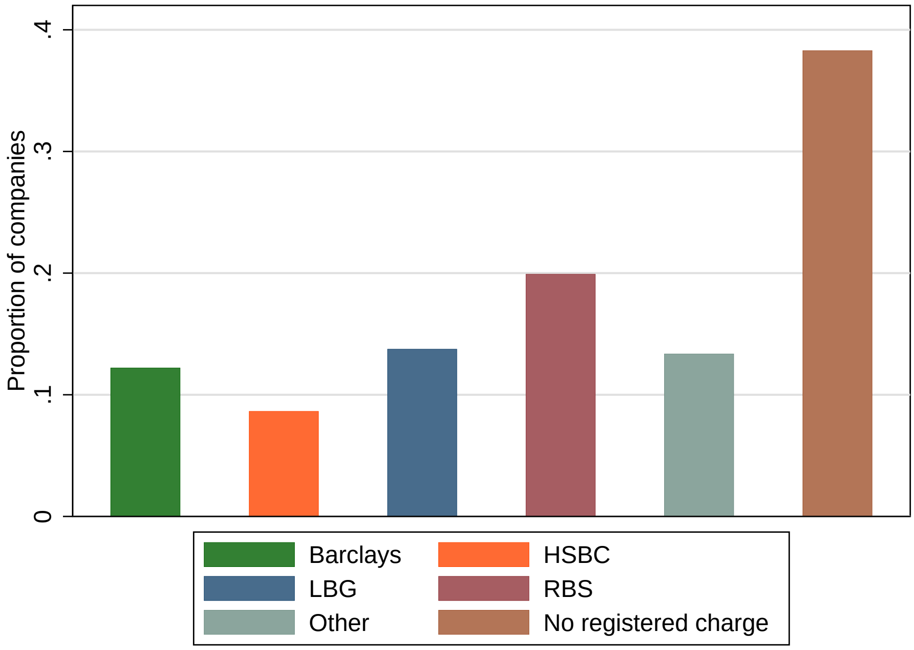

Figure 4.6: Charge durations and switching behaviour

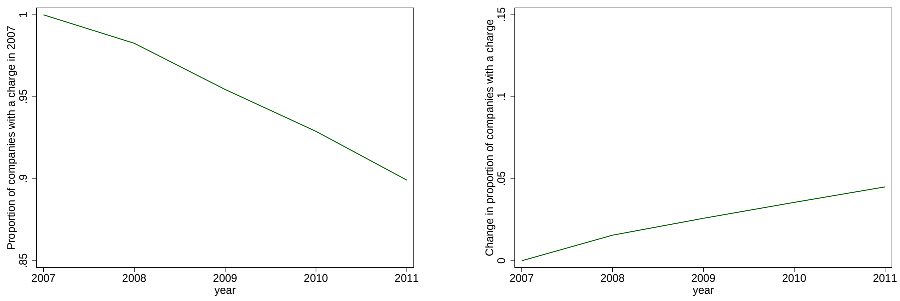

(a) Survival rate

(b) Switching rate

Figure 4.7: Corporate liquidations in England and Wales

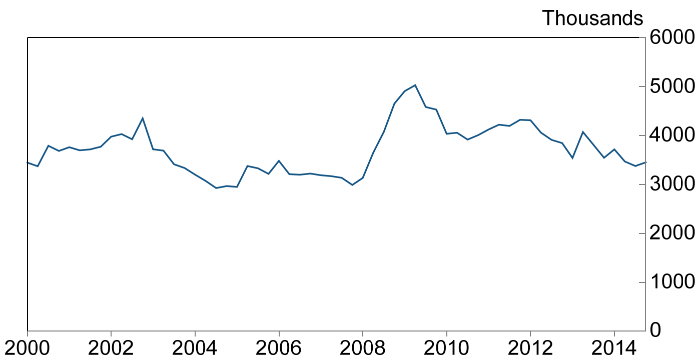

Table 4.1: Summary statistics, 2005 to 2011

Variable	Mean [St. Dev]	Median	Min	Max	Obs
Total debt (£'000s)	43,709 [1,500,246]	563	1	753,000,000	354,332
Capital (£'000s)	15,629 [351,876]	290	1	43,100,000	460,195
Capital per worker (£'000s)	259 [4,305]	18	0	453,333	153,726
Turnover per head (£'000s)	340 [2,655]	127	0	291,235	117,173
Remuneration per head (£'000s)	34 [32]	28	0	3,410	154,727
Number of employees	637 [7,023]	76	1	639,964	160,354

Table 4.2: Descriptive statistics of dependent variables in 2007

	Bank 1	Bank 2	Bank 3	Bank 4
Tangible assets ('000)	23,869 [302,440]	22,188 [184,333]	18,529 [114,892]	28,947 [483,617]
Employment	351 [1,428]	302 [1,438]	490 [2,542]	413 [6,049]
Turnover ('000)	45,703 [145,564]	35,213 [114,438]	68,354 [322,248]	56,310 [833,066]
Turnover per head	379 [4,954]	240 [737]	237 [810]	239 [582]

Table 4.3: First stage regression, baseline sample

	(1)	(2)	(3)	(4)
VARIABLES	unbalanced 2008	unbalanced 2009	balanced 2008	balanced 2009
justbarc_07	-0.078***	-0.074***	-0.113***	-0.074***
	[0.01]	[0.01]	[0.01]	[0.01]
justhsbc_07	-0.071***	-0.089***	-0.107***	-0.089***
	[0.01]	[0.01]	[0.01]	[0.01]
justlbg_07	-0.074***	-0.076***	-0.102***	-0.078***
	[0.01]	[0.01]	[0.01]	[0.01]
justrbsnw_07	-0.059***	-0.053***	-0.084***	-0.054***
	[0.01]	[0.01]	[0.01]	[0.01]
any2or3or4_07	-0.047***	-0.035***	-0.073***	-0.029**
	[0.01]	[0.01]	[0.01]	[0.01]
othermfi07	-0.053***	-0.045***	-0.078***	-0.041***
	[0.01]	[0.01]	[0.01]	[0.01]
chargenonmfi07	-0.052***	-0.057***	-0.083***	-0.058***
	[0.01]	[0.01]	[0.01]	[0.01]
Observations	54,432	41,537	38,365	38,365
R-squared	0.006	0.006	0.011	0.006
F statistic	12.61	9.449	16.71	9.692
p value of regression	0	0	0	0
p value of all banks	0	0	0	0

Table 4.4: Turnover per head, baseline sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt07	0.004	0.779***	0.007	0.454***	0.002	0.684***	0.008	0.571***
	[0.01]	[0.13]	[0.01]	[0.16]	[0.01]	[0.12]	[0.01]	[0.18]
Observations	14,952	14,952	12,501	12,501	11,435	11,435	11,435	11,435
R-squared	0.011	-1.698	0.012	-0.670	0.014	-1.414	0.011	-1.080
Sargan		5.971		10.87		7.806		6.709
Sargan pval		0.426		0.0925		0.253		0.349
F statistic		8.156		2.848		9.058		2.911
F pval		6.00e-10		0.00574		0		0.00484

Table 4.5: Firm capital per head, baseline sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt07	0.134***	0.739***	0.124***	1.050***	0.135***	0.732***	0.128***	0.958***
	[0.01]	[0.09]	[0.01]	[0.23]	[0.01]	[0.08]	[0.01]	[0.20]
Observations	20,345	20,345	16,597	16,597	15,518	15,518	15,518	15,518
R-squared	0.035	-0.595	0.032	-1.243	0.038	-0.597	0.034	-0.994
Sargan		14.16		8.090		7.626		7.264
Sargan pval		0.0279		0.232		0.267		0.297
F statistic		17.10		4.198		19.49		4.826
F pval		0		0.000124		0		1.92e-05

Table 4.6: Total assets, baseline sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt07	0.206***	0.740***	0.222***	0.257***	0.228***	0.738***	0.226***	0.194**
	[0.00]	[0.07]	[0.01]	[0.09]	[0.01]	[0.05]	[0.01]	[0.09]
Observations	48,077	48,077	35,675	35,675	32,953	32,953	32,953	32,953
R-squared	0.085	-0.427	0.079	0.077	0.099	-0.366	0.082	0.081
Sargan		14.21		13.23		13.40		14.94
Sargan pval		0.0274		0.0395		0.0371		0.0208
F statistic		21.85		10.68		41.10		10.87
F pval		0		0		0		0

Table 4.7: Total employment in heads, baseline sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt07	0.043***	-0.055	0.067***	-0.513***	0.046***	-0.081*	0.066***	-0.506***
	[0.00]	[0.04]	[0.00]	[0.15]	[0.00]	[0.04]	[0.00]	[0.14]
Observations	21,006	21,006	17,110	17,110	16,017	16,017	16,017	16,017
R-squared	0.017	-0.033	0.032	-1.258	0.018	-0.067	0.031	-1.213
Sargan		6.630		10.77		5.693		10.10
Sargan pval		0.356		0.0956		0.458		0.120
F statistic		17.23		3.947		17.94		4.521
F pval		0		0.000258		0		4.77e-05

Table 4.8: Labour remuneration per head, baseline sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt07	0.012***	0.939***	0.004	1.060***	0.011***	0.909***	0.004	1.030***
	[0.00]	[0.10]	[0.00]	[0.21]	[0.00]	[0.09]	[0.00]	[0.20]
Observations	20,160	20,160	16,635	16,635	15,536	15,536	15,536	15,536
R-squared	0.005	-4.600	0.012	-7.289	0.006	-4.357	0.012	-7.026
Sargan		3.123		10.06		4.880		11.26
Sargan pval		0.793		0.122		0.559		0.0808
F statistic		16.41		3.971		18.32		4.196
F pval		0		0.000241		0		0.000125

Table 4.9: Firm survival logit regression, baseline sample

		(1)	(2)	(3)	(4)
EQUATION	VARIABLES	2008	2009	2010	2011
failedever	predicted_debt	-7.360***	-7.200***	-8.094***	-8.005***
		[0.93]	[0.69]	[0.69]	[0.93]
	Observations	82,624	82,709	82,725	82,725
	pseudo R-squared	0.239	0.236	0.217	0.196

Table 4.10: Turnover per head, 2005 sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2006	2006	2007	2007	2006	2006	2007	2007
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt05	-0.012**	1.075*	0.004	0.257	-0.013**	0.439	0.002	0.235
	[0.00]	[0.65]	[0.00]	[0.23]	[0.01]	[0.54]	[0.00]	[0.23]
Observations	15,388	15,388	12,302	12,302	11,425	11,425	11,425	11,425
R-squared	0.003	-3.958	0.004	-0.269	0.004	-0.706	0.003	-0.230
Sargan		1.270		8.179		5.030		7.356
Sargan pval		0.973		0.225		0.540		0.289
F statistic		0.512		1.163		0.246		1.213
F pval		0.826		0.320		0.973		0.292

Table 4.11: Capital per head, 2005 sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2006	2006	2007	2007	2006	2006	2007	2007
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt05	0.127***	-0.164	0.145***	-0.489	0.135***	0.066	0.152***	-0.444
	[0.01]	[0.51]	[0.01]	[0.39]	[0.01]	[0.38]	[0.01]	[0.39]
Observations	19,804	19,804	16,189	16,189	15,182	15,182	15,182	15,182
R-squared	0.032	-0.116	0.041	-0.629	0.037	0.028	0.045	-0.554
Sargan		11.40		7.784		7.200		10.79
Sargan pval		0.0767		0.254		0.303		0.0950
F statistic		0.372		1.300		0.580		1.214
F pval		0.919		0.246		0.773		0.291

Table 4.12: Capital, 2005 sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2006	2006	2007	2007	2006	2006	2007	2007
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt05	0.216***	0.196	0.281***	0.448***	0.239***	0.322	0.286***	0.330**
	[0.00]	[0.32]	[0.01]	[0.12]	[0.01]	[0.22]	[0.01]	[0.14]
Observations	52,310	52,310	40,171	40,171	36,495	36,495	36,495	36,495
R-squared	0.087	0.086	0.123	0.083	0.104	0.092	0.129	0.126
Sargan		17.04		8.700		8.546		11.68
Sargan pval		0.00913		0.191		0.201		0.0694
F statistic		0.751		6.284		1.534		4.642
F pval		0.628		2.17e-07		0.151		3.30e-05

Table 4.13: Employment in heads, 2005 sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2006	2006	2007	2007	2006	2006	2007	2007
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt05	0.054***	0.429	0.089***	-0.108	0.057***	0.516*	0.088***	-0.003
	[0.00]	[0.38]	[0.00]	[0.22]	[0.00]	[0.31]	[0.00]	[0.23]
Observations	20,381	20,381	16,669	16,669	15,624	15,624	15,624	15,624
R-squared	0.019	-0.592	0.038	-0.103	0.020	-0.904	0.038	0.006
Sargan		3.515		8.824		1.118		9.043
Sargan pval		0.742		0.184		0.981		0.171
F statistic		0.343		1.495		0.693		1.332
F pval		0.934		0.164		0.678		0.230

Table 4.14: Labour remuneration per head, 2005 sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2006	2006	2007	2007	2006	2006	2007	2007
VARIABLES	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
lndebt05	-0.005*	0.772	-0.004	0.353*	-0.005	0.166	-0.005	0.487*
	[0.00]	[0.55]	[0.00]	[0.21]	[0.00]	[0.24]	[0.00]	[0.25]
Observations	19,561	19,561	16,074	16,074	15,047	15,047	15,047	15,047
R-squared	0.003	-3.987	0.004	-1.039	0.003	-0.186	0.004	-1.985
Sargan		3.147		5.808		16.96		2.761
Sargan pval		0.790		0.445		0.00942		0.838
F statistic		0.375		1.199		0.505		1.143
F pval		0.917		0.300		0.831		0.333

Table 4.15: Turnover per head, small v large sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	SME	large	SME	large	SME	large	SME	large
lndebt07	0.900***	0.286**	0.250	0.577**	0.668***	0.375**	0.381**	0.471***
	[0.20]	[0.11]	[0.16]	[0.23]	[0.15]	[0.15]	[0.18]	[0.16]
Observations	11,576	3,376	9,592	2,909	8,428	2,639	8,428	2,639
R-squared	-2.178	-0.288	-0.206	-1.358	-1.336	-0.486	-0.491	-0.869
Sargan	7.370	10.09	7.403	3.137	10.76	12.05	8.530	3.087
Sargan pval	0.288	0.121	0.285	0.791	0.0962	0.0609	0.202	0.798
F statistic	4.160	3.865	1.931	1.370	5.407	2.885	2.004	2.468
F pval	0.000139	0.000338	0.0606	0.213	3.34e-06	0.00529	0.0508	0.0159

Table 4.16: Capital per head, small v large sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	SME	large	SME	large	SME	large	SME	large
lndebt07	0.482***	0.816***	0.503*	1.091***	0.447***	0.740***	0.373*	0.681***
	[0.15]	[0.11]	[0.27]	[0.33]	[0.13]	[0.12]	[0.20]	[0.20]
Observations	11,863	8,482	12,459	4,138	8,577	3,699	8,577	3,699
R-squared	-0.166	-0.931	-0.186	-1.675	-0.136	-0.719	-0.067	-0.524
Sargan	10.98	13.80	12.18	2.626	3.457	6.634	6.731	5.236
Sargan pval	0.0889	0.0320	0.0581	0.854	0.750	0.356	0.346	0.514
F statistic	5.390	12.19	1.664	2.083	6.993	8.064	2.550	3.011
F pval	3.49e-06	0	0.113	0.0421	2.45e-08	9.20e-10	0.0127	0.00372

Table 4.17: Capital, small v large sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced 2008	unbalanced 2008	unbalanced 2009	unbalanced 2009	balanced 2008	balanced 2008	balanced 2009	balanced 2009
VARIABLES	SME	large	SME	large	SME	large	SME	large
lndebt07	-0.004	0.827***	-0.237	0.434***	0.205	0.777***	-0.081	0.346***
	[0.17]	[0.06]	[0.22]	[0.09]	[0.14]	[0.05]	[0.20]	[0.09]
Observations	25,528	22,549	21,701	13,974	16,429	12,242	16,429	12,242
R-squared	-0.003	-0.702	-0.213	0.021	0.073	-0.391	-0.054	0.072
Sargan	6.556	1.593	11.71	8.141	7.214	7.016	17.15	9.474
Sargan pval	0.364	0.953	0.0688	0.228	0.302	0.319	0.00875	0.149
F statistic	2.607	28.88	2.558	11.72	3.787	41.24	2.690	11.89
F pval	0.0109	0	0.0124	0	0.000411	0	0.00876	0

Table 4.18: Total employment in heads, small v large sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	SME	large	SME	large	SME	large	SME	large
lndebt07	-0.027	-0.070	-0.380*	-0.632***	-0.082	-0.087	-0.303*	-0.436***
	[0.08]	[0.05]	[0.21]	[0.22]	[0.08]	[0.06]	[0.16]	[0.15]
Observations	12,267	8,739	12,825	4,285	8,934	3,829	8,934	3,829
R-squared	-0.016	-0.063	-0.766	-1.852	-0.070	-0.087	-0.506	-1.002
Sargan	3.404	9.443	11.48	3.044	2.019	8.143	7.247	2.673
Sargan pval	0.757	0.150	0.0747	0.803	0.918	0.228	0.299	0.849
F statistic	5.312	12.91	1.499	2.234	5.863	8.784	2.218	3.252
F pval	4.42e-06	0	0.162	0.0288	8.24e-07	9.38e-11	0.0299	0.00191

Table 4.19: Labour remuneration per head, small v large sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	SME	large	SME	large	SME	large	SME	large
lndebt07	1.025***	0.841***	0.623***	0.878***	0.844***	0.789***	0.653***	0.732***
	[0.19]	[0.10]	[0.23]	[0.25]	[0.15]	[0.11]	[0.21]	[0.18]
Observations	11,746	8,414	12,514	4,121	8,624	3,678	8,624	3,678
R-squared	-5.478	-3.722	-2.581	-5.109	-3.794	-3.000	-2.891	-3.619
Sargan	3.957	8.091	29.40	0.924	12.85	5.673	23.54	2.284
Sargan pval	0.682	0.231	5.12e-05	0.988	0.0455	0.461	0.000636	0.892
F statistic	4.795	12.67	1.400	2.130	6.021	8.592	1.891	2.936
F pval	2.11e-05	0	0.200	0.0374	5.07e-07	1.74e-10	0.0666	0.00457

Table 4.20: Turnover per head, young v old sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	young	old	young	old	young	old	young	old
lndebt07	0.641***	0.580***	0.556**	0.161	0.760***	0.542***	0.567**	0.369*
	[0.19]	[0.12]	[0.22]	[0.15]	[0.22]	[0.11]	[0.23]	[0.20]
Observations	4,759	10,193	3,767	8,734	3,439	7,996	3,439	7,996
R-squared	-0.785	-1.239	-0.764	-0.097	-1.287	-1.102	-0.802	-0.537
Sargan	6.637	11.97	4.650	13.66	1.357	13.57	4.473	8.388
Sargan pval	0.356	0.0627	0.589	0.0337	0.968	0.0348	0.613	0.211
F statistic	3.516	5.889	2.149	1.576	3.126	6.730	2.126	1.402
F pval	0.000908	7.58e-07	0.0357	0.137	0.00272	5.61e-08	0.0378	0.200

Table 4.21: Capital per head, young v old sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	young	old	young	old	young	old	young	old
lndebt07	0.414***	0.812***	0.596**	1.124***	0.594***	0.779***	0.546**	1.106***
	[0.13]	[0.11]	[0.26]	[0.30]	[0.14]	[0.10]	[0.23]	[0.29]
Observations	5,898	14,447	4,602	11,995	4,261	11,257	4,261	11,257
R-squared	-0.043	-0.925	-0.187	-1.700	-0.200	-0.850	-0.137	-1.615
Sargan	22.44	6.564	4.777	9.993	10.46	5.779	6.033	8.651
Sargan pval	0.00101	0.363	0.573	0.125	0.107	0.448	0.419	0.194
F statistic	6.842	12.14	2.298	2.479	6.608	13.14	2.640	2.769
F pval	4.03e-08	0	0.0245	0.0153	8.61e-08	0	0.0101	0.00710

Table 4.22: Capital, young v old sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	young	old	young	old	young	old	young	old
lndebt07	0.340***	0.832***	-0.069	0.569***	0.574***	0.801***	-0.143	0.534***
	[0.11]	[0.07]	[0.18]	[0.11]	[0.09]	[0.06]	[0.19]	[0.11]
Observations	20,657	27,420	14,187	21,488	12,924	20,029	12,924	20,029
R-squared	0.073	-0.861	-0.048	-0.173	-0.011	-0.716	-0.113	-0.120
Sargan	10.38	6.620	9.005	12.07	11.99	6.389	15.80	10.32
Sargan pval	0.109	0.357	0.173	0.0603	0.0622	0.381	0.0148	0.112
F statistic	6.995	22.60	3.806	7.840	12.66	30.89	3.558	8.384
F pval	2.36e-08	0	0.000389	1.62e-09	0	0	0.000794	2.86e-10

Table 4.23: Total employment in heads, young v old sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	young	old	young	old	young	old	young	old
lndebt07	-0.010	-0.125**	-0.159	-0.646***	-0.103	-0.100**	-0.177	-0.644***
	[0.07]	[0.05]	[0.16]	[0.22]	[0.08]	[0.05]	[0.15]	[0.20]
Observations	6,260	14,746	4,831	12,279	4,484	11,533	4,484	11,533
R-squared	-0.005	-0.163	-0.153	-2.261	-0.089	-0.114	-0.179	-2.199
Sargan	3.965	6.232	17.24	10.26	5.777	8.037	14.76	10.74
Sargan pval	0.681	0.398	0.00845	0.114	0.449	0.235	0.0222	0.0969
F statistic	7.503	11.38	2.488	2.109	6.415	11.60	2.771	2.460
F pval	5.05e-09	0	0.0150	0.0392	1.56e-07	0	0.00712	0.0161

Table 4.24: Labour remuneration per head, young v old sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	unbalanced	unbalanced	unbalanced	unbalanced	balanced	balanced	balanced	balanced
	2008	2008	2009	2009	2008	2008	2009	2009
VARIABLES	young	old	young	old	young	old	young	old
lndebt07	0.704***	0.879***	0.613***	0.927***	0.863***	0.861***	0.589***	0.873***
	[0.13]	[0.11]	[0.20]	[0.25]	[0.16]	[0.10]	[0.19]	[0.22]
Observations	5,779	14,381	4,552	12,083	4,217	11,319	4,217	11,319
R-squared	-1.723	-5.090	-1.823	-6.601	-2.683	-4.835	-1.701	-5.999
Sargan	14.37	10.48	12.13	17.09	5.535	8.829	14.69	18.66
Sargan pval	0.0258	0.106	0.0591	0.00896	0.477	0.183	0.0228	0.00478
F statistic	6.045	11.82	2.209	2.266	5.896	12.51	2.235	2.597
F pval	4.80e-07	0	0.0307	0.0265	7.75e-07	0	0.0288	0.0112

Bibliography

Acemoglu, D. and A. Scott (1997). Asymmetric business cycles: Theory and time-series evidence. *Journal of Monetary Economics 40*(3), 501 – 533.

Amiti, M. and D. Weinstein (2013). How much do bank shocks affect investment? evidence from matched bank-firm loan data. *Federal Reserve Bank of New York Staff Reports 604*.

Angrist, J. D., Ò. Jordà, and G. Kuersteiner (2013). Semiparametric estimates of monetary policy effects: String theory revisited. Working Paper 19355, National Bureau of Economic Research.

Auerbach, A. J. and Y. Gorodnichenko (2011). Fiscal multipliers in recession and expansion. NBER Working Papers 17447, National Bureau of Economic Research, Inc.

Bank of England (2013, October). Trends in lending. Technical report, Bank of England.

Barnett, A., S. Batten, A. Chiu, J. Franklin, and M. Sebastia-Barriel (2014). The uk productivity puzzle. Quarterly bulletin, Bank of England.

Bentolila, S., M. Jansen, G. Jimnez, and S. Ruano (2013). When credit dries up: Job losses in the great recession. Discussion Paper 9776, CEPR.

Berger, D. and J. Vavra (2012). Consumption dynamics during recessions. *mimeo*.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008). An equilibrium model of 'global imbalances' and low interest rates. *American Economic Review 98*(1), 358–93.

Chirinko, R. S. (2008, June). σ: The long and short of it. *Journal of Macroeconomics 30*(2), 671–686.

Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008-09 financial crisis. *The Quarterly Journal of Economics 129*(1).

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45.

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2012). Innocent bystanders? monetary policy and inequality in the u.s. NBER Working Papers 18170, National Bureau of Economic Research, Inc.

Cover, J. P. (1992). Asymmetric effects of positive and negative money-supply shocks. *The Quarterly Journal of Economics 107*(4), pp. 1261–1282.

Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of economics and statistics 80*(4), 549–560.

Edgerton, J. (2012). Credit supply and business investment during the great recession:evidence from public records of equipment financing. *mimeo*.

Eggertsson, G. and N. Mehrotra (2014, 4). A model of secular stagnation. mimeo.

Flannery, M., E. Giacomini, and X. Wang (2013). The effect of bank shocks on corporate financing and investment: Evidence from 2007-2009 financial crisis. note.

FSA (2011). The failure of the royal bank of scotland: Financial services authority board report.

Gali, J. (2008). Introduction to monetary policy, inflation, and the business cycle: An introduction to the new keynesian framework. In *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.

Garcia, R. (2002). Are the effects of monetary policy asymmetric? *Economic Inquiry 40*(1), 102–119.

Giglio, S. and T. Severo (2012). Intangible capital, relative asset shortages and bubbles. *Journal of Monetary Economics 59*(3), 303–317.

Gilchrist, S. and E. Zakrajsek (2012a). Credit spreads and business cycle fluctuations. *American Economic Review 102*(4), 1692–1720.

Gilchrist, S. and E. Zakrajsek (2012b). Credit spreads and business cycle fluctuations. *American Economic Review 102*(4), 1692–1720.

Granger, C. and T. Terasvirta (1994). Modelling nonlinear economic relationships. *International Journal of Forecasting 10*(1), 169–171.

Greenstone, M. and A. Mas (2012, November). Do credit market shocks affect the real economy? quasi-experimental evidence from the great recession and 'normal' economic times. Technical Report 12-27, MIT Department of Economics Working Paper.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long-run implications of investment-specific technological change. *The American Economic Review 87*(3), pp. 342–362.

Hansen, A. (1938). *Full Recovery Or Stagnation?* W. W. Norton, Incorporated.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, 65–70.

IMF (2014, April). World economic outlook. Technical report, IMF.

Ingves, S. (2014, 11). Interview with dagens nyheter. Interview with Dagens Nyheter.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review 95*(1), 161–182.

Karabarbounis, L. and B. Neiman (2014). The global decline of the labor share. *The Quarterly Journal of Economics 129*(1), 61–103.

King, M. and D. Low (2014, February). Measuring the "world" real interest rate. Working Paper 19887, National Bureau of Economic Research.

Laubach, T. and J. C. Williams (2003). Measuring the Natural Rate of Interest. *Review of Economics and Statistics 85*(4), 1063–1070.

Lo, M. C. and J. Piger (2005). Is the response of output to monetary policy asymmetric? evidence from a regime-switching coefficients model. *Journal of Money, Credit and Banking 37*(5), 865–86.

Long, J. B. d., L. H. Summers, N. G. Mankiw, and C. D. Romer (1988). How does macroeconomic policy affect output? *Brookings Papers on Economic Activity 1988*(2), pp. 433–494.

Marx, K. (1867). *Capital: A Critique of Political Economy.* Number v. 1 in Capital: A Critique of Political Economy. Penguin Books Limited.

Montiel Olea, J. L., J. Stock, and M. W. Watson (2012). Inference in vars with external instruments.

Ongena, S., J. Peydro, and N. v. Horen (2013). Shocks abroad, pain at home? bank-firm level evidence on the international transmission of financial shocks. Discussion Paper 2013-040, Tilburg University, Center for Economic Research.

ONS (2013). Business population estimates. Technical report, ONS.

Oulton, N. and M. Sebastia-Barriel (2013, Jan). Long and short-term effects of the financial crisis on labour productivity, capital and output. CEP Discussion Papers CEPDP1185, CEP.

Paravisini, D., V. Rappoport, P. Schnabl, and D. Wolfenzon (2011, April). Dissecting the effect of credit supply on trade: Evidence from matched credit-export data. Working Paper 16975, National Bureau of Economic Research.

Pattani, A., G. Vera, and J. Wackett (2011, December). Going public: Uk companies' use of capital markets. *Bank of England Quarterly Bulletin*.

PCBS (2013). An accident waiting to happen: the failure of hbos.

Peek, J. and E. S. Rosengren (1997). The international transmission of financial shocks: The case of jap the international transmission of financial shocks: The case of japan an. *The American Economic Review*, 495–505.

Peek, J. and E. S. Rosengren (2000). Collateral damage: Effects of the japanese bank crisis on real activity in the united states. *American Economic Review*, 30–45.

Piketty, T. (2011). On the long-run evolution of inheritance: France 18202050. *The Quarterly Journal of Economics 126*(3), 1071–1131.

Piketty, T., G. Postel-Vinay, and J.-L. Rosenthal (2014). Inherited vs self-made wealth: Theory & evidence from a rentier society (Paris 18721927). *Explorations in Economic History 51*(C), 21–40.

Ramey, V. A. and S. Zubairy (2014). Government spending multipliers in good times and in bad: Evidence from u.s. historical data.

Reinhart, C. and K. Rogoff (2012). *This Time is Different: Eight Centuries of Financial Folly*. publisher.

Rey, H. (2014, 1). Panel on secular stagnation. Remarks by Helene Rey at Panel on Secular Stagnation.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Rose, A. K. and T. Wieladek (2014). Financial protectionism? first evidence. *The Journal of Finance 69*(5), 2127–2149.

Smets, F. and G. Peersman (2001). Are the effects of monetary policy in the euro area greater in recessions than in booms? Working Paper Series 052, European Central Bank.

Stock, J. H. and M. Yogo (2002). Testing for weak instruments in linear iv regression.

Summers, L. (2013, 11). Crises yesterday and today. Remarks by Larry Summers at 14th Annual IMF Research Conference, Nov 8 2013.

Summers, L. H. (2014, Apr). U.s. economic prospects. *Bus Econ 49*(2), 65–73.

Thoma, M. A. (1994). Subsample instability and asymmetries in money-income causality. *Journal of Econometrics 64*(1-2), 279–306.

Van Reenen, J. (1996). The creation and capture of rents: Wages and innovation in a panel of u. k. companies. *The Quarterly Journal of Economics 111*(1), 195–226.

Vavra, J. (2013). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *mimeo*.

Weise, C. L. (1999). The asymmetric effects of monetary policy: A nonlinear vector autoregression approach. *Journal of Money, Credit and Banking 31*(1), 85–108.

Whelan, K. (2000). A guide to the use of chain aggregated nipa. *FEDS Working Paper 2000*(35).

Woodford, M. (2011). *Interest and prices: Foundations of a theory of monetary policy.* princeton university press.

Yang, F. (2009). Consumption over the life cycle: How different is housing? *Review of Economic Dynamics 12*(3), 423 – 443.