ESSAYS IN MACROECONOMICS: CYCLES AND FRICTIONS

Mario Lupoli

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

2023

Identifiers to use to cite or link to this thesis:

DOI: https://doi.org/10.17630/sta/528
http://hdl.handle.net/10023/27865

Essays in Macroeconomics: Cycles and Frictions

Mario Lupoli

University of St Andrews

This thesis is submitted in partial fulfilment for the degree of
Doctor of Philosophy (PhD)
at the University of St Andrews

February 2023

Abstract

This thesis contributes to different literatures in macroeconomics with frictions. In the first two Chapters I consider imperfections in the credit market and how these can amplify monetary policy shocks. I start from a purely empirical model, which identifies monetary policy shocks and then I develop a structural model with an explicit market for mortgage loans intermediated by a banking sector. Households and banks are each facing a different optimisation program. I show that this model better captures the volatility of macroeconomic aggregates than alternative frictionless cases. This richer modelling setting assigns a more complicated role to the monetary authority, as the policy rate influences asset prices, nominal debt and bank profitability in addition to intertemporal consumption. The Third Chapter is concerned with wage rigidity and how to measure it. We define it as relative to the wage one would expect under Nash bargaining. Then we develop a statistic for wage rigidity, the Nash Wage Elasticity (NWE) by regressing actual wages on the Nash bargained wage. Most of our calibrations yield a NWE between 0 and 0.1, signifying that actual wages are very rigid and that the Nash wage is a poor description of the business cycle. We calibrate a search and matching model to match our estimated NWE, showing how this modification translates into greater cyclical fluctuations. In the fourth Chapter I analyse the causal relation linking index investment to commodity future prices. I show that standard Granger causality results cannot be taken at face value given the extraordinary movement in prices during the Great Financial Crisis. I apply instead a Time-Varying Granger test apt to gauge the evolution of the causal relation, showing how future prices are endogenous to index investment flows at particular points in time, generally supporting the hypothesis of financialization in the commodity market.

Dayadhvam: I have heard the key
Turn in the door once and turn once only
We think of the key, each in his prison
Thinking of the key, each confirms a prison
Only at nightfall, aethereal rumours
Revive for a moment a broken Coriolanus

(T.S. Eliot - The Waste Land)

Acknowledgements

I am indebted to Matthew Knowles, without him I would not have even started this PhD, let alone finished it.

I wish to thank my dear PhD colleagues, Nayha Mansoor, Zhiting Wu, Nicolò Bandera, who welcomed me in St Andrews and the ones who joined after me, in order appearance: Lazaros Zografopoulos, Ciarán MacDomhnaill, Jacob Rees Stevens and Kairat Umargaliyev.

I owe special thanks to Nikolay Chernyshev and Rui Luo.

For Chapter I, I wish to thank Matthew Knowles, Ozge Senay, Alan Sutherland, Oliver Ashtari Tafti and Roderick McCrorie. I'm grateful to two anonymous referees of the International Journal of Central Banking and the Co-Editor Òscar Jordà for their comments.

For Chapter II, I am obliged to Ozge Senay, Matthew Knowles, Jonathan Swarbrick, Alan Sutherland, Jacob Stevens, Oliver Ashtari Tafti and Giacomo Candian. I'm grateful to David Cobham, James Moberly and the participants of the Money, Macro and Finance Research Society Annual PhD Conference (2022).

For Chapter III, we thank Alan Sutherland, Moritz Kuhn and seminar participants at the University of St Andrews, University of Cologne, Econtribute, Austrian Economic Society Annual Conference (2022), Royal Economic Society Annual Conference (2022), University of Rochester Stockman Conference (2022) and the Money, Macro and Finance Research Society Annual Meeting (2022) for very helpful suggestions and comments.

For Chapter IV, I wish to acknowledge Professor Christopher Gilbert for his kind support and suggestions, the participants at the Rimini Centre for Economic Analysis Conference on Recent Developments in Economics, Econometrics and Finance (2022) and the participants at the St Andrews Econometric Workshop (2022).

I am grateful to my dear friend Oliver for the fruitful discussions in the past few years and his helpful comments were instrumental to the completion of this Dissertation.

Declaration

Candidate's declaration

I, Mario Lupoli, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 80,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in August 2019.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date Signature of supervisor

Permission for Publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Mario Lupoli, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date　　　　Signature of candidate

Date　　　　Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Mario Lupoli, understand that by declaring that I have original research data or digital outputs, I should make every effort in meeting the University's and research funders' requirements on the deposit and sharing of research data or research digital outputs.

Date Signature of candidate

Permission for publication of underpinning research data or digital outputs

We understand that for any original research data or digital outputs which are deposited, we are giving permission for them to be made available for use in accordance with the requirements of the University and research funders, for the time being in force.

We also understand that the title and the description will be published, and that the underpinning research data or digital outputs will be electronically accessible for use in accordance with the license specified at the point of deposit, unless exempt by award of an embargo as requested below.

The following is an agreed request by candidate and supervisor regarding the publication of underpinning research data or digital outputs:

Embargo on all of electronic files for a period of 2 years on the following ground(s):

- Publication would preclude future publication

Supporting statement for embargo request

The digital output is a time series derived from public sources. I'd keep it embargoed in order to be able to publish it in an academic journal at a later stage.

Title and Description

- I require an embargo on the title and description

Date Signature of candidate

Date Signature of supervisor

FUNDING

This work was supported by the Economic and Social Research Council [grant reference number ES/P000681/1].

Publications

Some of the work described in this dissertation has been published previously or is under review for publication. The following list gives an overview of these publications.

Articles in Peer-Reviewed Journals

Chapter I of this Thesis has been published as:

Mario, L. 2022. 'Deleverage and Defaults in the United Kingdom'. *International Journal of Central Banking* 18 (5): 53–110

Articles with Co-Authors

Chapter III of this thesis, 'The Nash Wage Elasticity and its Business Cycle Implications' is co-authored with Dr. Matthew Knowles.

CONTENTS

List of Figures

List of Tables

INTRODUCTION

Almost the entirety of current macro-modelling literature develops on micro foundations[1] and it is confronted with the tasks of explaining business cycles fluctuations and provide a framework of analysis for the practical design of policy (Goodfriend and King, 1997).

However, oftentimes stylised models fail to capture the observed empirical moments on key macro-variables, falling short of their first task above and thus being of limited help with policy-making. Sometimes these shortcomings are well known, other times they are more elusive and become manifest around crises and turning points, bringing about a revitalised interest in the underlying dynamics behind critical macroeconomic phenomena.

The thorny point here is how the dealings of optimising agents can lead to the kind of volatility observed in data. This is often achieved by 'throwing a spanner' in the smooth works of representative agents with the introduction of frictions.

The overarching topic of this dissertation pertains to frictions perturbating the macro-economic environment. Those misalignments impede an efficient market allocation creating consequential dynamics visible in data and motivate policy trade-offs. The fundamental reason for including frictions is to add model mechanics that create internal propagation of imbalances besides the external shocks.

All four chapters focus on the observable features of markets that are grounded in data but are not satisfactorily reflected in workhorse models. This dissertation brings together frictions that exist in three key macro markets that all interact closely to determine how economies react to external shocks and disturbances. These three markets are the credit market, labour markets and commodity markets.

Credit market variables move at a lower frequency than the overlapping business cycle, hence strengthening the negative effects of a recession at the time of its

[1]Based on rationally optimising representative agents and profit-maximising firms.

reversal (e.g., when an economic crisis coincides with a credit crunch). This makes it even more important to understand the mechanisms at work in determining how credit market variables change and respond to external disturbances. It is an open question whether it is optimal to intervene on credit aggregates.

Similarly, contemporary macro-models fall short in explaining the cyclicality of labour market stocks. Searching and matching models describe wages as way more rigid than they appear from published statistics. This is the essence of the Shimer's critique, which we aim to address estimating a measure for wage elasticity.

Lastly, bubble trajectories in commodity index investments appear at particular points in time, signifying a conspicuous departure from pricing fundamentals casting doubts on the efficient market hypothesis. Bubbles originate in practically all asset classes and are a stable feat of the real estate and stock markets.

In this dissertation, we retain an empirical approach and concentrate on commodity prices to gain insights on change in causality. In particular, we apply statistic tests to detect them in real time. The dramatic increase in long-only index positions followed by the GFC and a generalised collapse was so glaring that prompted US regulators to establish trading limitations to prevent excessive speculation. Yet, the academic consensus still supports the allocative efficiency of markets. We will show that commodity markets have not been efficient during the GFC and at particular points in time.

The first two chapters show how monetary policy drives demand/supply dynamics for household loans and how interest rates decisions are transmitted to the wider economy via the banking sector. The main contribution of these two thesis chapters is to describe and model more precisely the mechanism of transmission associated with the credit channel of monetary policy in the presence of key frictions such as information asymmetry, capital requirements and sticky prices. In contrast to most of the relevant literature, these chapters explicitly account for the intensive and extensive margins of lending. Debt deleverage decisions happen on the intensive margin, capturing how intensely households make use of banking loans. Conversely, insolvency decisions are usually made along the extensive margin, i.e., whether households forcefully withdraw from borrowing – hence defaulting.

Chapter II, "Deleverage and Defaults in UK" is an empirical paper using the IV-VAR methodology. Applying a high-frequency instrument based on market data, I am able to identify monetary policy shocks in UK and trace their effects on the credit sector variables, as house prices, the household debt stock and insolvencies. By calculating impulse response functions to the said monetary policy shock, I am able to assess the reaction of the credit sector showing that individual insolvencies rise very fast - peaking at 3% - as opposed to the sluggishness of real debt adjustment.

This result, which holds under a wide array of different identifications, points to the existence of accelerating dynamics as house prices and debt are permanently reduced upon an interest rate shock. I therefore conjecture about 'accelerator-like' dynamics in the mortgage market, as the opportunity cost of lending rises excluding some risky projects from access to debt finance. This reduces house prices and housing investments at the time of the money-tightening and in subsequent periods as future housing holdings depend on current ones. Besides that, and on the extensive margin, the increased default frequency feeds into banks' balance-sheets further tightening credit and dampening asset prices.

Figure 0.1: Credit View of Monetary Transmission

Author's elaboration that expands on 'Credit view' channels

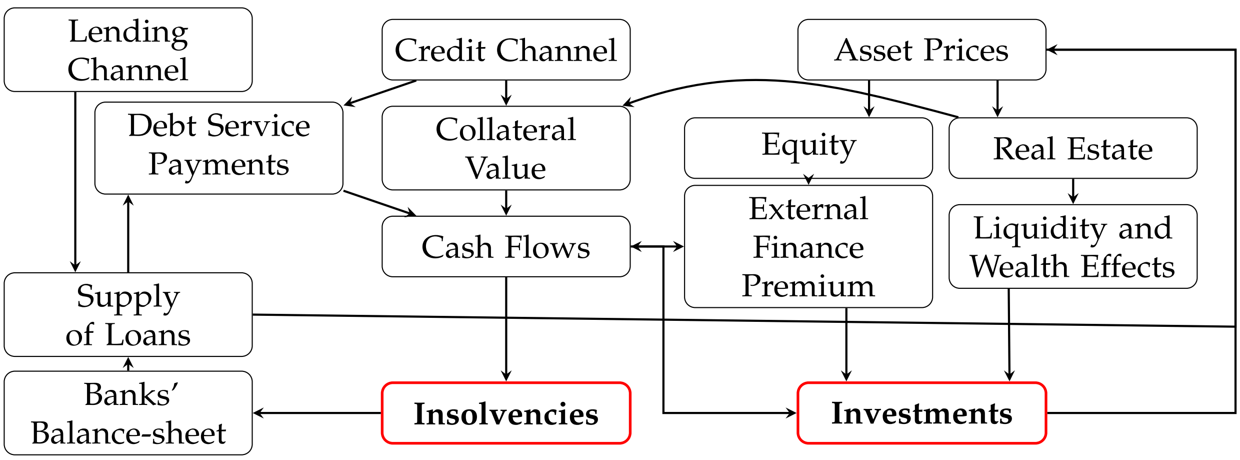

The empirical results are at odds with the large literature on 'leaning against the wind' – i.e., tightening of monetary policy to cool off an overheated housing market reducing house prices and household debt. Therefore, in the second chapter of this dissertation "Leaning against the Wind and the Default Channel of Monetary Policy" I explore asset prices targeting in a medium-scale tractable DSGE model. The modelling framework allows for monetary policy to overtly follow a Taylor rule that can be augmented to encompass credit aggregates.

The objective of this essay is twofold: I analyse a relevant policy question using a rich model that is faithful to the 'credit view' of the monetary transmission mechanism (see Fig. 1) and therefore is apt to trace structurally the channels described above at an aggregate level from the VAR presented in "Deleverage and Defaults in UK".

There, I present a standard New-Keynesian setting to model house prices dynamics where households debt is constrained by a time-varying loan-to-value ratio motivated by an informational asymmetry. I can disentangle between demand and supply of mortgages, as the external finance premium constrains the supply of loans and the net-worth channel boost the demand enriching the business cycle dynamics of an otherwise standard model.

While financial frictions are understood to provide a credible propagation mechanism internal to the model and produce meaningful fluctuations on shocks, standard labour market frictions are criticised for their inability to do so. Searching and matching models (SAM) fail to generate plausible fluctuations in unemployment on productivity shocks (a fact known as the 'Shimer Puzzle').

In "The Nash Wage Elasticity and its Business Cycle Implications" we address this critique by proposing a new measure of wage rigidity, the Nash Wage Elasticity (NWE). The NWE represents the percentage change in the actual wage rate when the wage that would occur under Nash bargaining changes by 1%. This result qualifies wage cyclicality relative to the Nash wage, as many models are concerned with wage procyclicality but without a benchmark is unclear what is the model-implied degree of stickiness. Providing an OLS and IV estimation of the NWE is the first contribution of this third Chapter. The second contribution is showing how calibrating the NWE in a standard SAM model can approximate the volatility of the actual wage.

We do so by building a broad modelling framework that nests discrete time versions of many different models from the search and matching literature. Combining equations and publicly available stocks and flows data, we can compute a model-implied Nash wage in a way similar to the 'business cycle accounting' literature. We find that the Nash wage is more procyclical than all commonly used actual wage measures (all employees, new hires wage and user cost of labour) and thus NWE is significantly less than one. These estimates indicate that, for both continuing workers and new hires there is a high degree of wage rigidity and Nash bargaining provides a poor description of wage setting.

We find that when we calibrate the NWE in a parsimonious model, it makes an huge difference to fluctuations in unemployment as we are able to match US business cycle moments. This means that aggregate wage rigidity is quantitatively important for macro-modelling and it can be reflected into general equilibrium models by a savvy calibration of the NWE.

The dissertation then turns to commodity markets, where the Great Financial crisis (GFC) put a stop to runaway commodity prices, which collapsed in 2008. The build-up of futures contracts was accompanies by an increase of index-investment, suggesting that the interest in commodity futures had expanded from professional hedgers to a broader class of long-only investors.

This anecdotal fact begs the question on whether financialization of commodities markets translates into artificially higher prices.

The fourth chapter delves on whether money in-flows associated with index investments help predicting commodity future prices. Applying a time-varying Granger causality test we are able to identify changing points in the index investment-future prices relation and gauge its strength. This result represents an addition on the literature on index investment and commodity prices. Whilst the causal relation has been previously found in selected commodities, as crude oil prices or certain agricultural commodities, there is no settled academic position on causality.

A second order of considerations for Granger-testing stems from the extreme movements in prices associated with the GFC, as existing statistical tests are not robust to structural breaks and non-stationarity and may produce spurious results. To this purpose we use a recently introduced time-varying Granger test based on a recursive-evolving algorithm.

We find whole-sample causality for non-ferrous metals whilst for some agricultural commodities and crude oil causality is just around the GFC. For non-ferrous metals this may be due to their inelastic supply and elastic demand, whereas for the other commodities this may suggest a certain degree of financialisaton associated with long-only investment positions.

We will now look at the first Chapter on 'Deleverage and Defaults in UK'.

Chapter One

Deleverage and Defaults in UK

Introduction

In the United Kingdom the number of household insolvencies has continuously risen from their post-Crisis trough in 2015, and, in 2018 the number stood as high as at 2010, close to 130,000 year-end new defaults. By the end of the same year, the real household debt stock exceeded its all time historical level previously set in 2008.[1] The Bank Rate, on the other hand, was unprecedentedly low at 0.1%. It is therefore topical to understand how a recessive monetary shock might impact financial stability in such a context.

The effect of monetary policy on financial variables has received great attention in recent works, with a number of authors arguing for a tightening of monetary policy in situations of rising house prices or rising debt (Borio and Lowe, 2011; Gambacorta and Signoretti, 2014), to '*lean against the wind*'. The advantage of such a stance appears to be particularly relevant for highly levered economies, where local policymakers might want to cool down debt accumulation and asset prices. Several empirical papers have indeed found an effect of monetary policy on debt showing a marked deleverage effect coupled with a decline in house prices (Hofmann and Peersman, 2017; Robstad, 2018; Laseen and Strid, 2018) (henceforth 'papers by HPRLS').

[1] Data sources are in Appendix A.1.

The aim of this paper is to extend this empirical framework to understand if the responses of personal insolvencies to a monetary policy shock warrant particular policy attention. To do so, I set up a Vector Autoregression (VAR) model akin to the ones present in the papers by HPRLS. The Hofmann and Peersman paper has investigated a panel of economies, whereas the latter two focused on Norway and Sweden respectively. Here I will concentrate on the United Kingdom. The elements of novelty of my work consist in including the number of individual insolvencies among the regressors and using an external instrument to identify the dynamic system. I highlight the role of defaults, which are only tangentially treated in the papers by HPRLS. In particular, my VAR analysis delivers a response on impulse of household insolvencies to a monetary shock, among to other variables common to the literature.

The first contribution of this paper is to show that households' credit quality makes up a separate channel of monetary transmission. A policy contraction produces some sudden and disorderly deleverage, thereby increasing the aggregated insolvency level. In the VAR, Insolvencies react much quicker than deleveraging and I find that a monetary tightening leads to an uptick in individual insolvencies, they peak at 2.2% after 8 quarters versus 0.36% debt reduction at the same horizon. I then conjecture that household insolvencies might be part of a *financial accelerator*-like mechanism feeding back to financial variables.

Moreover, the instrumental VAR model results also deliver policy relevant answers in regards of a flexible inflation targeting. Debt-to-GDP ratio in UK is not significantly different from zero upon a tightening and it is therefore an ineffective measure when targeting financial stability. UK Debt-to-Income ratio declines but the Granger causality test confirms that real debts are endogenous to house prices, and house prices as a policy target are the object of a vast literature.

The importance of households' credit risk in the monetary policy transmission has implications for macro-prudential policy. The HPRLS papers have generally assumed that debt deleverage would be orderly and neutral to households' credit quality. However theory (Bernanke et al., 1999) and empirical evidence suggest that more defaults happen in distressed environments. The papers by HPRLS do not reconcile this twofold aspect of debt deleverage, implicitly assuming that families either pay back their debts, stop rolling them over or renounce to take additional leverage after an interest rate tightening. Following this line of thought, a policy induced deleverage might even be desirable from a macro-prudential angle. But what if this produces more defaults?

'*Leaning against the wind*' stance postulates the use of interest rate to target financial variables, this translates into monetary policy what is a common macro-prudential principle: creating risk buffers at the cycle height to counter down swings (e.g. the Countercyclical Capital Buffer measure).[2] That raises the question of whether traditional macro-prudential policy would be better in achieving financial stability rather than interest rate policy. Given that disorderly insolvencies are an important part of the transmission of monetary policy then there is a potential welfare case for using the interest rate in lieu of other more apt instruments.

Since VAR models are sensitive to identification assumptions, I also explore various alternative sign-restriction identification schemes under a Bayesian approach in Section 1.4 as a robustness check. The result is that the baseline model inference continues to hold also when the shock is sign-identified.

I present the results of time-varying Granger causality tests to uncover the causal direction between two variables at the time (Section 1.5). Such tests identify the changing points of causal relations among variables, thereby addressing the discontinuity represented by the 2008 crisis. This Granger causality testing is performed on a reduced form version of the model, is not dependent on the modelling choices established in the first part of this paper. The policy rate Granger-causes insolvencies when it is high, ceasing to be relevant to bankruptcies from when it plummeted to 2%. House prices drive debt dynamics whilst the opposite only holds during recessions.

The paper is structured as follows: in the first section I present the literature behind household credit decisions and the transmission of monetary policy. I shall devote the second section to comparing different UK papers and how they have dealt with the identification challenge in retrieving structural innovations. My model is then presented in Section 1.3 with impulse response analysis. The remaining sections present the sign-restriction approach and the time-varying Granger test.

[2]'*the countercyclical capital buffer regime may also help to lean against the build-up phase of the credit cycle in the first place. In downturns, the regime should help to reduce the risk that the supply of credit will be constrained by regulatory capital requirements that could undermine the performance of the real economy and result in additional credit losses in the banking system.*' BIS description of Countercyclical Capital Buffer (CCyB) (underline mine).

1.1 Related Literature

The concept of credit risk features prominently in the seminal theoretical literature on '*financial market frictions*'. The fact that borrowers may fail in honouring their debts provides a micro-foundation for costly-state verification Bernanke et al. (1999) and collateral constraints Kiyotaki and Moore (1997) models. These efforts have established how lenders and borrowers optimising decisions can produce stronger fluctuations in production and investments through oscillation in firms' net worth in a New-Keynesian general equilibrium context.

In real life, economy-specific structural factors such as the proportion of adjustable rates over fixed and the average loan maturity dictate whether household would either take up more debt or deleverage on the back of shorter-term attrition in lending rates and house prices. This makes the theoretical impulse response functions bounded to their own model hypotheses and represents the reason why the problem at hand has been often approached from an empirical angle.

This paper retains an applied approach and is similar in spirit and methodology to three papers developed by authors affiliated to Central Banks (HPRLS papers). The aim is to shed light on how household finance responds to tight monetary policy shock and the methodology is a Vector Autoregression analysis. In this section I will mainly focus on these three papers with an eye on a few selected general equilibrium models that have discussed a '*leaning against the wind*' stance.

As the Swedish Riksbank '*leant against the wind*' to curb house prices through targeting the private debt stock, a discussion arose regarding the trade-offs of setting monetary policy in response to asset prices and debt variables. Gambacorta and Signoretti (2014) present such framework in a DSGE environment, finding that a mixed policy rule produces greater gains in a highly leveraged economy.

An opposite conclusion appears in the theoretical framework laid by Svensson (2014) , who argued that a rule responding to household debts has little effect on the overall stock since income reacts to policy adjustment faster than debts producing recessive consequences. Hence the cost of deviating from inflation targeting is higher than the benefit as it bears a disproportionate effect on output and inflation. [3] Laseen and Strid (2018)'s paper is a direct response to Svensson (2014) and finds a strong decline in real household debts and Debt-to-GDP ratio following a tightening.

[3] Using a calibration for the Swedish economy.

The IV-VAR presented below in this paper supports this household debt dynamics with UK data, although capturing no significant movement of Debt-to-GDP ratio.

Hofmann and Peersman (2017) takes a slight different angle, hinting at a *'debt service channel'* for monetary policy transmission by which interest and principal payments relative to the existing household debt stock become more onerous as lending rates increase with a monetary tightening. This makes the economies with a higher stock of household debts more prone to a deterioration on a interest rate contraction. My position is conceptually similar to theirs in arguing for a credit quality channel of monetary transmission. Not only a rate tightening impacts households' debt burden but pushes some into default. This aspect is lacking in Hofmann and Peersman (2017), who assume a benign debt deleverage, i.e. driven by principal repayments, a view generally common across HPRLS.

To summarise the literature up to this point: whilst the effect of tight monetary policy is well understood in regards of debts and house prices, there is no consensus on the gains in terms of financial stability. I therefore contribute to this debate by adopting the HPRLS VAR framework and supplement it with individual insolvencies. I also discard the Cholesky identification to avoid defending a particular recursive ordering, relying instead on an external series of shocks. Hopefully, this effort will help nuancing more the effects of a mixed policy aimed to stabilise credit aggregates.

The paper most similar to mine is Piffer (2018), who tries to reconcile the *'financial accelerator'* model (Bernanke et al., 1999) with an Instrumental VAR akin to the one proposed below. He specifically includes delinquencies in his analysis on US and investigates whether a policy easing shock causes more or less defaults. This research question stems from partial equilibrium models of the risk-taking channel of monetary transmission. In a lower interest environment, lenders may have the incentive of targeting riskier clients to increase their interest income. This may lead to a deterioration of lending portfolios and therefore the increase of non-performing loans. Piffer (2018) empirically finds that an increase in wealth dampens default. This finding is consistent with Bernanke, Gertler, and Gilchrist's DSGE model, which shows that positive net-worth effects prevail over risky lending pitfalls.

Nevertheless, debt might build up in periods of relative financial quietness. A prolonged period of low inflation may be conducive to a crisis (Borio and Lowe, 2011) as supply side developments may feed into an overly positive sentiment causing lending and asset price booms. Credible monetary policy reinforces the low risk perception and adds to the general exuberant feeling (Borio and Lowe, 2011). The loosening of credit standards coupled with yield compression often precedes the downturn and have the potential of exacerbating the ensuing crisis. This connects back to New-Keynesian DSGE models as many credit variables are pro-cyclical as net-worth is.

DSGE models do not account directly for defaults (Goodhart and Tsomocos, 2011; Gambacorta and Signoretti, 2014) but they are a normal feature of the economic cycle [4] and they increase in crises. Household insolvencies endogenously arise from net-worth down-movements, which are reinforced by falling house prices in downturn periods. Feedback effects from banks' balance-sheet may also result in a reduced credit supply and amplify the cyclical swing. The recessive potential of a monetary policy rule that purposely reacts to credit variables deteriorating household finances is therefore still to be fully investigated.

1.2 The Identification Challenge

1.2.1 Monetary Policy in UK

In this section I briefly outline the history of monetary policy in the UK, since this is relevant for the identification of monetary policy shocks. In recent history the Bank of England (BoE) has not been bound by a single monetary rule. It targeted the money supply from 1976 to transition to the exchange rate, at first informally tracking the Deutsche Mark (1987-88) and from '89 by maintaining a floating band around a fixed basket of ECU participating currencies within the of Exchange Rate Mechanism (ERM) (King, 1997). Following Black Wednesday and its withdrawal from the ERM, UK moved towards pure inflation targeting in October 1992. A change of monetary regime happened when the new Labour executive granted to BoE operational independence in 1997, although it did not change in the focus on inflation targeting. With the Bank of England Act of 1998, the Monetary Policy Committee (MPC) was given the responsibility of formulating monetary policy in lieu of acting on a target rate set by the Treasury. The main policy instrument is the Bank Rate but asset purchases were made as the Bank rate reached zero lower bound in March 2009.

[4] As Goodhart and Tsomocos (2011) note, very seldom the repayment rate is 100%.

Concomitantly to this policy shifting in the early '90s, the Bank of England underwent a series of structural reforms to improve the transparency of the decision making process (King, 1997). It published its first Inflation Report in August 1993 and set a fixed calendar for MPC meetings and the publication of the relevant minutes thereafter to counter-balance Treasury's discretionality and, to the extent possible, separate the rate-setting process from the Government political agenda. The management of expectations has become a separate channel of transmission and unconventional policy gaining prominence since. From March 2009, the MPC also voted on the size of assets purchase programmes. The Central Bank adopted an additional communication lever, a *'forward guidance'* policy aimed to clearly communicate under which conditions monetary policy is to be tightened and quantitative easing modified (Dale and Talbot, 2013).

1.2.2 Identification of Exogenous UK Policy Shocks

The policy regime is not irrelevant to VAR identification and bears powerful consequences on the model-implied conclusions. Interest rate is endogenous to the state of the economy therefore to assess the impact of shocks, one would need to find interest rate developments that are plausibly exogenous. The Cholesky identification is the most used strategy in structural VARs literature but presents a number of issues that I will discuss below. Because of its properties, it has been considered unreliable to retrieve UK policy shocks. I shall outline what I mean by identification and survey alternative approaches used in the British VAR literature.

Generally, identification boils down to performing a discretionary orthogonalisation of the time-regression residuals. Such transformation is needed to interpret errors as exogenous shocks originating outside of the system (Sims, 1980). This means that the researcher has to formulate and make clear some valid hypotheses to back the identification decision before estimating the VAR equations. Finding an economically suitable identification is per se a daunting task, [5] which requires careful pondering as it reflects assumptions on behaviour of the analysed economy and on causal chains linking the regressors.

A straightforward method to achieve full identification is to impose restrictions on contemporaneous reactions of macro-variables to monetary policy shock such

[5]*'The number of structural VARs is limited only by the inventiveness of the researcher'* (Stock and Watson, 2001). Indeed, many different identifications have been proposed so far, such as sign or long run restrictions". For a survey see Ramey (2016b)

that each variable respond to impulse with a time lag from the one ordered right before. This recursive identification is computationally inexpensive and is achieved by operating a Cholesky decomposition on the reduced for residual covariance matrix.

Such triangular system has a number of drawbacks: (1) progressive delayed reactions are difficult to defend with lower frequency data or including financial variables, which are likely to adjust simultaneously with the macro ones. (2) Cholesky-identified VARs tend to produce at times puzzling impulse response functions with results at odds with textbook theory. This may be due to the omission of forward looking variables that the Central Bank uses to inform its decision. An incorrect identification may pick up the endogenous component of interest rates, i.e. when the Monetary Authority moves the rate with a predictable rule, responding to developments in the other endogenous variables (Arias et al., 2019). (3) when different monetary regimes coexist within the same sample, instrumenting the interest rate in a Cholesky ordering may be incorrectly identifying policy shocks (Rusnak et al., 2013).

A key difference between the papers by HPRLS and the literature regarding the UK is that the former all use a Cholesky decomposition, [6] which has been openly impugned and discarded in many of the UK papers. A reason behind that choice might be that in the British cases researchers have endeavoured to achieve, either directly or indirectly, a double goal: trace the effects of a monetary policy shock and assess the transmission mechanism over a very long sample. The need for a different identification is dictated by the length of the period analysed and the breadth of the research questions tackled, almost assuming a historical perspective. [7]

We have seen in the previous paragraph that the shift to inflation targeting is a source of discontinuity in the data. Cloyne and Hürtgen (2016) address that by including in a VAR a novel narrative series as endogenous regressor, which means supplementing an otherwise standard system with new information. They find that the response of inflation to monetary innovations is similar if taken pre and post 1992. What changes is the volatility of exogenous shock series, which is significantly

[6] Although Robstad (2018) also proposes an alternative sign restriction identification and different Cholesky ordering.

[7] A tabular summary of key cited studies is presented in Appendix A.2 with a comparison of their research questions and sample periods.

reduced arguably thanks to a more attentive steering of monetary policy by the BoE.

Ellis et al. (2014) make explicit the historical dimension of their study as they set out to deal with different policy regimes analysing a sample from 1975 to 2005. A Factor Augmented VAR model is meant to mitigate the omitted variable problem by including factors from some 350 variables the Central Banker might react to. Structural changes in UK policy making show in time-varying impulse responses, as prior to 1992 monetary policy was neutral to inflation. After that date, monetary policy gained in efficiency producing clear responses in CPI and asset prices to a monetary tightening.

Analysing an overlapping time span (1974 - 2005), Mountford (2005) finds that monetary policy accounts for a limited variation of output. Monetary policy reaction to the other variables in the VAR are thus quantitatively more important than exogenous monetary shock, hence the title of the paper is *'Leaning into the wind'*.

So we have established some econometric issues when applying VAR analysis to UK: (1) There is a clear policy change in 1992, (2) monetary policy might endogenously react to variables that are either inside or outside the VAR, (3) previous UK studies have all been concerned in disentangling actual shock from the '*systematic component*' of monetary policy (as defined in Gerko and Rey (2017)).

My approach differ from the historical one, as I am estimating a VAR on a circumscribed time period, broadly coinciding with BoE reforms on adopting an inflation targeting. Nevertheless, the IV-VAR is apt to produce more reliable results with low-frequency data as opposed to a Cholesky decomposition, as it allows to disregard a battery of rather mechanical assumptions about the system ordering.

Gerko and Rey (2017) and Cesa-Bianchi et al. (2020) articles are more recent and they translate to the UK the Instrumental VAR methodology that I shall describe in the next paragraph and use for my analysis. Gerko and Rey (2017) finds a significant price and production puzzles when applying the Cholesky identification to 1982-2015 data which an instrumental identification mitigates. In that instance, a monetary tightening is neutral to RPIX and Industrial Production and drives up lending spreads. That weak response might again be due to the length of sample and policy heterogeneity. The significant pass-trough of the interest rate shock on corporate and mortgage spreads is shared with Cesa-Bianchi et al. (2020), who in turn find a significant decrease in economic activity measured by a rise in unemployment.

1.3 The Instrumental Vector Autoregression Approach

1.3.1 The Model

The model is an Instrumental Vector Autoregression (VAR-IV). Since monetary policy might be endogenous to the other variables, I use an external instrument to identify the interest rate equation. Following this stream of empirical research, I identify the shock using an index of daily surprises on the Sterling Deposit Future adopting the approach pioneered by Gertler and Karadi (2015) and Mertens and Ravn (2013), although with lower frequency data and applied to UK variables.

High frequency identification aims to isolate exogenous shocks which are not connected to the other time series in the VAR (Ramey, 2016b). To do so we need firstly a reduced form VAR that takes the following shape:

$$y_t = C + \sum_{j=1}^{p} \underbrace{A_p}_{B_0^{-1}B_j} y_{t-j} + \underbrace{u_t}_{B_0^{-1}w_t} \tag{1.1}$$

And then we need identifying restriction on the matrix B_0^{-1} to retrieve the monetary policy shocks. Here an instrument Z respecting the following conditions comes handy:

$$\mathbf{E}[Z_t w_t^{p'}] = \phi \tag{1.2}$$

$$\mathbf{E}[Z_t w_t^{q'}] = 0 \tag{1.3}$$

Z must be correlated to monetary policy shocks $w_t^{p'}$ and uncorrelated to the other structural shocks $w_t^{q'}$.

So, as in Mertens and Ravn (2013) and Gertler and Karadi (2015), I proceeded estimating a two stage regression (TSLS) following these steps:

1. Retrieve the error u_t from the reduced form representation.

2. Compute the following regression $u_t^p = a + xZ_t + e$, of which fitted values are $\hat{u_t^p}$.

3. Estimate $u_t^q = \frac{s^q}{s^p}\hat{u_t^p} + \xi$.

Where the first stage isolates the exogenous part dependant on the instrument Z_t and the second stage yields an estimate of the ratio $u_t^q = \frac{s^q}{s^p}$. The separated s^q and s^q can be obtained from partitioning of the structural coefficients matrix B and covariance matrix Σ given the restrictions $\Sigma = \mathrm{B}_0^{-1}\mathrm{B}_0^{-1'}$ and $u_t^q = \frac{s^q}{s^p}$.

$$B_0^{-1} = \begin{bmatrix} \beta_{11} & \beta_{12} \\ \beta_{21} & \beta_{22} \end{bmatrix} \qquad \Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \tag{1.4}$$

Were β_{11} and Σ_{11} are $k \times k$ instruments used (here a scalar) and β_{21} and Σ_{21} are then $k \times (n-k)$. The identification is thus provided by the closed form solution firstly derived by Mertens and Ravn (2013).

$$\beta_{21}\beta_{11}' = \frac{s^q}{s^p} \tag{1.5}$$

$$\beta_{12}\beta_{12}' = (\Sigma_{21} - \frac{\beta_{21}}{\beta_{11}})'\mathrm{Q}^{-1}(\Sigma_{21} - \frac{\beta_{21}}{\beta_{11}}\Sigma_{11}) \tag{1.6}$$

$$Q = \frac{\beta_{21}}{\beta_{11}}\Sigma_{11}\frac{\beta_{21}'}{\beta_{11}} - (\Sigma_{21}\frac{\beta_{21}'}{\beta_{11}} + \beta_{21}\beta_{11}\Sigma_{21}') + \Sigma_{22} \tag{1.7}$$

$$\beta_{11}\beta_{11}' = \Sigma_{11} - \beta_{12}\beta_{12}' \tag{1.8}$$

The first column of Σ can then be used to compute the impulse response functions for the monetary policy shock.

1.3.2 Stationarity and Data

In my baseline VAR specification I use UK Bank Rate, GDP, GDP Deflator, House Prices, Real Household Debt, Individual Insolvencies in this order. Data are taken in log-levels and are quarterly, spanning from Q1 1987 to Q4 2018 for a total of 128 data points.[8] An element that differences the present work from previous UK studies and the papers by HPRLS is the inclusion of Individual Insolvencies, which are compiled by the UK Insolvency Service and composed of Individual Voluntary Arrangements, Debt Relief Orders and Bankruptcies. Real debt series comes from ONS' Households Loans series, which includes secured debt (mortgages and equity

[8]Sources and Charts are reported in Appendix A.1.

releases) and unsecured debt (as credit cards and student loans). House Prices series is the UK average house price. This series follows exactly the same dynamics of the house price index, which is calculated normalising the average house price, and has the advantage of being measured in GBP.

I include 2 lags in accordance with the Bayesian Information Criterion (BIC), which is both consistent and parsimonious in the lag selection. The VAR system is stationary being the eigenvalues of the companion-form matrix outside the unit circle.

Cheng et al. (2019) deal with potential non-stationarity of series in a IV-VAR estimation finding that for the estimated coefficient the error is '*asymptotically negligible*'. In presence of non-stationarity, IRFs are asymptotically normal with the covariance matrix depending on the persistence of each series. Cheng et al. (2019) hence derive a GMM estimator for IRFs with an optimal weighting matrix based on a consistent covariance estimator which enables the computation of IRFs that are robust to non-stationarity of regressors. I have used that method to derive non-stationarity robust IRFs as part of my robustness checks (reported in Appendix A.4.3).

The external instrument Z_t I use to pin down the exogenous component of reduced form residuals spans from 1997 to the end of the sample. It is calculated around specific monetary policy events from a handpicked dataset. In accordance with the literature, my dataset of policy events includes three macro-categories of BoE appointments: announcements, MPC Minutes disclosures and Inflation Report publication.

Monetary policy is announced roughly every six weeks by BoE and the MPC meeting minutes are disclosed on the following day. In terms of communication, BoE has been publishing the Inflation Report since August 1993 and the minutes of monthly MPC meetings since August 1996 no later than six weeks after the meeting (two weeks from 1998). From 2015 MPC minutes and the Inflation Report have been disclosed on the meeting day. In November 2019, the Infation Report changed name into Monetary Policy Report and now carries more background information on the overall economic conditions underpinning the monetary policy decision.

1.3.3 The Instrument

Following the existing high-frequency IV-VAR literature on UK Cesa-Bianchi et al. (2020); Gerko and Rey (2017), I use the ICE LIFFE Three Month Sterling (Short) Future.[9] This future contract captures the three-month ahead interest rate and thus is a forward looking measure of interest rate surprises. According to papers mentioned, the instrument can capture the surprises associated with unconventional monetary as the publication of Inflation Reports and MPC Minutes update the expectations of the public with fresh information on the state of the economy and on what motivated the policy decision (Gerko and Rey, 2017).

The instrument is then calculated as follows:

$$Z_t^{daily} = -(P_{t,\tau+1}^{daily} - P_{t,\tau}^{daily}) \tag{1.9}$$

Since the Sterling Future is quoted at discount ($P_t = 100 - InterestRate$), the minus sign before the parentheses in Equation 1.9 denotes that positive monetary surprises corresponds to an increase in the interest rate. The subscript τ is the day of the relevant policy event and $\tau + 1$ is the day after.

In their paper Cesa-Bianchi et al. calls their surprise index 'daily' or 'high frequency', whereas here I reserved the label 'daily' to my indicator. (Cesa-Bianchi et al., 2020) it is more of a 'trading time' indicator, being constructed on a database of tick-by-tick data around monetary policy events (exactly 10 minutes before and 20 after). My indicator uses the daily difference in settlement prices for that derivative contract, thus it constitutes a lower frequency instrument than what is normally used in the literature. The contract settles at 11.00 a.m. therefore daily differences capture the money surprises as the announcement is disclosed at 12.00 a.m.

My first-stage regression (See 2) displays a F-Statistic (1,85) of 41.56 and R-squared is 0.32, meaning that the instrument is a strong one. These results exceed the 10 F-Statistic threshold Stock and Yogo (2005) a rule of thumb under which the power of the instrument is deemed weak.[10]

Similarly to Gertler and Karadi (2015), I derived a monthly and quarterly series by cumulating and differencing the rough surprises series in the following fashion:

[9] Intercontinental Exchange Website.

[10] I have used Gerko and Rey (2017) and Cesa-Bianchi et al. (2020) instrument in my baseline specification finding that the former is not a useful measure in my context [F-stat(1,69) = 0.63, R^2 = 0.01], whereas the latter makes a strong instrument [F-stat(1,70) = 21.74, R^2 = 0.24].

1. I have calculated the daily surprise in Eq. 1.9 as at the days in which took place a relevant policy making decision (meaning monetary policy committee announcements, minutes or inflation report disclosures),

2. I cumulated them and

3. took a 31 days rolling average.

The monthly indicator is then the end of month first difference of the series obtained with step (3). Similarly, the quarterly surprises series that I have used as an external instrument in my baseline specification took a 3-period sum of monthly surprises.

Figure 1.1 represents the three instrumental variables side by side both in their monthly formulation (top pane) and quarterly aggregated (bottom pane). In Cesa-Bianchi et al. (2020), the largest surprise is the one associated with the interest rate cut from 5% to 0.5% from September 2008 to March 2009. Gerko and Rey (2017) choose to omit policy rate announcements from their dataset. This is because they think that announcement press releases do not provide any new information due to their brevity.

In this paper I include base rate announcements and, due to these differences, my monthly surprise series is closer to the Cesa-Bianchi et al. one, though being more volatile. Monetary 'surprises' that are only present in my dataset are in March and May 2018, when the market started to price July 2018 tightening, updating its expectation thanks to policy's forward guidance. In general, I detect a slight increase in volatility from 2017 probably due to general markets' expectations of an upcoming policy normalisation after an extended period of low interest rate and the Brexit vote induced rate cut of 2016.

Figure 1.1: UK Money Surprises

I derived a daily frequency indicator of monetary surprise as in Cesa-Bianchi et al. (2020) (solid blue line). In my case money surprise is the change in price for a 3-month sterling derivative future during the day of a monetary policy announcement.

1.3.3.1 Instrument Robustness

Cesa-Bianchi et al. (2020) propose a Sargan-Hansen over-identification test to control for non-monetary information potentially 'contaminating' the instrument. Under the null hypothesis there is no correlation between instruments and reduced form residuals (i.e. the instruments are both valid). Since this statistical testing strategy requires more instruments than endogenous variables, I then leverage on the Cesa-Bianchi et al. (2020) dataset using their high frequency indicator alongside with Cloyne and Hürtgen (2016) narrative series as joint excluded instruments (quarterly re-sampled).

I perform this test twice, coupling my baseline instrument separately with both the externally available series. In both cases I cannot reject the null hypothesis with 0.01 significance level, concluding that the daily instrument derived in the above section is apt to identify the exogenous monetary shocks. This result is particularly important when using the Cloyne and Hürtgen series, which is based on a narrative approach and explicitly excludes other factors influencing monetary policy (Cesa-Bianchi et al., 2020).

To further gauge the robustness of my baseline model, I have tried a variety of instruments as alternatives to the bank rate either in the VAR or as excluded instruments for monetary policy surprises. These instruments include: the 5, 10 and 20 years UK ZCB rates, 3 month, 2, 5 and 10 years nominal par yields, the 3 months Libor swap rate and 3 months GBP/USD forward rate. Shorter rates produce similar results, with lower first stage statistics than the combination of policy rate and instrument I end up using in the baseline model.

1.3.4 Impulse Response Functions

Figure 1.2: Structural Impulse Responses of the Baseline IV-VAR(2) model on UK Data.

Solid line represents point estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times

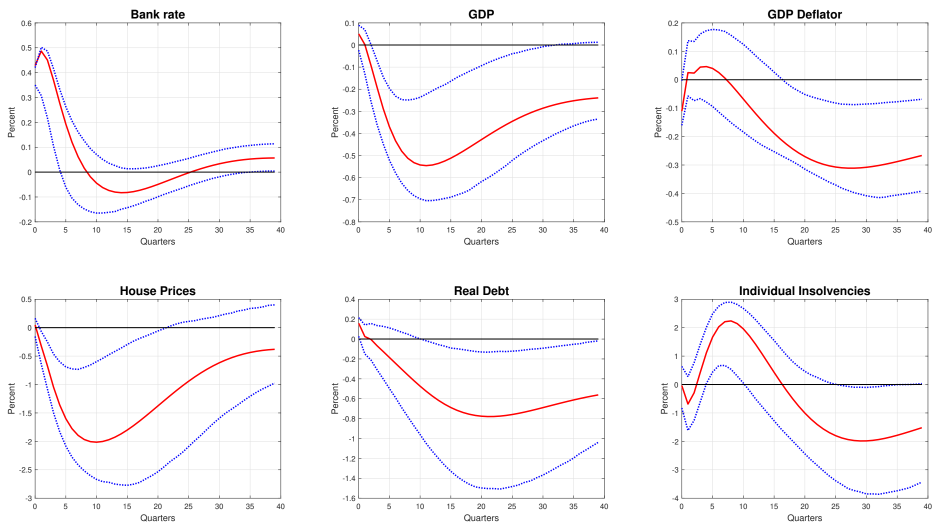

The purpose of this section is to trace the effects of a monetary tightening - a one standard deviation monetary surprise - to the six variables in the dynamic system, providing intuition for the different channels at play (Figure 1.2). Confidence bands are derived using a wild bootstrap method as originally proposed in Gonçalves and Kilian (2004) and later widely adopted in the IV-VAR literature (e.g. Mertens and Ravn (2013); Gertler and Karadi (2015)).

The responses on a monetary impulse of GDP and inflation are consistent with textbook macro-models, with a rate hike reducing investments and price level on the back off demand-side developments. If seen through DSGE lenses [11], house prices and real debt responses are conditioned by frictions in the provision of credit, supporting the institutional views (as in HPRLS) that a money tightening impacts house prices and real debts.

The decrease in house prices and real debt may be due to 'accelerator-like' dynamics that involves on one hand an increase in the cost of borrowing, and the opportunity cost of lending vis-à-vis the higher base rate, and on the other hand the households' net worth. This mechanism is captured in theoretical models (Bernanke et al., 1999) and it is self-reinforcing as a contraction depresses current period investments having lasting effects on the future price of capital, further dampening investments and net worth. This puts strains on the availability of external finance besides debt-servicing costs, as households are likely to pledge housing properties as collateral when entering into recourse debt contracts. Hence the fall in house prices leads to a fall in real debt. This amplification mechanism feeds into consumption and output, exacerbating the downturn.

Insolvencies are anti-cyclical, increasing in downturns and tapering in benign periods across the business cycle. Qualitatively, the hump-shaped response I obtain of insolvencies to a recessive shock is consistent with Bernanke et al. (1999). As per their model, defaults are rising following a decrease in capital, here represented by housing. Capital acquisition is proportional to net worth, so a shock that reduces the return to capital transmits to wealth and raises the default probability. Insolvencies are highly correlated with the unemployment rate (excluded from the baseline VAR) as they are connected to the level of economic activity.

My VAR specification features a decrease in GDP, house prices and debts. Both in the Cholesky specification (Appendix A.4.1) and in the instrumental variable approach, insolvencies are rising following a monetary policy shock (within a 90% confidence interval). The Cholesky decomposition does not yield any counter-intuitive puzzling response in that case, just a stronger and more persistent positive response in the GDP Deflator, otherwise being qualitatively consistent with the IV-VAR.

Shock's contractionary effects on real GDP and house prices persist after as many as more than 30 quarters. The decline in house prices is somehow comparable to what

[11] Piffer (2018) retained a similar approach comparing his VAR findings with general equilibrium models featuring financial market imperfections.

has been found by Robstad (2018) whereas debt deleverage dynamic is stronger in terms of magnitude and more persistent. It shows the through after 20 quarters with signs of recovery thereafter, but after 40 quarters it is still significantly below zero. The GDP deflator response is somewhat weak in the aftermath of the policy decision and becomes significantly negative after 17 quarters.

House prices' response starts very close to zero, highlighting that even without a strict zero restriction, there is no contemporaneous reaction to a monetary policy shock. Individual Insolvencies show a significant uptick before the tenth quarter, with a peak at 2.2% 8 quarters after the fundamental shock. They fully revert to zero prior to the twentieth quarter after the shock and then are significantly below null at a 90% significance level.

In Figure 1.3 I plot the effect of a monetary policy shock to key consumption aggregate series, individually substituting them for GDP in the baseline model. Household Total Final Expenditure (consumption) quickly decreases from a near zero response at the time of the shock. Durable consumption (house goods, vehicles) instantaneously falls by -0.5%, when non-durable consumption (food, drinks) spikes at time 0 to decay to nil within the first quarter. This illustration may offer a view on an accelerator-like effect on the transmission mechanism involving household debt and insolvencies appearing in the data. There is a quick and persistent demand side reduction of investment and consumption upon a tightening. Hence insolvency may happen on the back of a reduction in collateral value and tighter borrowing constraints. This finding is consistent the Monacelli's DSGE model, which attributes the slump in durable consumption to collateral constrains becoming tighter after a rate hike.

The policy takeaway is therefore that insolvencies play a role in the monetary transmission mechanism as an exogenous tightening has sizable short run effects on the level of defaults, causing their surge in the immediate wake of the relevant decision. Real debts show a sluggish response, arriving at their lowest level much slower than delinquencies. By the time defaults arrive at their peak in 8 quarters, debt has reduced only by 0.36%. The response on impulse of insolvencies is hump-shaped and becomes significantly negative after its spike, signalling that tight monetary policy can achieve a modicum of financial stabilisation in the longer-run.

Figure 1.3: Structural Impulse Responses of Consumption Aggregates in an IV-VAR(2) model on UK Data.

The figure represents the response of impulse of consumption aggregates when individually substituted to GDP in the baseline VAR. Solid line represents point estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times

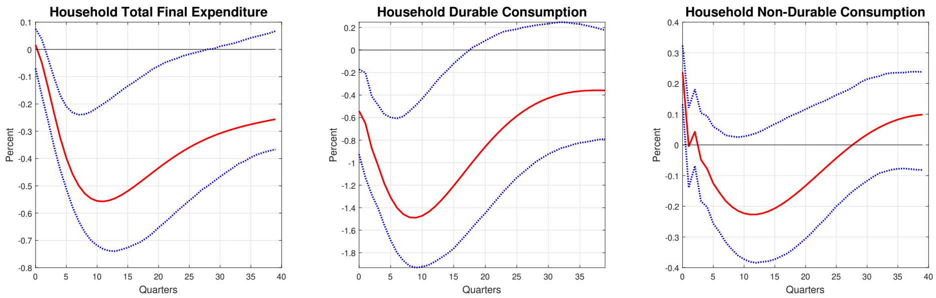

1.3.4.1 Policy Relevance and Comparison with papers by HPRLS

My results regarding house prices and debt are broadly comparable with the papers by HPRLS (Table 4.3), which have made use of the same regressor in spite of the different countries in analysis. It shows however a stronger response of both variables on inpulse. This may be due to the different identification strategy, which here pins down exogenous shocks with the help of an external instrument.

One of the reasons behind the empirical modelling in researching the matter at hand, is the lack of agreement on what the theoretical response of debts is on a shock. Svensson argued through a DSGE example that in Sweden the Debt-to-Income response to a policy tightening can be positive because of the short tenor of loans and the low prepayment rate (Svensson, 2014) and the same applies to Debt-to-GDP. In some cases, DSGE models that are assessing the benefits of a '*leaning against the wind*' policy stance are also ambiguous on stating the costs (Svensson, 2017).

I have controlled other potential policy targets in the VAR by substitution Real Debt (Appendix A.4.4) with alternative regressors. Tight monetary policy does not result in a meaningful change of the Debt-to-GDP ratio, which appears to raise following a tightening shock but is never significantly different from zero. This result is shared with Robstad (2018) and contradicts Laseen and Strid (2018).

When used in my specification, Debt-to-Income ratio follows an undetermined path up to the sixteenth quarter and then is briefly significantly negative (Appendix A.4.5). A Granger-causality test highlights that the causality relation goes from house prices to debt and therefore Debt-to-Income as monetary policy target might work indirectly through the steering of real-estate prices.[12]

A specification more similar to Cesa-Bianchi et al. (2020) is presented in Appendix A.4.6. It includes unemployment as a measure of economic activity and mortgage and corporate rates beside the other variables from the baseline model. This specification is insightful in highlighting the policy rate close to 1-to-1 pass-through on the quoted household mortgage after two quarters since the shock, whereas the corporate rate response is weaker and noisier. In all cases insolvencies respond similarly than the baseline model.

Table 1.1: Comparison with Previous Studies

Peak Response to a standard deviation (or 1%, when marked with an asterisk) monetary policy shock on House Prices and Real Debt. From papers' text body and visual inspection of impulse response charts.

Authors	Country	Method	Identification	Peak House Prices Response	Peak Debt Response
Laseen and Strid (2018)	Sweden	Bayesian (Litterman Prior)	Recursive	-0.20%	-0.20%
Robstad (2018)	Norway	Bayesian (Inverse Wishart Prior)	Recursive	-3.00%*	-1.00%*
Hofmann and Peersman (2017)	Across Panel	OLS	Recursive	-1.70%*	-1.20%*
Mario Lupoli	UK	OLS	Daily Frequency	-2.00%	-0.77%

1.3.4.2 Monetary Policy and Information Shocks

A key consideration for the success of the instrumental identification is that the instrument is uncorrelated with shocks in variables other than the one directly instrumented (Eq. 1.2 and 1.3). In this paper the risk of a spurious identification is greater than in the rest of the literature due to the reliance on daily surprises, an indicator sampled at a lower frequency than trading time.

I address this concern through two separate interventions: I sign-identify a pure monetary shock using the Jarociński and Karadi (2020) method and I test the instrument relevance with a Sargan-Hansen over-identification test (see above, in Section 1.3.3.1).

Jarociński and Karadi (2020) devised an empirical strategy to identify the information shocks and separate them from the pure monetary one. The methodology exploits two high frequency series: monetary surprises and stock price surprises, but instead of using them as excluded instruments, they are included as endogenous

[12] Whereas it can be directly impacted by macroprudential policy as in the form of LTV ratios or capital adequacy requirements.

variables in a Bayesian VAR. The system is then sign-identified by imposing ex-post restrictions on impulse response functions.

Here, the inclusion of stock market surprises is fundamental to separate two different shocks. Stock market surprises are deemed to move in the same direction of money surprises following an information shock, i.e. when the Central Bank discloses additional positive information about the state of the economy together with a monetary policy decision. In the case of a pure monetary shock, the stock market surprises will move oppositely to money surprise.

In practice, I enforce a sign restriction scheme which allows only the two high-frequency indicators to move simultaneously and I impose zero restrictions to the contemporaneous response of other variables (exactly as in Jarociński and Karadi (2020)). This identification is based on the block recursive scheme presented below in Section 1.4. I calculated the surprises in the daily FTSE 350 Index around monetary policy decisions in the same way I computed money surprises (Section 1.3.3).

In Figure 1.4 I compare the monetary policy shock to the positive information shock. This alternative identification strategy is instrumental to provide a qualitative benchmark to my baseline IV-VAR. A pure monetary shock continues to produce insolvencies even under these stricter identification assumptions. Defaults are more front-loaded than in the baseline instrumental identification, peaking after 5 quarters. This continues to suggest that pure monetary surprises are relevant to households' credit quality.

The caveat here is that this exactly identified scheme is based on stronger identification assumptions than the baseline IV-VAR as it is a mix of sign and zero restrictions. Also, in order to use the derivative high frequency instruments as endogenous variables in the VAR, I throw away their missing values, effectively running this VAR on a subset of 87 observations, making the inference less stable.

Figure 1.4: Structural Impulse Responses of a sign-restricted block recursively-identified VAR(2) model on UK Data identified as in Jarociński and Karadi (2020). The solid line represents the median response and dotted lines are the 68% percentile bands associated with the Monetary Policy Shock. Dashed line is the median response to the Information Shock.

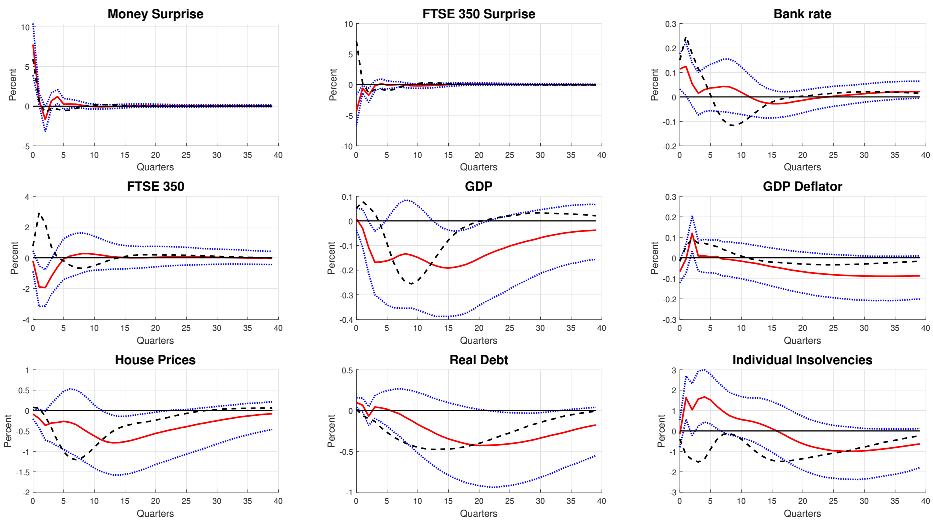

1.4 Robustness Check: Sign Restrictions

As a robustness check for the instrument identified VAR presented in Section 1.3, I considered a sign identification, as pioneered by Uhlig (2005) and applied to UK data by Mountford (2005).

This specification can achieve an identification of the monetary policy shock by restricting impulse responses to be either a positive or negative for a number of periods after the shock. This eliminates puzzles by construction producing impulse responses that match the textbook knowledge on what the qualitative consequence of a shock is.

1.4.1 The Bayesian VAR Model

Consider a VAR model as in Eq. 1.1:

$$y_t = C + A_1 y_{t_1} + A_2 y_{t_2} + \ldots + A_p y_{t_p} + u_t \tag{1.10}$$

In which $u_t \sim N(0, \Sigma_u)$. It can be re-written in a compact form of:

$$y = X\beta + U \tag{1.11}$$

Were $X = (I_n \otimes X)$ and $\beta = vec(A_1, A_2, ..., A_p, C)$. The VAR model is estimated through a Bayesian approach. The Bayes theorem enables us to approximate the posterior density given a sampling distribution and prior beliefs. In particular, the chosen prior is the Inverse-Wishart, the conjugate of the multivariate normal covariance matrix:

$$\beta | \Sigma_u \sim \quad \text{i.i.d.} \quad N(\beta, \Sigma_u \otimes \beta) \tag{1.12}$$

and

$$\Sigma_u \sim IW(\Psi, \upsilon) \tag{1.13}$$

The Inverse-Wishart is an informative prior parametrised by a semi-definite Ψ matrix and υ degrees of freedom. The conjugacy implies a posterior distribution of the same family of the prior allowing simpler estimation of the parameters.

1.4.2 The Structural Form

The VAR model in Equation (1.10) is a reduced form of a model where $A_i = B_0^{-1} B_i$ and the model errors are a weighted average of structural shocks $u_t = B_0^{-1} w_t$, as in the underbrace of Eq. 1.1.

Differently from the case illustrated above, the Bayesian setting entails embracing a priori beliefs on the parameters of B_0 (Miranda-Agrippino and Ricco, 2018) as the selected prior is informative.

I have tried different forms of identification in order to recover the structural monetary shock, belonging to the following categories:

1. Partial Identification;

2. Exact Identification.

The first identification procedure doesn't attempt to identify all structural shocks but only a monetary policy one. Conversely, other identification schemes do by means of exact identification of the system. In both cases a mix of sign and exclusion restrictions is imposed over the parameters of B_0^{-1} to overcome any potential counter-intuitive response to monetary shocks, the 'price puzzle'.

1.4.3 Partial Identification

According to Uhlig (2017) a useful heuristic is to verify the reasonableness of restrictions by only imposing restrictions justifiable by textbook economic theory and remain 'agnostic' on the variables which response to a shock is to be investigated.

So the first sign-restiction specification that I have tried is the most parsimonious one, in the spirit of Uhlig (2005). I only try to retrieve the first vector of the covariance matrix imposing three restrictions to the structural policy shock, which I am interested in identifying. As a benchmark for using sign restrictions to control a VAR, I have used the recent paper by Cantore et al. (2020), which has the advantage of establishing straightforward sign-identification rules for a monetary policy shock with the aim of imposing as few restrictions as possible and to do so in accordance with known macro-models. I deem a monetary policy shock to:

- interest rate increases upon a monetary shock;
- decrease of GDP upon a monetary shock;
- decrease of the GDP deflator upon a monetary shock;

In this case restrictions on the covariance matrix Σ have the form:

$$\begin{pmatrix} u^{IR} \\ u^{GDP} \\ u^{Defl} \\ u^{HP} \\ u^{Debts} \\ u^{Ins} \end{pmatrix} = \begin{array}{c} \begin{array}{cccccc} w_m & w_y & w_3 & w_4 & w_5 & w_6 \end{array} \\ \begin{bmatrix} + & * & * & * & * & * \\ - & * & * & * & * & * \\ - & * & * & * & * & * \\ * & * & * & * & * & * \\ * & * & * & * & * & * \\ * & * & * & * & * & * \end{bmatrix} \end{array} \tag{1.14}$$

where the asterisk denotes unrestricted coefficients and the $+/-$ signs indicate the restricted direction on a shock impact. As a difference with Cantore et al. (2020), I do not impose sign restrictions up to the second time-period, only limiting the contemporaneous responses, hence being more sparing with the number of assumptions. Throughout this section and the next I have used Arias et al. (2014) algorithms rather than the Uhlig (2005)'s ones. The former are based on finding an orthogonal rotation matrix through the QR decomposition of a randomly generated matrix of normal numbers, which have the uniform Haar distribution. The impulse responses are shown in Figure 1.5.

Figure 1.5: Structural Impulse Responses of a sign-restricted identified VAR(2) model on UK Data.

The solid line represents the median response and dotted lines are the 68% percentile bands.

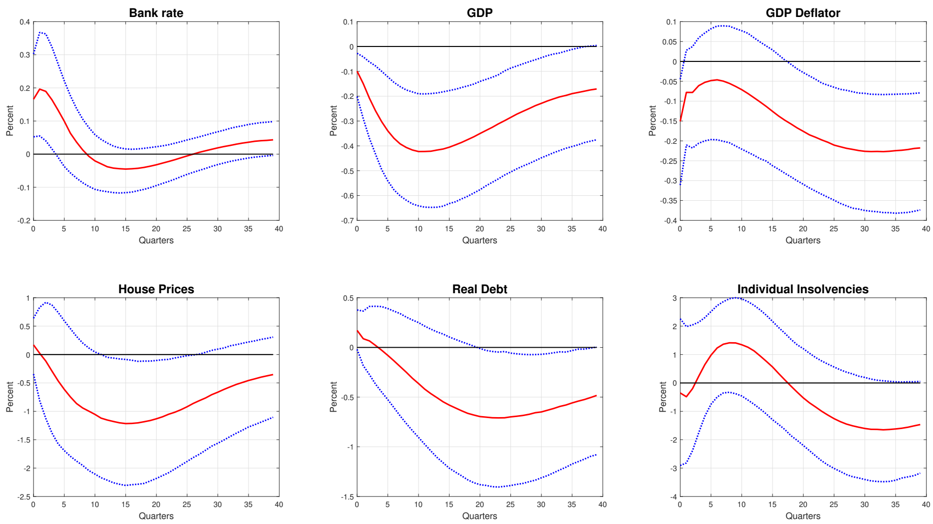

The impulse response functions are not very informative. Whilst maintaining a similar shape to the IV-VAR they are weaker, displaying very wide posterior density percentile bands. To correct that specification and to improve the shock retrieval, I have used the approach of Arias et al. (2019) in imposing restrictions on the '*systematic component*' of monetary policy.

This identification has the benefit of only restricting the interest rate equation of the VAR system and boils down to two sets of restrictions:

1. The interest rate only contemporaneously responds to GDP and price level;

2. The contemporaneous response of interest rate to GDP and price level is positive.

The Restriction (1) allows the interest rate setting process to be consistent with a standard Taylor Rule and Restriction (2) captures the endogenous component of a given policy decision, i.e. the Central Bank hikes the rate simultaneously to an increase of output and prices. An important feature of this approaches is that it does not force GDP and Deflator to be negative at a given horizon, but pins down their response in assuming to what aggregates the Central Bank reacts to.

This identification scheme yields clearer IRFs that are again similar to the baseline IV-VAR model (in Figure 1.6). GDP, House Prices are immediately declining, whereas the Deflator is significantly negative in the longer term. Insolvencies show a short-term hump, increasing to the 2% on impulse, falling in the same ballpark as in the unrestricted baseline IV-VAR. The covariance matrix presents in this case three zero restrictions (0 in the scheme below - as per restriction 1), as the interest rate is deemed not to react to house prices, debts and insolvencies within the same period.

$$\begin{pmatrix} u^{IR} \\ u^{GDP} \\ u^{Defl} \\ u^{HP} \\ u^{Debts} \\ u^{Ins} \end{pmatrix} = \begin{array}{c} \begin{array}{cccccc} w_m & w_y & w_3 & w_4 & w_5 & w_6 \end{array} \\ \begin{bmatrix} + & + & + & 0 & 0 & 0 \\ * & * & * & * & * & * \\ * & * & * & * & * & * \\ * & * & * & * & * & * \\ * & * & * & * & * & * \\ * & * & * & * & * & * \end{bmatrix} \end{array} \tag{1.15}$$

Figure 1.6: Structural Impulse Responses of a restricted identified VAR(2) model on UK Data identifying the systematic component of policy.

The solid line represents the median response and dotted lines are the 68% percentile bands.

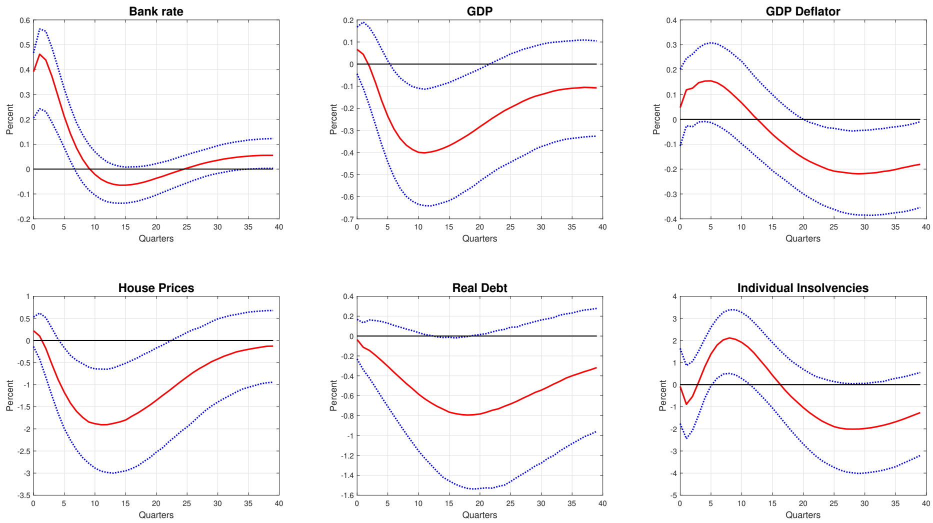

1.4.4 Exact Identification

The second battery of sign restrictions hinges on exactly identifying the whole model imposing $n \times (n-1)/2$ restrictions on the structural impact matrix in order to recover the structural shocks that are not a linear combination of others.

In the discussion on the Cholesky decomposition above, we have seen that the recursive restriction pattern holds justifiable under an economic standpoint as it is seen as a way to establish a causal chain among variables. In this section I have used the same ordering as in Section 1.3 to implement sign-restrictions in two different recursive systems.

The first one is a standard Cholesky system, with Σ upper triangular. It revolves around the standard assumption that the variables are affected by a monetary policy shock according to their ordering, in this case: GDP, inflation rate, house prices followed by real debts and insolvencies. Sign restrictions are imposed up to the second period of the impulse response function.[13]

The second identification scheme provides a block-recursive identification, the

[13] I have omitted the scheme of this identification and its IRFs as they are very similar to the block-recursive ones below.

variables are grouped in two separate blocks. The first 3×3 block represents the main macroeconomic variables ordered as stated at the beginning of this section. Last three variables constitutes the household debt market. The second block variables' shocks do not feed into the first block macroeconomic aggregates meaning that they do not have a contemporaneous effect on the first block. The relevant impulse responses are shown in Figure 1.7.

This intuition behind that scheme is that the credit variables react with some lag on monetary impulse and contemporaneously among themselves, being house prices, real debts and insolvencies interrelated. This evidence is also supported by the baseline model, where house prices and Insolvencies responses started very close to 0 without imposing exclusion restrictions on their respective coefficients.

The system is exactly identified as the coefficients associated with the macro variables (represented as dots in the below scheme) are dictated by starting covariance matrix and are not affected by further QR rotations.

$$\begin{pmatrix} u^{IR} \\ u^{Prod} \\ u^{Defl} \\ u^{HP} \\ u^{Debts} \\ u^{Ins} \end{pmatrix} = \begin{array}{c} \begin{array}{cccccc} w_m & w_y & w_3 & w_4 & w_5 & w_6 \end{array} \\ \begin{bmatrix} \bullet & \bullet & \bullet & * & * & * \\ 0 & \bullet & \bullet & * & * & * \\ 0 & 0 & \bullet & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * & * \end{bmatrix} \end{array} \tag{1.16}$$

Figure 1.7: Structural Impulse Responses of a sign-restricted block recursively-identified VAR(2) model on UK Data. The solid line represents the median response and dotted lines are the 68% percentile bands.

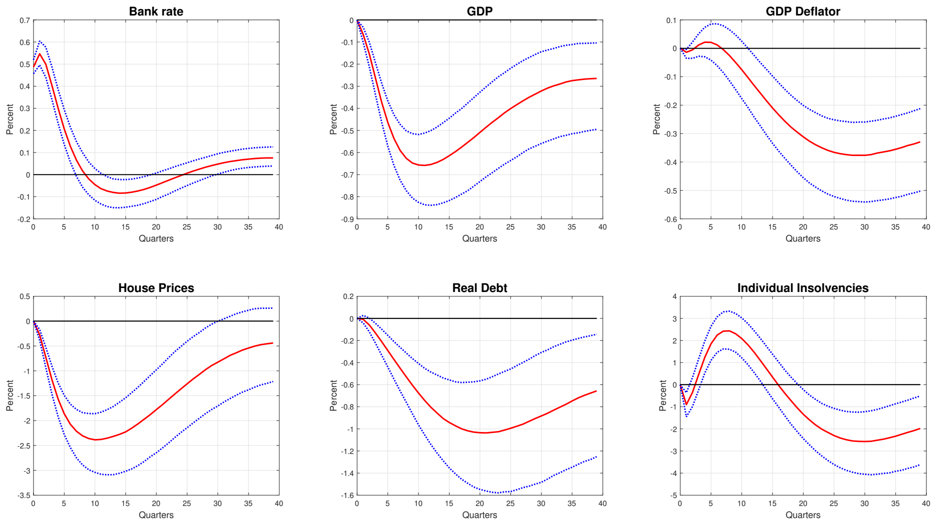

1.5 Causal Inference

I test for causality using a Time-Varying Granger-causality test. Granger tests are widely used in the analysis of VAR as they enable the researcher to understand which variable makes a useful predictor of others within the same system. They test p zero constraints to the coefficients matrix. When we fail to reject the null hypothesis of no Granger-causality from a regressor to another, we infer that the former is a good predictor for the latter.

A standard Granger test based on Wald statistics is reported in Lütkepohl (2005) and Shi et al. (2018)[14]:

$$W = [R\,vec(\hat{A})]'[R((X'X)^{-1} \otimes \hat{\Sigma})R')]^{-1}[R\,vec(\hat{A})] \tag{1.17}$$

and

$$W \sim \chi^2\,(p) \tag{1.18}$$

Where $\hat{A}$ represent the matrix of reduced form VAR coefficients and $\hat{\Sigma}$ the estimated

[14] In Shi et al. (2018) the matrix of coefficient is row-vectorised, in Eq. 1.17 I report a version with column-vectorisation.

covariance matrix. X is the matrix of lags and R is a $[n \times (k^2 p + k)]$ constraints selection matrix where p are the lags, k the VAR dimension and n the number of restrictions to be tested.

Shi et al. (2018) have recently proposed an alternative way to carry out the static test in Eq. 1.17. Computing the Wald statistic over the span of the entire VAR averages the information and potentially produces misleading inference. In particular, such a test would not reveal shifts in Granger-causality relations with the relevant changing points.

They hence base their time-varying testing strategies on a series of nested computation of the Wald statistics on data sub-samples. Starting from the first data point, the Wald statistic is computed on an arbitrary long sub-sample, which is then rolled one period ahead. At each iteration forward, a number of ancillary regressions is calculated expanding the sample backwards until it includes the first observation. The relevant statistic is then a *Supremum Norm* of the set of Wald statistics (SW) calculated for each iteration forward. When the SW exceeds a certain critical value for the first time a changing point in causality relation is identified.

I use this test to address the 2008 Crisis discontinuity in the dataset, during which variables showed extreme behaviour. I also use the Shi et al. (2018) version of the test that is robust to conditional heteroscedasticity (in Eq. 4.4) given that it is applied to reduced form residuals of Eq. 1.1, which I have identified as endogenous in the first part of this paper. An implication of the Shi et al. (2018) paper is that the asymptotic distribution of the Wald test should hold when there is not cointegration, given that the VAR is stationary.

$$W = T_w [\mathrm{R}\, vec(\hat{A})]' [\mathrm{R}((\mathrm{V}^{-1} \hat{\Omega} \mathrm{V}^{-1}) \mathrm{R}')]^{-1} [\mathrm{R}\, vec(\hat{A})] \tag{1.19}$$

Where $\mathrm{V} = \hat{\mathrm{Q}} \otimes \mathrm{I}_n$ and $\hat{Q} = \frac{1}{T_w} \sum_{t=Tf_1}^{Tf_2} x_t x_t'$,and $\hat{\Omega} = \sum_{t=Tf_1}^{Tf_2} \hat{\xi}_t \hat{\xi}_t'$ with $\hat{\xi} = x_t \otimes \hat{\varepsilon}_t$.

Figure 1.8: Time-Varying Heteroscedastic Granger Causality Test

Critical Values are derived from the 95% percentile of the SW statistic on a bootstrapped sample of the VAR

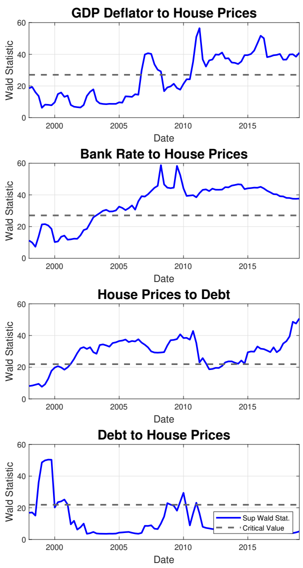

1.5.1 Results the Time-Varying Granger-Causality Test

In Figure 1.8 I present the results from an evolving recursive heteroscedastic Wald test as in Eq. 4.4 where coefficients $\hat{A}$ are calculated from a reduced form of the IV-VAR presented above. I use 2 lags in accordance with the Bayesian criterion.

The objective of this exercise is to uncover potential structural changes involving the baseline variables. I find the predictability test in object useful as it enables further inference on the dynamic relations among variables. Some regressors are good predictor of others only for a limited period of the sample and this is not immediately evident from a whole-sample Granger test. This permits to extend the scope of this when it comes to identifying the channels of monetary policy transmission.

There is a data evidence of a debt deflation channel impacting on borrowers. House prices and real household debts are well predicted by the GDP deflator in 2007-2008 and more recently. This is consistent with the view that stable and low inflation with positive GDP developments may be conducive to leverage (Borio and Lowe, 2011). The GDP deflator Granger-causes the abnormal build up of individual insolvencies from 2016, emphasising the role of the price level on household decisions.

Bank rate Granger-causes house prices almost across the entire sample period and it is a useful predictor of the debt stock from September 1999 to December 2010. Practically, the Bank rate ceases to be relevant to debt once that has reached its peak in late 2010. There is a clear change in the volatility of the SW of interest rate to insolvencies when the rate approaches zero lower bound. The bank rate bears no impact on insolvencies throughout the last decade but it predicts them during the first part of the sample. The Bank rate Granger-causes insolvencies intermittently for two years in 2004-2006, when there are numerous tightening episodes. This seems to suggest that a monetary policy tightening has on defaults a different effect than an easing. An interesting expansion of the present work could be using non-linear VAR models to account for potential differences in how insolvencies respond to either tight or easy monetary policy.

Household debt is endogenous to house prices, supporting the notion of co-movement of these two variables (Borio, 2014). It is interesting that the causal relation goes from house prices to real debt and not vice versa, this reinforces the understanding of Broadbent (2019) of house prices driving the credit expansion. Inversion of that relation follows on periods of house market decline, maybe due to debt overhang dynamics, e.g. around the financial crisis.

1.6 Conclusions

I present an IV-VAR model with household insolvencies showing that a policy-induced debt deleverage also corresponds to an increase in default levels. This finding is new as insolvencies have not been taken into account by previous papers investigating debt reduction and monetary policy in other countries. I find that households' credit quality acts as a transmission mechanism for monetary policy by deteriorating fast in response to a contractionary monetary shock. This view is consistent with financial frictions DSGE models such as those featuring the '*financial accelerator*'.

This paper has policy implications for both monetary policy and financial stability. Monetary authorities may wish to steer rates attentively in presence of highly levered households. Asset prices rallies and increase in debt in a benign environment can be quickly reversed by a rate hike. Thus there appears to be trade-offs between inflation and household conditions.

Optimal monetary policy and welfare implications of different policy rules in presence of insolvencies and high household debt are outside of the scope of this empirical paper, but their investigation in a canonical DSGE setting represents an interesting and relevant research program for financial stabilisation. Such research could build on the stylised facts regarding deleveraging and default here presented.

Central Banks that deviate from pure inflation targeting to factor in financial stability will wish to be careful that the policy rule is effective. Trying to trigger a debt reduction with monetary policy instruments might be detrimental to households and therefore not achieve its intended objective, adding to imbalances instead of steering the economy clear of a recession.

APPENDIX A

APPENDIX TO CHAPTER 1

A.1 Data Sources

Table A.1: Data Sources

Time Series that weren't originally adjusted have been seasonally transformed and deflated

	Variable	Datastream Ticker	Source	Deflated	Deseasoned
In the VAR					
a	Bank Rate	UKPRATE	BoE	NA	NA
b	GDP Chained Volume	UKGDP...D	ONS	NA	Y
c	GDP Deflator	NA	g/b	NA	NA
d	Average House Price	UKNWALLP	Nationwide	N	N
e	Household Debt	UKNIWKQ.A	ONS	N	N
f	Individual Insolvencies	UKAIHK..P	Insolvency Service	N	N
GDP					
g	GDP at Market Prices	UKGDP...B	ONS	NA	Y
Additional Variables					
h	Gross Disposable Income	UKPERDISD	ONS	N	Y
i	Annualised Income	NA	Four Quarters Rolling Sum of i	NA	NA
j	Unemployment	UKUN%O16Q	ONS	NA	Y
k	Debt to Income	NA	e/i	NA	NA
l	Debt to GDP	NA	e/Four Quarters Rolling Sum of b	NA	NA
m	Mortgage Rate	NA	BoE [1]	NA	NA
n	Corporate Rate	NA	BoE [16]	NA	NA
o	Household Final Consumption Expenditure	NA	ONS	Y	Y
p	Total Durable Goods	NA	ONS	Y	Y
q	Total Non Durable Goods	NA	ONS	Y	Y
r	FTSE 350 Index	FTSE350	Refinitiv	NA	NA
Z	Instrument	NA	Own Calculations	NA	NA

[1] Mortgage Rate up to Q4 2016 from 'A millennium of macroeconomic data' dataset, then extrapolated from 'UK Secured Loans,New Advances, Floating Rate' (DS Ticker UKZ6JT..R.). Corporate Rate up to Q4 2016 from 'A millennium of macroeconomic data' dataset, then extrapolated from UK corporate benchmark yields across all maturities and ratings (DS Series TRBC).

Figure A.1: Baseline Model Time Series and Debt Ratios

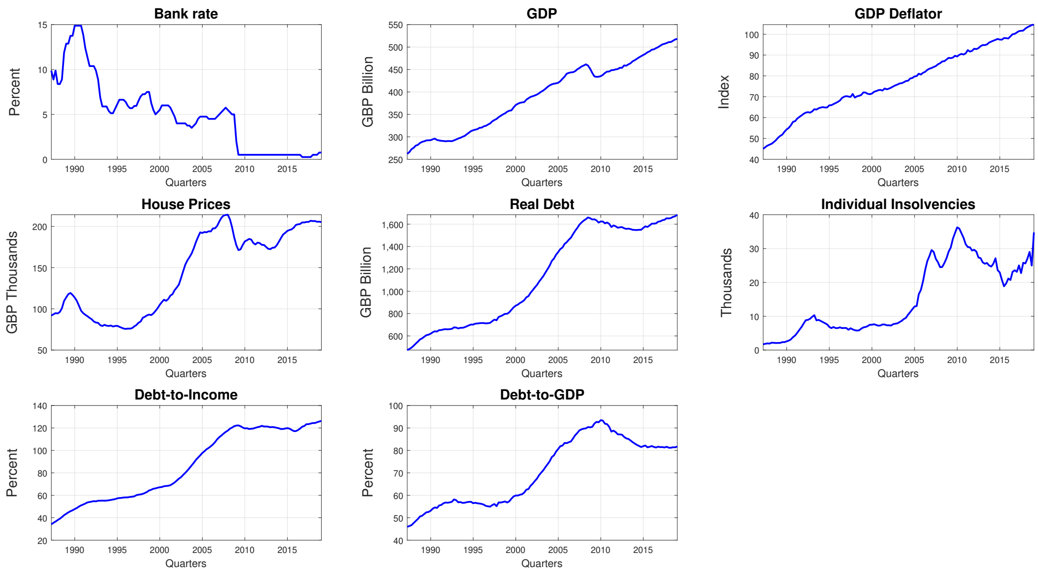

A.2 Summary of Cited Studies

Table A.2: Cited Studies

Authors	Country	Time Sample	Research Question/Goal	Identification	Regressors
Hofmann and Peersman (2017)	Across Panel	1985 - 2008	Is there a debt service channel of monetary transmission?	Recursive	Real GDP, GDP Deflator, Real House Prices, Policy Rate, Real Private Credit and a Debt-Servicing Ratio.
Laseen and Strid (2018)	Sweden	1995 - 2013	[to investigate] the relation between the shorter-term dynamics of debt and the effects of monetary policy on debt	Recursive	Trade-weighted Foreign GDP, Foreign CPIF, Foreign Short Term Rate, Repo Rate, Domestic Real GDP, Domestic CPIF, House Prices and Real Debts.
Robstad (2018)	Norway	1994 - 2013	to quantify the effect of a monetary policy shock on household credit and house prices in Norway	Recursive	GDP, CPI-ATE, Policy Rate, FX Rate, House Prices and Real Household Credit.
Mountford (2005)	UK	1974 - 2005	to investigate the effects of UK monetary policy [shocks]	Sign Restriction	GDP, Bank Rate, M0, FX rate, GDP Deflator, Oil Price.
Ellis et al. (2014)	UK	1975 - 2005	to investigate changes in the transmission mechanism of economic shocks in the UK	Sign Restriction	Bank Rate and 2 Factors from 350 UK data series.
Cloyne and Hürtgen (2016)	UK	1975 - 2007	estimat[ing] the effects of monetary policy	Narrative	Industrial Production, RPIX, Commodity Prices, Narrative Surprises.
Gerko and Rey (2017)	US and UK	1982 - 2015	How does [the importance of financial markets] affect the effectiveness of monetary policy?	Instrumental	5yr GILT Rate, RPIX, Industrial Production, Corporate Spread, Mortgage Spread, VIX, FX Rate.
Cesa-Bianchi et al. (2020)	UK	1992 - 2015	how monetary policy transmits to the broader economy	Instrumental	1yr GILT rate, CPI, unemployment rate, FX rate, mortgage spread, corporate bond spread

A.3 Exogenous Instruments

A.3.1 Instrumental Variable Sources

- Bank of England Directory of MPC Minutes: https://www.bankofengland.co.uk/sitemap/minutes;
- Bank of England Directory of Inflation Report publications: https://www.bankofengland.co.uk/sitemap/inflation-report;
- Bank of England 'A millennium of macroeconomic data' dataset: https://www.bankofengland.co.uk/statistics/research-datasets;
- Monetary Policy Committee Voting History Spreadsheet: https://www.bankofengland.co.uk/monetary-policy2.

Figure A.2: UK Money Surprises Instruments

Monthly and Quarterly Surprises proxies from Gerko and Rey (2017) and Cesa-Bianchi et al. (2020). Pearson correlation coefficient with my surprises is overlay-ed.

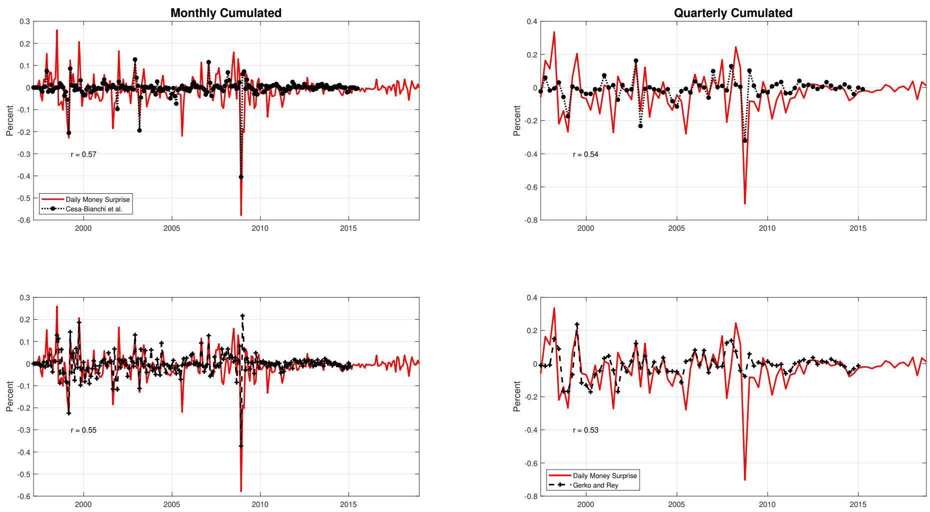

Table A.3: This Paper First Stage Regression Results

R-Squared = 0.32; F-statistic vs. Constant Model = 41.56

	Coefficient	Standard Error	t-Stat	p-Value
Intercept	0.05	0.03	1.42	0.16
Money Surprises (Instrument)	1.65	0.26	4.45	0.00

Figure A.3: Baseline Model Interest Rate Equation Residuals and Instrumental Surprises

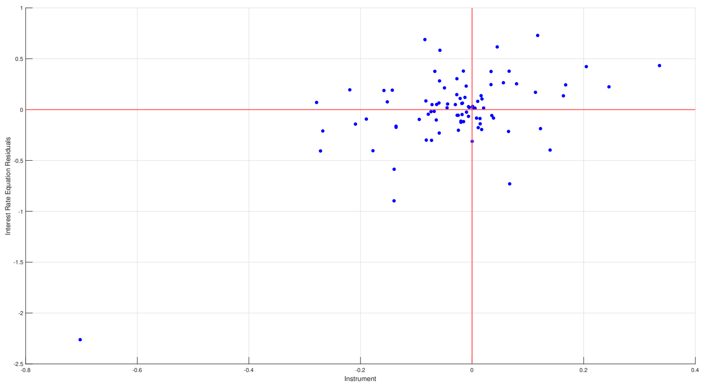

A.4 Alternative Specifications

A.4.1 Cholesky SVAR

Figure A.4: Structural Impulse Responses of a Cholesky SVAR

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

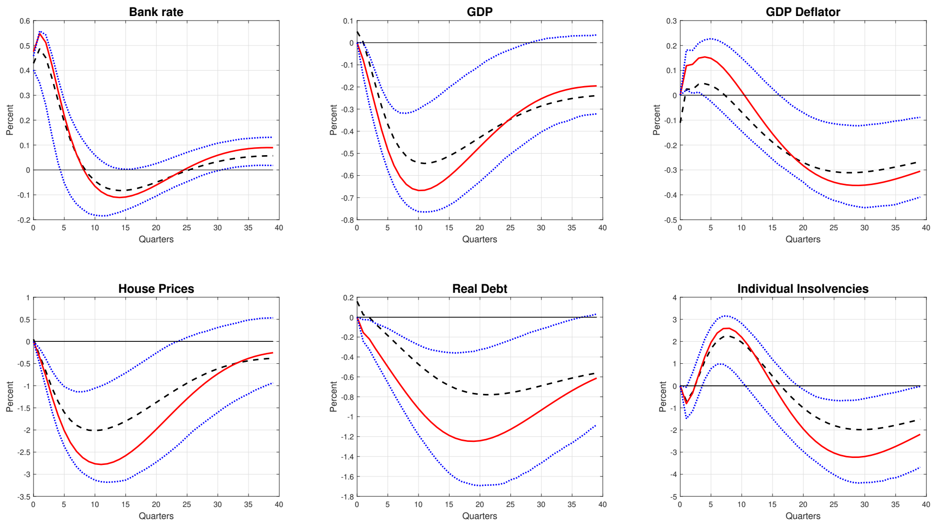

Figure A.5: Forecast Error Variance Decomposition of a Cholesky SVAR

If analysed with a forecast error variance decomposition, in the Cholesky setting, my findings are different from Mountford's as interest rate explains at the least 30% of variation of a GDP shock after 40 quarters and 70% of its own variation, thereby not 'leaning into the wind'.

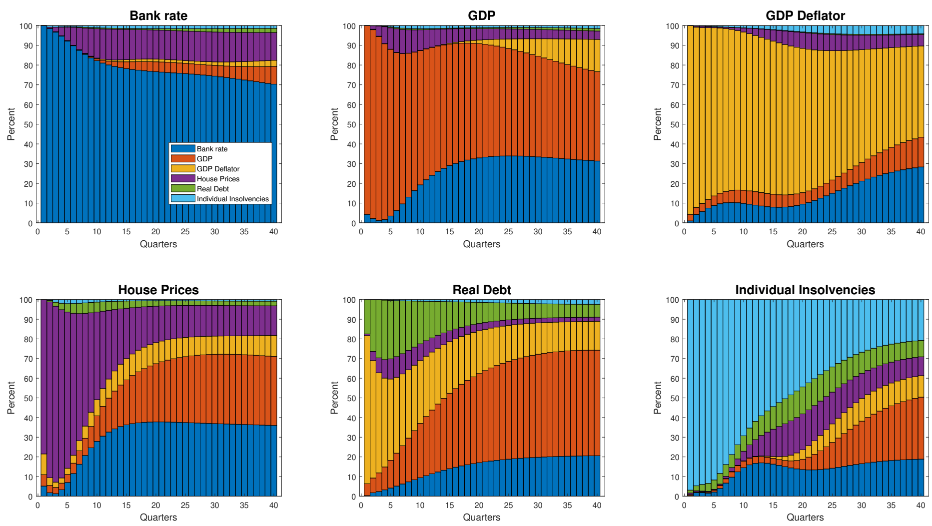

A.4.2 Non-Stationarity Robust

Figure A.6: Structural Impulse Responses of a IV-VAR(2) model computed with Cheng et al. (2019) GMM estimator and consistent covariance in case of non-stationarity.

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

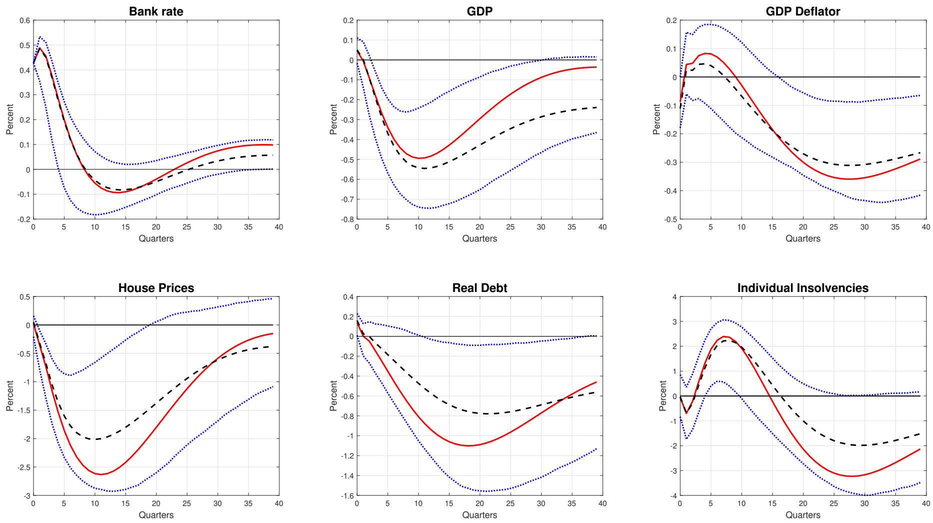

A.4.3 With Consumption Aggregates

Figure A.7: Structural Impulse Responses of a IV-VAR(2) with Household Total Final Expenditure

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

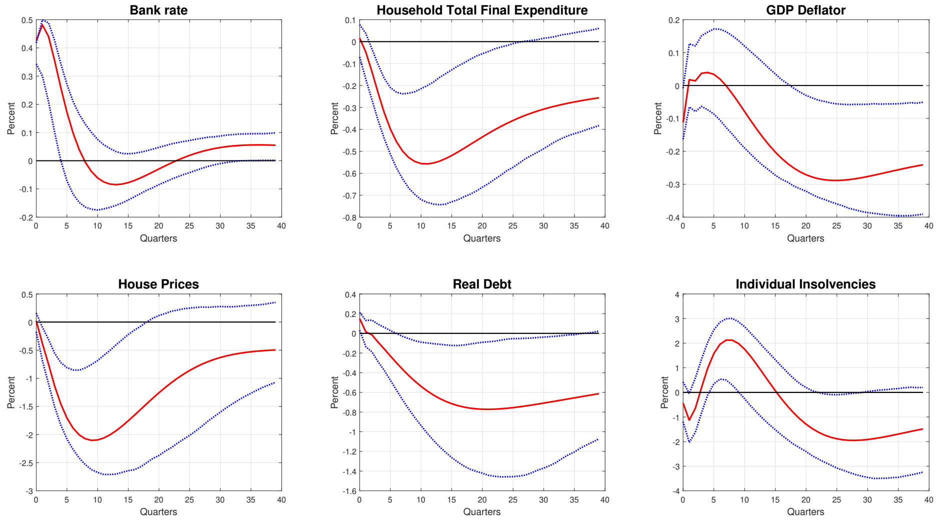

Figure A.8: Structural Impulse Responses of a IV-VAR(2) with Household Durable Consumption

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

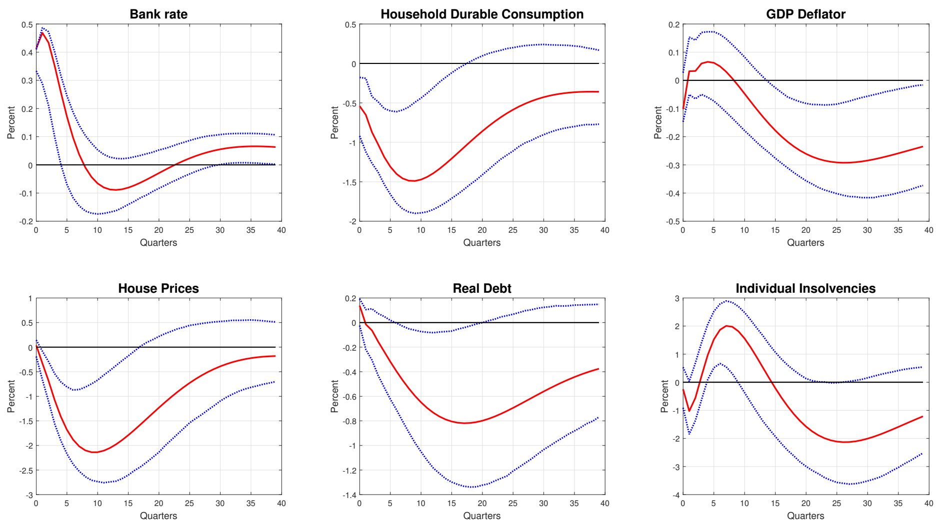

Figure A.9: Structural Impulse Responses of a IV-VAR(2) with Household Non-Durable Consumption

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

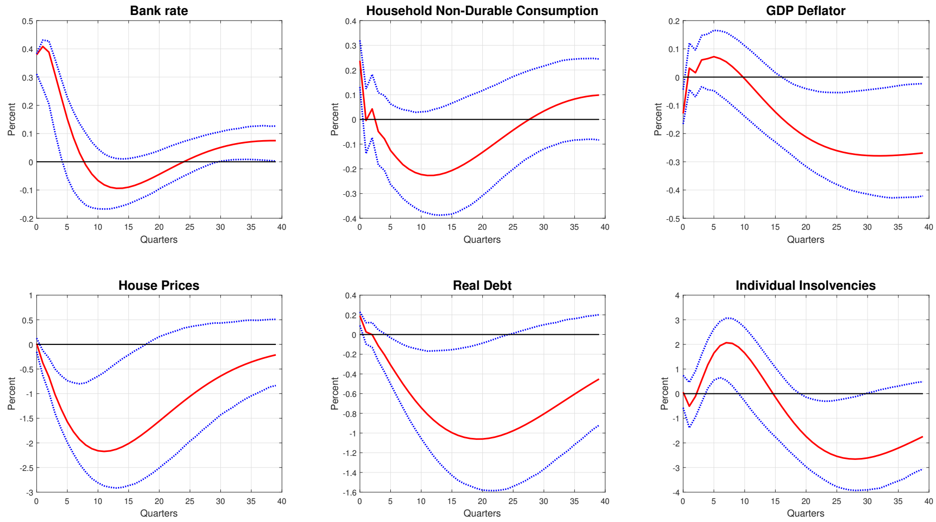

A.4.4 Debt-to-GDP

Figure A.10: Structural Impulse Responses of a IV-VAR(2) model with Debt-to-GDP

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

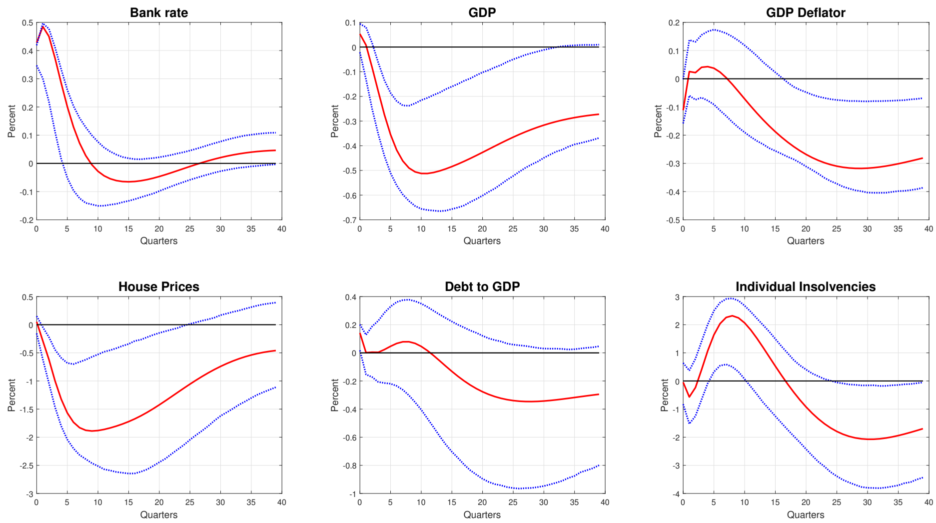

A.4.5 Debt-to-Income

Figure A.11: Structural Impulse Responses of a IV-VAR(2) model with Debt-to-Income

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

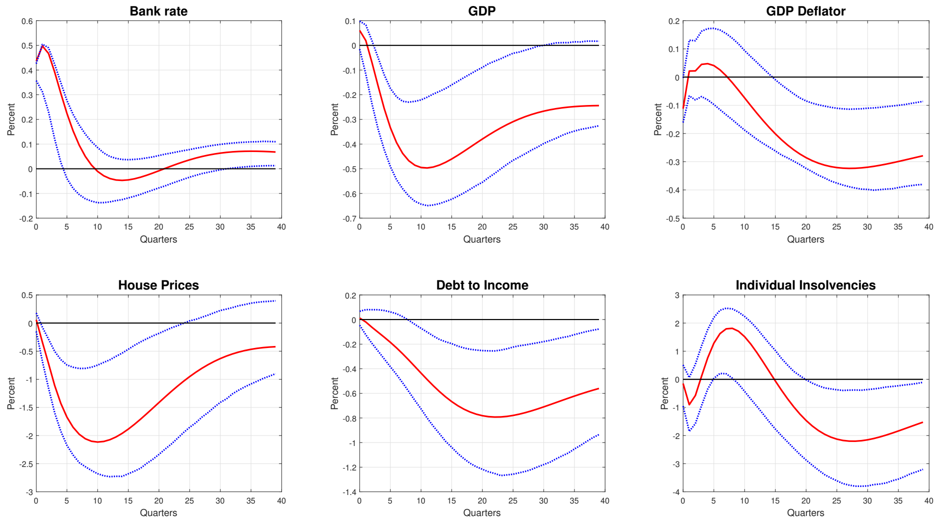

A.4.6 With Lending Rates

Figure A.12: Structural Impulse Responses of a IV-VAR(2) model with Credit Spreads

Solid line represents point-estimates. 90% confidence bands (dotted lines) are obtained simulating artificial data and re-sampling the residuals 5,000 times. The dashed line shows the baseline model IRFs

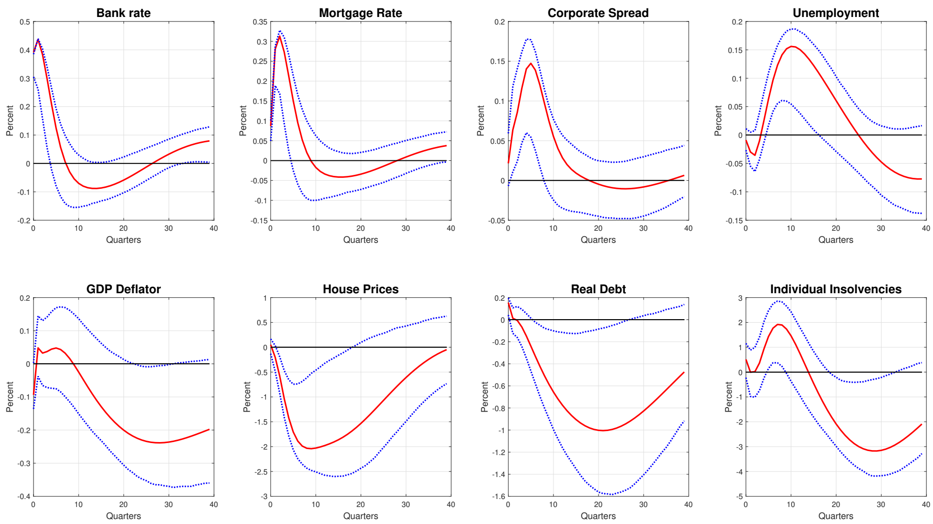

Chapter Two

Leaning Against the Wind and the Default Channel of Monetary Policy

Credit-market conditions - sharp increases in insolvencies and bankruptcies, rising real debt burdens, collapsing asset prices, and bank failures – are [. . .] themselves a major factor depressing economic activity

Bernanke, Gertler, and Gilchrist

2.1 Introduction

An open question in macroeconomics is whether the Central Bank should deviate from inflation targeting to devote particular attention to credit variables. The case for mixed monetary policy rules becomes periodically more salient, when swings in the credit cycle produce large deviations of lending from its baseline. As this happens, debt-to-income ratio tends to be elevated and assets overvalued. This situation manifests until the reversal of the credit cycle brings about a wave of defaults and an economy-wide lending contraction.

The key empirical stylised facts of the credit cycle are that a) it moves at a lower frequency than the business cycle (Aikman et al., 2015) and b) correlates with lending conditions, credit spreads and loan approvals are strongly pro-cyclical and c) a reversal of the credit cycle brings about defaults.[1] Moreover, the credit cycle interactions with the business cycle add to its persistence. Low price growth and low interest rates may facilitate debt accumulation and in turn reinforce booming asset prices in the build-up phase, making the default phase more pronounced. This circularity between interest rates, asset prices and credit quality is associated with weaker financial stability (Borio and Lowe, 2011).[2] Almost always the build-up of the credit cycle correlates with heightened systemic risk and its burst coincides with a recession.

Facts (a) and (b) mentioned above pertain to the intensive margin of lending - debt/deleverage decisions of households and banks - while fact (c) is on the extensive margin - if the lending relationship survives the economic strain. The model I present below is connected to the first chapter of my thesis, which complements the empirical literature on the intensive margin (borrowers deleverage), with evidence on the extensive margin (defaults).

In this paper I propose a simple New-Keynesian model geared to capture the stylised facts of a realistic credit market. Its primary purpose is to offer a modelling environment to tackle monetary and macroprudential policy experiments with a set of micro-founded frictions. Secondly, it constitutes a methodological contribution as I consolidate the key features of different models in the literature into a single one.

The circular nature of the relationship between monetary policy, asset price movements and the dynamics of household debt constitutes the key reinforcing mechanism of the lower frequency credit cycle. These interactions have been long overlooked in many papers on conventional monetary policy. However, the recent emphasis of policymakers on counter-cyclical monetary policy, namely leaning against the wind (LATW), has made it more important than ever to incorporate a more realistic credit dimension in models of monetary policy.

The motivation behind this work is that the benefits of LATW are not clear. An important critique levelled to New-Keynesian models advocating for mixed policy rules is that they rely on counterfactual modelling assumptions (Svensson, 2014, 2017; Laseen and Strid, 2018). In other words, real life debt-deleverage dynamics

[1] See Appendix B.1.

[2] In my paper 'Deleverage and Defaults' I offered a causal analysis finding that in the UK GDP deflator Granger-caused house prices and debt in the years after the Great Recession.

depend on a range of core assumptions that are often subjective. Is debt fixed or floating? Has it a short maturity or a longer term? Is it anchored to inflation or payable in real terms? This lies at the core of the Svensson's critique: the degrees of freedom are so many that is always possible to build a Dynamic Stochastic General Equilibrium Model (henceforth 'DSGE') model where LATW is optimal.

A second and more fundamental critique is that DSGE papers in the LATW research program have almost all abstracted from default, ignoring that financial distress accompanies debt/deleverage phenomena (i.e. the extensive margin of a credit relation).[3]

Since capturing debt/deleverage dynamics in a theoretical model requires taking a stand on a number of variables, there are many empirical papers (Hofmann and Peersman (2017); Robstad (2018); Laseen and Strid (2018)) that keep their assumptions to a minimal level and try to stay agnostic on the exact nature of incentives in the credit market. Two key observations stand out from reviewing them. The bulk of empirical literature uses VAR models to express the view that a (conventional) monetary policy surprise could be used to dampen the credit cycle by i) cooling off house prices and ii) reduce the overall household leverage. In this paper, I aim to show that there's a third unsought result to such a stance: iii) increase in the individual insolvency level (shown in Fig. 2.1).

The model I present below comprises credit constrained mortgagors (households), depositors, and banks intermediating credit between the two agents and facing a regulatory capital constraint. Accounting for two levels of frictions (on both the demand and supply of loans), allows for a richer calibration of credit aggregates and a more realistic set-up for policy experiments. The key areas of focus for our analysis arising from the empirical literature are the following: (1) credit quality represents an important channel of monetary policy transmission; (2) using rate setting to target credit variables might not be as innocuous as hinted elsewhere in the literature.

As central banks are facing a path of monetary policy tightening, it is very important to understand what that would imply for households that are coming from a decade of record low interest rates and having endured the Covid pandemic by taking up further leverage.

[3]An exception is: Christiano et al. (2014) model allows for an estimation of a time-varying steady state default rate, recognising that borrower's credit quality moves with the business cycle.

In Section 2.2 I will outline a simple model with the housing market and a constrained household sector. In Section 3.1 the model will be calibrated to match UK credit ratios and I will show impulse responses to a technology shock and a monetary policy shock. Then, I offer a discussion on Section 2.5.

Figure 2.1: Impulse Responses of an IV-VAR for UK

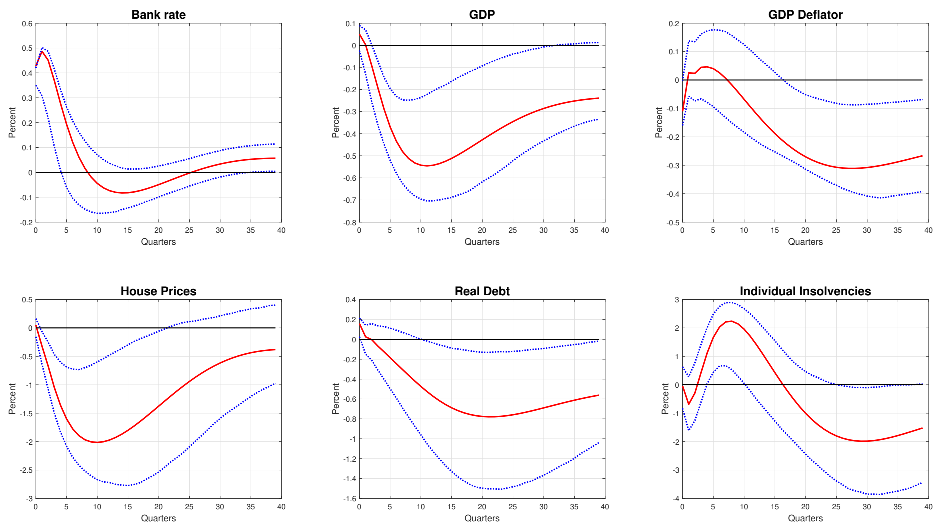

2.1.1 Leaning into the Credit Cycle and New Keynesian Modelling

LATW is shorthand for counter-cyclical conventional monetary policy and refers to the deployment of interest rate rules inversely responding to house prices or credit aggregates. The metaphor was coined by FED Chairman (1951-1970) William McChesney Martin to refer to the necessity of leaning against inflation (Romer and Romer, 2004a), but the concept has evolved to refer more generally to an activist management of money aggregates by the Central Bank (Friedman and Schwartz, 1963; Sargent, 1979). In the current sense, LATW is almost exclusively used to indicate a policy stance that factors the credit cycle in a standard Taylor rule.

The case for central banks to target asset prices is therefore neither a recent nor a settled debate. It periodically rose to prominence around cycle reversals, as for instance in the case of the Greenspan put, a backstop policy to arrest equity markets (late '80s). The FED failure to react to the dot-com bubble re-ignited the discussion around asset price targeting (Blinder and Reis, 2005). The LATW debate gained momentum again after the Great Financial Crisis in 2008-9 (GFC), this time with a focus on housing prices, in the wake of the real-estate and financial bubble collapse.

And indeed the reversal of the credit cycle has been characterised as the burst of an asset bubble: a negative shock causes a supply-sided reduction in lending, which causes in turn an economy-wide spending contraction (Carney, 2020).[4]

The focal point of the LATW lies in the premise that reacting to asset price can depress economic activity immediately but inaction risks to cause a worse recession later when a tightening is warranted by inflation. Therefore policy action involves making an informed decision about the timing of the effects of the policy (even more relevant in regards of the potential role to be played by expectations formation). It is important to note that this literature challenges the pre-crisis consensus as early papers with financial frictions do not find a case to deviate from inflation targeting Bernanke and Gertler (1995); Iacoviello (2005)

LATW proponents see value in trying to preempt a systemic crisis by reacting to asset prices. For instance, the Swedish Riksbank targeted the debt-to-income ratio with the interest rate to deflate a then ongoing housing bubble (2011-2014) Riksbank (2014), a move that was harshly contested at the time. During the past decade, the Bundesbank has also been critical of the ECB low interest rate policy as it claims the policy created asset bubbles in Germany.[5]

Among others, the Bank for International Settlements (BIS) has been an early proponent of using interest rates to curb the credit cycle (Borio and Lowe, 2011). BIS papers are empirical (Borio and Lowe, 2011; Borio, 2014), whereas the DSGE modelling effort around LATW stems from the GFC.

Modelling a credit channel in a New-Keynesian setting involves allowing for some degree of heterogeneity among representative agents. There has to be a class of lenders and a class of borrowers, secondly, there needs to be a type of micro-founded friction that affects market completeness. Without frictions, lenders and borrowers would be able to achieve on the credit market their preferred optimal allocation and the mere existence of a credit market will not affect the model dynamics.

Financial frictions are introduced in DSGE models in two ways: by introducing an

[4] The former governor of Bank of England said: "*Borrowers reduce spending, or in extremis, default. These responses make the economic downturn much deeper and more prolonged*", in a sentence similar to the quote at the beginning of this paper.

[5] FT Article

informational asymmetry (Bernanke et al., 1999) or by incorporating a collateral constraint Kiyotaki and Moore (1997). The former modelling approach stems from the Townsend (1979) costly state verification ('CSV') partial equilibrium model, later factored in an array of general equilibrium models, as Bernanke and Gertler (1989) and Carlstrom and Fuerst (1997). CSV postulates that the borrowers' (firms) default probability increases in their leverage ratio. Upon default, the lender has to pay a proportional cost to audit the firm and know the ex-post value of firm's assets, which he repossess. Hence, the rate of return on risky lending is a function of the firm's degree of leverage and will deviate from the economy riskless rate. In this paper I implement a CSV friction (Section 2.2.2). I choose this friction to link up borrowers' housing wealth to the business cycle, so obtaining a pro-cyclical default rate.

An alternative form of introducing frictions, the collateral constraint methodology emphasises the limited enforceability of credit contracts. If a debtor decides not to honor his loan, the creditor cannot force him to pay, instead, he has to rely on collateral and hence 'hard' assets as land or real estate can serve the dual role of good and collateral to loans.[6] Here credit frictions limit the purchase of capital, reducing overall investment and output in later periods following a negative shock. Iacoviello (2005) extends this framework to a new-Keynesian DSGE model with entrepreneurs pledging real estate to back their loans up to an exogenous loan-to-value (LTV) ratio.

Iacoviello (2005), in the simplest declination of his model, presents financially constrained firms borrowing from patient households. Firms use real estate in their production function and households consume housing value and a consumption good, while demanding real money balances. Iacoviello and Neri (2010) extend that framework to a Bayesian estimation aimed to match US housing market data. This enriched model features a separate production function for housing, with land and labour as inputs together with entrepreneurial capital.

Guerrieri and Iacoviello (2017) introduce the modification that a collateral constraint might be slack, thus allowing households to self-insure in benign economic periods. Similarly, Bluwstein et al. (2018) analyse credit frictions with an occasionally binding constraint.

[6] Liu et al. (2013) provide a good example of a DSGE model in which land price drives investment fluctuations, as real estate is a common collateral to business loans.

Building on the collateral constraint literature, Monacelli (2009) shows in a DSGE setting that credit constrained households can shed light on the transmission mechanism as a tightening prompts borrowers to substitute durable goods with non-durable ones. This explains the conventional VAR impulse responses to a monetary policy shock: household debt persistently declines and there follows a collapse in non-durable consumption larger than that in durables, making a case for models with financial frictions.

A more complex model allowing for credit intermediation is the Gerali et al. (2010), which is built on Iacoviello (2005). Here, banks collect deposits from households, accumulate profits into their equity and lend to firms at a rate higher than the policy one due to the cost they face to deviate from an exogenous capital adequacy ratio and to manage their capital position. Firms are collaterally constrained when they borrow to purchase capital.

Gambacorta and Signoretti (2014) use this set-up to test different monetary policy rules testing the gains of 'leaning against the wind' stance. They find it useful especially in the high leverage calibrations of their model, where it is apposite to reduce the level of firm debt and the price of capital. On the other hand, Gelain et al. (2018) show that when mortgages are longer in maturity, debt-to-GDP targeting is less effective.

Financial frictions have come to be seen as a main channel of shock propagation in Christiano et al. (2014), it is shown that shocks to the cross-sectional variance of ex-post capital realisation have the potential to explain US recent history business cycle fluctuations.

Matters get thornier when banking intermediaries are introduced in the model. Clerc et al. (2015) propose a model with banks allocating loans to mortgages and business loans, these three agents raise finance in an imperfect market and they are subject to default risk. This model re-affirms the centrality of constraints on banking intermediaries' capital as excessive leverage translates into an excessive volatility of aggregates whereas too tight capital requirements constrain credit too much.

As far as the empirical performance is concerned, Brzoza-Brzezina et al. (2013) have directly compared the CVS framework with the collateral constraint one, making the necessary tweaks to work with two identical models but for the way financial frictions are introduced. They conclude that the CVS methodology produces impulse response functions that are closer to a VAR model, whereas under the collateral constraint approach, the responses are more front-loaded and short-lived than in the case of the empirical model.

I draw from the CSV literature by modelling an idiosyncratic shock that affects the ex-post realisation of housing projects. This is the framework introduced in Bernanke et al. (1999) but applied to housing as in Forlati and Lambertini (2011) and Lambertini et al. (2017b). The structure of the economy follows Iacoviello (2005), but in my model households are the borrowing party and discount future utility more. The aggregate supply is fairly standard, with final and intermediate goods producers. The difference here is that intermediate entrepreneurs smooth their consumption by investing in 1-period sight deposits at banks. Banks have to comply with an exogenous capital adequacy ratio or else they have to pay a cost to deviate from it.

The new insight that the model below offers is that by modelling a richer set of transmission that inform debt and deleverage decisions. Collateralised borrowing will have effects on consumption smoothing as asset prices enter the borrowing constraint. In the absence of these micro-foundations, the whole 'credit view' of the monetary transmission mechanism would be absent.

2.2 The Model

2.2.1 Model Agents

To capture the key features of the credit market, I resort to a model with heterogeneous agents, whom are divided in hand-to-mouth borrowers, entrepreneurial savers and banking intermediaries. To summarise the key relations among the agents presented above, it is useful to picture the credit market in a static fashion in the following sectoral T-accounts:

Households		Banks		Entrepreneurs	
Wages	Loans	Loans	Capital	Deposits	Wages
Housing	Net-Worth		Deposits	Housing	Net-Worth
Profits					

The agents populating the model are households, banks and entrepreneurs. Households and entrepreneurs trade housing in a perfectly competitive market. The bank is forced to hold a certain fraction of loans as regulatory capital. Net-worth is defined as the difference between assets and liabilities (for banks it is defined as banking capital). The above scheme also tracks the different type of assets available to agents, noting the fact that the economy is cashless, the entrepreneurs save with return-yielding deposits and households only hold dwellings.

The presence of a capital constrained bank introduces a wedge in the residential market. Deposits and loans are lower than in a steady state without capital requirements exactly because $D \neq L$. Thus in steady state, both households and entrepreneurs consume less than what they would without banking regulation.

The second wedge stems from the risky mortgage friction. The presence of a dead-weight monitoring cost implies a foreclosure cost that is proportional to housing value, thus an increase in defaults causes a destruction of existing residential real estate.

Dynamically, shocks affect households' and entrepreneurs' net worth, causing a re-allocation of housing value between them. For example, a monetary policy shock affects inter alia the cost of borrowing and the demand for housing, causing a de-accumulation of residential housing and an accumulation of commercial housing. A technological shock has the opposite effect, with different persistence due to the auto-regressive parameter.

The construction sector is anti-cyclical. New real estate is built by the entrepreneur, so this agent is able to modulate construction when his demand is high.[7] Furthermore, since monitoring costs are modelled as a dead-weight loss of housing value, new constructions increase on impact of an adverse shock as agents try to replenish the housing lost on the back of foreclosures.

2.2.2 Risky Mortgages and Collateral Constraint

In this section I will illustrate the derivation of the collateral constraint that will feature prominently in the Household Problem (Section 2.2.3) presented below. The key feature of Lambertini et al. (2017b) risky mortgages friction set-up is the addition of a constraint that impedes consumption smoothing and makes

[7] This is similar to Iacoviello (2005), who does not feature a construction sector but has a similar re-allocation dynamics due to fixed supply of housing.

loan supply dependent on an endogenous loan-to-value ratio that reflects lending conditions.

All housing projects are equal ex-ante, but then they are individually scaled by an idiosyncratic shock ω_{t+1}. The Collateral Constraint postulates that the payoff of 1-period adjustable rate mortgage equals the expected value of the collateral $E_t[H^h_{t+1}P^h_{t+1}]$ upon the realisation of an idiosyncratic shock $\hat{\omega}_{t+1}$. After the realisation of the individual shock, if the shock decreases the housing value below a certain endogenous threshold $\hat{\omega}$, then the borrower defaults. The lender has to pay a monitoring cost μ to observe the realisation of the idiosyncratic shock. After the shock, the lender seizes the residual value of the housing stock.

$$L_{t+1}(1+R^z_{t+1}) = \hat{\omega}_{t+1}(1-\delta^h)H^h_{t+1}P^h_{t+1} \qquad \text{(Default Threshold)}$$

The Default Threshold identifies the equilibrium $\hat{\omega}$ as the threshold that equates the risky loan returns (the left-hand side) with the housing value scaled by the external shock (right-hand side). This equality can intuitively represent a negative equity condition, so that if the value of the house falls below the value of the mortgage payment, then the borrower will default. δ^h is the calibrated depreciation rate of the housing stock.

The zero profit condition below equates the gross riskless rate on the loan portfolio to the expected risky returns in case of regular mortgage payment plus the expected collateral realisation upon default net of monitoring costs. Hence, the above Default Threshold is state contingent as $1+R^z_{t+1}$ adjust to make the below Riskless Rate true. The average rate of return $(1+R_{L,t})$ is pre-determined and payable on loans taken at time t that become due at time t+1.

$$(1+R_{L,t})L_{t+1} = \int_{\hat{\omega}_{t+1}}^{\infty}(1+R^z_{t+1})L_{t+1}f_{t+1}(\omega)d\omega + (1-\mu)(1-\delta^h)H^h_{t+1}P^h_{t+1}\int_0^{\hat{\omega}_{t+1}}\omega f_{t+1}(\omega)\,d\omega \qquad \text{(Riskless Rate)}$$

Where $ft+1(\omega)$ is the probability density function of the random variable ω.

Substituting $(1+R_{L,t})L_t$ in the Riskless Rate using the definition in Default Threshold yields the collateral constraint, which forces loans to be equal to a certain fraction $\Phi[\hat{\omega}]$ of expected housing value. Hence $\Phi[\hat{\omega}]$ is the expected share of housing value net of monitoring costs going to the lender and thus the fraction of housing value that the lender is willing to finance. It follows that in this set-up the Lender's share can be intuitively interpreted as the loan-to-value ratio applied by lenders.

$$L_{t+1}\left(1+R_{L,t}\right)=(1-\delta^{h})E_t[\Phi\left[\hat{\omega}_{t+1}\right]H_{t+1}P^{h}_{t+1}] \qquad \text{(Collateral Constraint)}$$

Adding the Borrowing Constraint to the household problem is similar in spirit to the optimal contracting problem set out in Carlstrom and Fuerst (1997) and Bernanke et al. (1999), where the borrower maximises his share subject to the lender's share. Here the household problem is formulated to allow households to maximise their utility subject to lender's zero profit condition besides the usual budget constraint. In this setting it is natural to interpret the lender's share as the dynamic LTV ratio ($\Phi[\hat{\omega}]$ in the equation below).

$$\Phi[\hat{\omega}_{t+1}]=\hat{\omega}\underbrace{\int_{\hat{\omega}_{t+1}}^{\infty}f(\omega_{t+1})\,d\omega}_{(1-F[\hat{\omega}_{t+1}])}+(1-\mu)\underbrace{\int_{0}^{\hat{\omega}_{t+1}}\omega f(\omega_{t+1})\,d\omega}_{G[\hat{\omega}_{t+1}]} \qquad \text{(Loan-to-Value)}$$

The presence of monitoring costs proportional to the housing value introduces a dead-weight cost associated with defaulting.

The other useful definitions for the following section are: the probability of default $F[\hat{\omega}]$ and the borrower's share of housing value conditional on default $G[\hat{\omega}]$, i.e. the value of housing conditional on the shock being less than the cutoff $\hat{\omega}$.

2.2.3 Household Problem

The model economy is populated by infinite living households. Households belongs to a risk-sharing family. Households are identical before facing an idiosyncratic shock to their housing value. After the realisation of such shock, some households decide to default and some not, according to the Default Threshold defined above.

In this configuration, the family head maximises an utilitarian welfare function that can be conceived as the aggregation across all households and thus is representative of the average household. Housing delivers utility prior to the idiosyncratic shocks and all households have the same ex-ante utility.

For the formal maximisation problem I assume that the family head maximises average utility of all households as:

$$E_0 \sum_{t=0}^{\infty} \beta^t U \left(C_t^h, N_t, H_{t+1}^h \right) \tag{2.1}$$

The period utility takes the functional form:

$$U \left(C_t^h, N_t, H_{t+1}^h \right) = \frac{(C_t^h)^{1-\sigma}}{1-\sigma} + \frac{\varsigma H_{t+1}^{h,1-\chi}}{1-\chi} - [N_t^{c(1+z)} + N_t^{h(1+z)}]^{\frac{1+\phi}{1+z}} \frac{u}{\phi+1} \tag{2.2}$$

The utility function means that households draw utility from a consumption good (denoted by C^h) and real housing (H^h). Working yields them a dis-utility and they can choose how many hours to work in the goods sector or in the construction sector (N^c and N^h respectively). E_0 is the expectation operator and β a discount factor. σ and χ are the coefficients of constant relative risk aversion. ς represents a measure of housing preferences.

The housing stock H^h is predetermined and household chooses it at the beginning of period t and uses it in $t+1$.

Labour demand in the two sectors is homothetic and the sectors are not-perfect substitutes. The parameter z represents the elasticity of substitution among labour types and u is the dis-utility from working.

A final caveat on the functional form of the utility function is that it does not include money, nor money are present in the rest of the model. Hence the model economy is cashless as there is no particular reason to hold money balances, this is consistent with an economy with seamless electronic payments (Woodford, 2003b).

Utility in Eq. 2.12 is maximised subject to a flow of budget constraints and collateral constraints. The flow budget constraint is expressed in real terms (already deflated by the price level) and it is:

$$C_t^h + P_t^h H_{t+1}^h + (1 - F[\hat{\omega}_t])(1 + R_t^z) L_t = L_{t+1} + N_t^c W_t^c + N_t^h W_t^h + (1 - G[\hat{\omega}_t]) P_t^h H_t^h + F_t \tag{2.3}$$

Households accumulate real estate (H^h scaled by its price P^h) which is used as dwellings and is subject to a depreciation rate δ^h. At each time period it pays the risky rate with probability $(1 - F[\overline{\omega}])$. F are the profits from owning firms. This can be intended as the sum of budget constraints of all households. It is consistent with the definitions of $F[\hat{\omega}_t]$ and $G[\hat{\omega}_t]$, which are integrals over the distribution of shock ω.

It is important to note at this point that both households and entrepreneurs (see next section) are accumulating real estate. To facilitate the comprehension of the model, we deem that the total stock of housing comprises of dwellings and offices. Households accumulate dwellings, of which services yields them utility and entrepreneurs accumulate offices, which are used to produce new real estate.

Timing Assumptions

At end of period t, the household takes a loan from the bank of value L_{t+1} in order to buy H_{t+1}^h housing at price P_t^h. At start of period $t+1$, the household gets the ω_{t+1} shock, it can then observe the value of other aggregate shocks and therefore infer the impact of the price P_{t+1}^h on housing value scaled by the idiosyncratic shock and decides accordingly whether to default on loans L_t taken in period $t-1$. Hence, there are defaults in equilibrium. Loans are payable at an adjustable rate $(1 + R_t^z)$, which is determined after the realisation of the idiosyncratic shock.

Reduced Budget and Collateral Constraints

Using the definition laid out above in Default Threshold and Loan-to-Value, Equations 2.3 simplifies to:

$$H^h_{t+1}P^h_t + C^h_t - L_{t+1} = -\Big(H^h_t P^h_t \mu \mathrm{G}\left(\hat{\omega}_t\right) + L_t\left(R_{L,t}+1\right)\Big) + (1-\delta^h)H^h_t P^h_t + N^c_t W^c_t + N^h_t W^h_t + F_t \tag{2.4}$$

Households maximise utility also subject to the following flow Collateral Constraint shown in section 2.2.2 above:

$$L_{t+1}(1+R_{L,t}) = (1-\delta^h)E_t[\Phi[\hat{\omega}_{t+1}]H^h_{t+1}P^h_{t+1}] \tag{2.5}$$

This additional constraint embeds the financial friction within the model as it postulates that households are limited in their borrowing capacity by their existing housing holdings.

2.2.3.1 First Order Conditions

The household decides hours worked, housing investment and loans. Maximising lifetime utility subject to the two constraints yields the following first order conditions (FOCs):

The labour supply schedule for construction worker is the following:

$$\frac{W^h_t}{C^c_t} = {N_t}^{h(z)}[{N^h_t}^{(1+z)} + {N^c_t}^{(1+z)}]^{\frac{\phi-z}{1+z}} \tag{2.6}$$

It is homothetic to the labour supply in the goods sector. The ratio of hours worked in the two sectors is:

$$\frac{W^c_t}{W^h_t} = \left(\frac{N^c_t}{N^h_t}\right)^z \tag{2.7}$$

Housing Demand is the FOC with respect of H_t:

$$\varsigma\beta \frac{{C^h_t}^{\sigma}}{{H^h_t}^{\chi} P^h_t} = \frac{P^h_{t-1}}{P^h_t}\left(\frac{C^h_t}{C^h_{t-1}}\right)^{\sigma} - \beta(1-\delta^h) + \beta\mu G[\omega_t] - (1-\delta^h){C^h_t}^{\sigma}\lambda^c_{t-1}\Phi[\omega_t] \tag{2.8}$$

And the Lagrange multipliers for the budget and borrowing constraint are:

$$\lambda_t = C_t^{h-\sigma} \tag{2.9}$$

$$\lambda^c{}_{t-1} = \frac{\lambda_{t-1}}{(1+R_{L,t-1})} - \beta\lambda_t \frac{(1+R_{L,t})}{(1+R_{L,t-1})} \tag{2.10}$$

From the above two FOCs stems the fact that if the households' discount rate differ from the entrepreneurs' one, then the collateral constraint is positive in steady state and it is binding. This means that households are net borrowers in steady state.

Finally, the binding collateral constraint implies that households' future consumption is linked to the period-to-period expansion of loans.

$$E_t\left\{\left(\frac{C_{t+1}^h}{C_t^h}\right)^\sigma\right\} = E_t\left\{\frac{L_{t+1}-L_t}{L_{t+1}}\frac{(1+R_{L,t+1})}{(1+R_{L,t})}\right\} \tag{2.11}$$

There's no consumption smoothing because by construction the collateral constraint is binding in the steady state and in the neighbourhood of it. So future consumption is linked to the expansion of loans. The household could in theory smooth consumption by under-borrowing but this would require setting the model such that the constraint could only occasionally be binding.

2.2.3.2 Entrepreneur's Problem

The entrepreneur only maximises lifetime utility drawn from the consumption of the consumption good subject to a flow of budget constraints. The lifetime utility is:

$$E_0 \sum_{t=0}^{\infty} \gamma^t U\left(C_t^e\right) \tag{2.12}$$

And the functional form for utility is:

$$U\left(C_t^e\right) = \frac{(C_t^e)^{1-\sigma}}{1-\sigma} \tag{2.13}$$

The entrepreneurial sector accumulates properties, which when owned by entrepreneurs is used commercially as offices (H^e). Offices are also used to produce new real estate (H^n) in a Cobb-Douglas production function. Entrepreneur's discount rate is γ.

The entrepreneur produces also a consumption good, using a production function linear in hours worked. The entrepreneur saves in banking deposits D.

$$H^e_{t+1}P^h_t + C^e_t + N^c_t W^c_t + N^h_t W^h_t + D_{t+1} = (1-\delta^h)H^e_t P^h_t + \frac{Y_t}{X_t} + (1+R^d_{t-1})D_t + H^n_t P^h_t \quad (2.14)$$

The new housing production function is:

$$H^n_t = A^h_t {H^e_{t-1}}^{(\nu)} N^{(1-\nu)}_{h,t} \quad (2.15)$$

The consumption good production function is:

$$Y_t = A^h_t N^c_t \quad (2.16)$$

From which we can infer the labour demand curves:

$$N^h_t W^h_t = (1-\nu)H^n_t P^h_t \quad (2.17)$$

$$N^c_t W^c_t = Y_t / X_t \quad (2.18)$$

From the inter-temporal utility maximisation problem stems the following Euler Equation, which determines the path of entrepreneurial consumption.

$$(1+R^d_t) = E_t\left[\frac{C^e_{t+1}}{\gamma C^e_t}\right] \quad (2.19)$$

The housing supply is:

$$\frac{P^h_t}{C^e_t} = E_t\left[\gamma\left(1-\delta^h+\nu\frac{H^n_{t+1}}{H^e_t}\right)\frac{P^h_{t+1}}{C^e_{t+1}}\right] \quad (2.20)$$

The overall real estate stock comprises both dwellings and offices:

$$H_t = H_t^e + H_t^h \tag{2.21}$$

It is a state variable which evolves in the following fashion:

$$H_t = H_t^n + (1-\delta^h)H_{t-1} \tag{2.22}$$

2.2.4 The Banking Sector and Risk Weighted Assets

A Bank matches deposits with loans and is bounded to comply with an exogenous inverse leverage ratio ξ (the capital adequacy ratio or CAR). It pays a quadratic adjustment cost when it deviates from it. The banking sector is modelled following Gambacorta and Signoretti (2014), who simplify the framework laid out in Gerali et al. (2010).

The starting point is the following identity, the loan portfolio is equal to deposits plus banking equity.

$$D_t + K_t^b = L_t \tag{2.23}$$

The banking sector holds fully diversified 1-period adjustable rate mortgage loans (ARM) on behalf of deposit-holders. The banking profit J_t^b is determined by the net interest margin plus the quadratic cost of deviating from target capital adequacy ξ. Capital accumulates with the retention of last period's earnings:

$$J_t^b = L_t\,(1+R_{L,t}) - D_t\,(1+R_{D,t}) - \frac{1}{2}\theta K_t^b \left(\frac{K_t^b}{RW_t L_t} - \xi\right)^2 \tag{2.24}$$

Maximising Banking profits subject to the bank budget identity yields the loan spread.[8]

[8]Both papers assume there is a continuum of commercial banks solving the same static optimisation problem. Imposing symmetry the first order condition is the same as postulating that there is a single bank extending loans, as I do in this paper.

$$(1+R_{L,t}) = \left(1+R_{D,t)}\right) - \frac{\theta}{RW_t}\left(\frac{K_t^b}{L_t}\right)^2\left(\frac{K_t^b}{RW_tL_t} - \xi\right) \tag{2.25}$$

Without risk-based weighting ($RW_t = 1$) the capital adequacy ratio is computed on the basis of actual assets (L). Factoring in risk-weights, the *effective* capital adequacy ratio is function of the mortgage default risk.

This is consistent with the Basel Framework Agreement, banks are bound to maintain high quality capital (shared and retained earnings) at the least above the 8% of threshold of assets on a risk-weighted basis.[9] The rationale is to provide banks with loss-absorption capacity in the event of a crisis.

Banking equity represents a state variable in the model and evolves at rate:

$$K_t^b = (1-\delta^k)K_{t-1}^b - \delta^k J_t^b \tag{2.26}$$

We can now postulate a simple risk-weighting rule, inspired by the Basel framework. The risk weight is equal to the expected credit loss, defined as the default rate reduced by 1 minus the recovery rate (i.e. the loss given default):

$$EL_t = PD_t(1-RR_t) \tag{2.27}$$

$$PD_t = F[\hat{\omega}_t] \tag{2.28}$$

$$RR_t = \int_0^{\hat{\omega}_t} \frac{(1-\mu)(1-\delta^h)H_t^hP_t^h}{F[\hat{\omega}]L_t(1+R_t^z)}\omega f(\omega)\,d\omega \tag{2.29}$$

The default rate is, as defined above, $F[\hat{\omega}_t]$, whilst also the recovery rate can be derived as the value of recovered collateral over defaulted assets (as in Candian and Dmitriev (2020)). Specifically, the last 'Recovery Rate' equation crystallises the fact that lenders recover a fraction $(1-\mu)$ of housing value upon default. Hence, the recovery rate is the expected asset value conditional on the probability of default being lower than the default threshold. Simplifying the equation using the Default Threshold expression yields the following expression, where the integral is $G[\hat{\omega}_t]$:

[9] Or, to use regulatory definitions, the total capital ratio has to be greater than 8% . See https://www.bis.org/fsi/fsisummaries/defcap_b3.htm

$$RR_t = \frac{(1-\mu)}{F[\hat{\omega}_t]\hat{\omega}_t} \int_0^{\hat{\omega}_t} \omega f(\omega)\, d\omega \tag{2.30}$$

The risk weight is normalised to be 1 in steady state, it is therefore divided by steady state expected loss ($\overline{EL}$) The parameter ρ_{rw} is a smoothing constant that captures the speed of adjustment of the risk-weight to ensure that there is not excessive volatility in the risk-weight itself.

$$RW_t = \left[PD_t\,(1-RR)\,\overline{EL}^{-1}\right]^{\rho_{rw}} \tag{2.31}$$

$$RW_t = \left[F[\hat{\omega}_t]\left(1-(1-\mu)\frac{G[\hat{\omega}_t]}{F[\hat{\omega}_t]\hat{\omega}_t}\right)\overline{EL}^{-1}\right]^{\rho_{rw}} \tag{2.32}$$

It is to be noted that the Basel Framework's rationale behind risk-weighting of assets is to ensure comparability among credit institutions and have the capital adequacy indicators reflect the inherent riskiness of the loan book. This latter motive is the dominant one: in case of losses, banking capital has to be enough as to absorb losses preventing catastrophic bank failures.

This marks a significant point of departure between the modelling device adopted above and the reality of the Basel Framework. Here the banks' credit risk is fully mitigated by the incentive compatibility constraint (stated in Riskless Rate), which allows the banks to charge a loan rate equal in expectations to the loan portfolio returns.[10]

Losses do not follow credit decisions, but stem from capital management decisions and the deviation from the regulatory CAR. Hence, the RW is an automatic macroprudential tool that limits credit expansion to dampen the risky lending financial friction. If defaults are increasing, the risk weight would also increase. This artificially raises the *effective* CAR and generates an extra cost for banks to lend, thus reducing the loan expansion.

A second point for consideration is the subtle difference between capital requirements and more active macroprudential policy. Capital ratios have to comply with the minimum regulatory standard at all times, as to mitigate systemic risk.

[10] Benes and Kumhof (2015) relax this assumption, allowing for credit losses in a Bernanke et al. accelerator framework.

However, the regulatory minimum of 8% can be at times raised stepwise by an add-on called the 'countercyclical capital buffer' (CCyB). This regulatory tool is the main policy lever in the hands of macroprudential authorities. CCyB can be increased to cool off periods of abnormal credit growth.[11]

Naturally, CCyB makes a standalone policy tool and can be modelled to obey a different rule than the above 'endogenous' risk weight. Benes and Kumhof (2015) model the countercyclical capital buffer in a model with financial frictions. Acosta-Smith et al. (2021) model two risk-weighting rules postulated in Basel III: the internal rating based (IRB) and the output floor.

2.2.5 Retailers and Sticky Prices

Sticky prices are modelled according to the popular Calvo setting.[12] I introduce them as in Bernanke et al. (1999) as a modelling device to have achieved a staggered adjustment of retail prices. At each time period, a fraction $1-\Theta$ of firms adjust their prices.

The aggregate price index P evolves according to:

$$P_t = [\Theta P_{t-1}^{1-\varepsilon} + (1-\Theta)(P_t^*)^{(1-\varepsilon)}]^{1/(1-\varepsilon)} \tag{2.33}$$

Whilst individual firm z optimises its expected profit when re-adjusting the sale price $P_t^*(z)$.

$$\sum_{k=0}^{\infty} \left[\Theta^k \left(\beta \frac{C_t}{C_{k+t}} \right) \left(\frac{P_t^*(z)\, Y_{t+k}^*(z)}{P_{k+t}} - \frac{X\, Y_{t+k}^*(z)}{X_{k+t}} \right) \right] = 0 \tag{2.34}$$

The optimality condition for P_t^* postulates that the equilibrium price equates expected marginal revenue to the expected marginal cost applying the stochastic discount rate.

Hence firms resetting their prices will do so as to charge a desired markup X_t on final goods' marginal cost and since the firms cannot reset prices at every period, they have to set the price to be consistent with future expected marginal cost. This

[11] In note 2 of the previous chapter I already commented on how CCyB encapsulates the Basel consensus as LATW is explicitly mentioned in its definition.

[12] For a textbook example see Galí (2009)

implies a steady state mark-up $X = \varepsilon/(1-\varepsilon)$ and profits rebated to households $F = (1-1/X)Y$.

The log-linearised combination of Eq. 2.33 and 2.34 is yields the textbook forward-looking New-Keynesian Phillips curve linking mark-up to inflation. The mark-up is the inverse of the marginal cost, thus firms revise their price to track the cyclicality of marginal costs (or $-X_t$ - the mark-up).

2.2.6 Monetary Policy

The monetary authority sets the interest rate according to the log-linear Taylor Rule of the type set out in Clarida et al. (1998).

$$r_t^n = (1-\rho_r)\left(\rho_l l_{t-1} + p_{t-1}^h \rho_{ph} + (1+\rho_\pi)\,\pi_{t-1} + \rho_y y_{t-1}\right) + \rho_r r_{t-1}^n \tag{2.35}$$

Where the nominal interest rate r^n responds proportionally to the deviations from inflation rate, output, loans and house prices. This rule implies a gradual adjustment (smoothing) of the interest rate at persistence given by the parameter ρ_n.

This rule encapsulates a shorter run dynamics of the nominal interest rate towards is desired level.[13]

Longer-run response coefficients are ρ_π, ρ_y and ρ_{ph}. The fact that the inflation weight $(1+\rho_\pi)$ is greater than 1 accommodates the Taylor principle, i.e. that the Central Bank needs to react to inflation sufficiently strongly to assure model determinacy.

This rule can easily nest various exceptions. For example, when setting $\rho_l = \rho_{ph} = 0$, it reduces to a standard inflation-targeting Taylor rule, which I adopt as the baseline rule. If $\rho_l > 0$ and $\rho_{ph} > 0$, then I can incorporate a degree of LATW in the model, whereby the Central Bank reacts to either the loan deviations from steady state, or to house prices. These alternative specifications with asset price targeting will be tested in Section 2.6.

There are a couple of comments to be made on the above Taylor rule and its

[13]The full-blown Taylor rule is, where the desired interest rate is r^* :

$$r_t^* = \rho_l l_{t-1} + p_{t-1}^h \rho_{ph} + (1+\rho_\pi)\,\pi_{t-1} + \rho_y y_{t-1}$$

$$r_t^n = (1-\rho_r) r_t^* + \rho_r r_t^n$$

econometric estimation. The first is that it is more usually presented in a forward-looking form: the Central Bank reacts to future inflation expectations rather than realised inflation. Same thing applies to output gap, since it is understood that the Central Bank has to factor in as certain policy lag as monetary policy does not quickly transmits to the economy at large, hence both price level and output targets have to be forward-looking (Goodhart, 2005). This poses an empirical problem, as it is not possible to estimate Eq. 2.35 with OLS. The estimation of the monetary policy parameters in the reaction function can be performed by IV-OLS or GMM, hinging on an exogenous instrument or on the retrieval of the exact information set available to monetary policymakers at specific points in time (Goodhart, 2005; Cobham and Kang, 2013). This is a thorny issue, as in reality Central Banks cannot perfectly react to contemporaneous inflation as it becomes known with a measurement lag. On the other hand, Central Banks have forecasting departments specifically tasked with the job of informing the decision making process. This latter consideration justifies forward-looking Taylor rules in rational expectations setting, signifying a perfect ability of the Central Bank to forecast future inflation. A purely backward looking Taylor rule is easier to estimate empirically as it provides a neat set of zero restrictions that are translatable 1-to-1 to time series methods, but these restrictions come at a price. The pitfall is, of course, the risk of ending up with a mis-specified model since this approach fails to disentangle the interest rate from its endogenous component.

A high interest smoothing coefficient can be the by-product of mis-specification as, for the UK, it does not tally with the official communication of the Bank of England. BoE has maintained that it do each time the full interest adjustment, without allowing any graduality, whereas step-wise patterns are visible in macro-data (Goodhart, 2005). A high coefficient of interest rate smoothing (0.9) indeed appears in recent empirical estimations for the UK (Kapetanios et al., 2019; Finocchiaro and Von Heideken, 2013), meaning that either the Central Bank reaction function factors in interest rate smoothing or the parameter stems from endogeneity or from the interest rate being locked at zer-lower bound in the most recent years.

In this paper I stick to the conventional wisdom that has informed much of the DSGE literature, sticking to a conventional Taylor rule of the type stated above.

2.3 Log-Linearised Base Model

The model is of 25 equations in 25 variables. I consider 3 exogenous shocks: a monetary policy shock to the Taylor rule, a technology shock affecting the goods production function and a cost-push shock in the Phillips curve. Elasticities are set to $\sigma = \phi = \chi = 1$. Capitalised variables without time subscript represent steady state values. Lowercase variables are log-deviations from steady state.

2.3.1 Aggregate Demand

Where Equations 2.36, 2.37 and 2.39 are the household flow of funds, the borrowing constraint and housing demand. The final two 2.40 and 2.41 are the labour supply schedules of the 2 sectors. Equation 2.36 is obtained combining the household resource constraint with the labour supply FOCs. 2.39 is derived from the household problem FOC with respect of housing and the Lagrange multiplier for the collateral constraint. I report the derivations of key equations in Appendix B.3.1.

$$\zeta_1 C^h c_t^h = \frac{Y_t}{H^h P^h} y_t + h_{t+1}^h + p_t^h - \mu \mathrm{G}(\overline{\omega})(h_t^h + p_t^h + \tilde{\eta}\omega_t) + (1-\delta^h)(h_t^h + p_t^h) + \\ (1-\nu)\frac{H^n}{H^h}(h_t^n + p_t^h) + (1-\delta^h)\frac{\Phi(\overline{\omega})}{1+RR}\left[l_{t+1} - l_t - RR(l_t + r_t^l)\right] \quad (2.36)$$

$$l_{t+1} = E_t[h_{t+1}^h + p_{t+1}^h] + E_t[\omega_{t+1}]\overline{\omega}\iota - r_t^l \quad (2.37)$$

$$\tilde{\lambda}_{t-1}^c = \frac{\gamma}{\gamma-\beta}(-c_{t-1}^h - r_{t-1}^l) - \frac{\beta}{\gamma-\beta}(r_t^l - r_{t-1}^l) \quad (2.38)$$

$$\varsigma\beta\zeta_1(c_t^h - h_t^h - p_t^h) = (p_{t-1}^h - p_t^h + c_t^h - c_{t-1}^h) + \beta\mu G[\overline{\omega}](\eta\omega_t) - (1-\delta^h)C^h\lambda^c\Phi[\overline{\omega}](c_t^h + \tilde{\lambda}_t^c + \iota\omega_t) \quad (2.39)$$

$$n_t^c = w_t^c - c_t^h \quad (2.40)$$

$$n_t^c\,(1-z) + z n_t^h = w_t^h - c_t^h \quad (2.41)$$

Where the logarithmic derivatives of $\Phi[\overline{\varpi}]$ an da $G[\overline{\varpi}]$ are:

$$\iota = \frac{\Phi'[\overline{\varpi}]]}{\Phi[\overline{\varpi}]} \qquad\qquad \eta = \frac{G'[\overline{\varpi}]]}{G[\overline{\varpi}]}$$

and they are arising from the log-linearisation around the steady state:

$$\ln f(x_t) - \ln f(X) \approx \frac{f'(X)}{f(X)}(x - X) \tag{2.42}$$

2.3.2 Aggregate Supply

The aggregate supply block includes the Entrepreneur's flow of funds 2.43, Euler equation for savers 2.44, an housing prices equation 2.45 and a Phillips curve 2.46. There are two production functions for housing units 2.47 and goods 2.48. Wages in the two sectors are 2.49 and 2.50.

$$\begin{aligned}\zeta_2 c_t^e = \nu \frac{H^n}{H^e}(h_t^n + p_t^h) - h_{t+1}^e - p_t^h + (1-\delta^h)(h_t^e + p_t^h) + \\ (1-\xi)(1-\delta^h)\frac{H^h}{H^e}\frac{\Phi(\overline{\varpi})}{1+RR}[d_t - d_{t+1} + RR(r_{t-1}^d + d_t)]\end{aligned} \tag{2.43}$$

$$r_t^d = E_t[c_{t+1}^e] - c_t^e \tag{2.44}$$

$$p_t^h - E_t[p_{t+1}^h] - c_t^e + c_{t+1}^e = (E_t[h_{t+1}^n] - h_t^e)\frac{\delta^h \nu}{\delta^h \nu + \left(1-\delta^h\right) H^e} \tag{2.45}$$

$$\pi_t = \gamma E_t[\pi_{t+1}] - \frac{(1-\Theta)(1-\gamma\Theta)}{\Theta}(x_t - s_t) \tag{2.46}$$

$$h_t^n = \nu h_{t-1}^e + n_t^h\,(1-\nu) \tag{2.47}$$

$$y_t = n_t^c + a_t \tag{2.48}$$

$$n_t^h + w_t^h = h_t^n + p_t^h \tag{2.49}$$

$$n_t^c + w_t^c = y_t - x_t \tag{2.50}$$

2.3.3 Banking Sector

The banking sector accumulates equity at the rate 2.51. The lending spread is determined by 2.52. RWA are 2.53 and profits are given by 2.54.

$$k_t = \left(1 - \delta^k\right) k_{t-1} + \delta^k j_t \tag{2.51}$$

$$r_t^l = r_t^d + \frac{\theta \, \xi^3}{RR} \left(l_t - k_t + rw_t\right) \tag{2.52}$$

$$rw_t = \rho_{rw} \, \aleph \, \overline{\omega} \omega_t \tag{2.53}$$

$$j_t = \frac{l_t + d_t \, (\xi - 1) - (\gamma - 1) \left[r_t^l + r_t^d \, (\xi - 1)\right]}{\xi} \tag{2.54}$$

2.3.4 Monetary Policy

The nominal rate of interest is set according to the Taylor Rule in 2.55. The real rate of interest is r^n net of inflation in 2.56.

$$r_t^n = (1 - \rho_r) \left(p_{t-1}^h \rho_{ph} + (1 + \rho_\pi) \, \pi_{t-1} + \rho_y \, y_{t-1}\right) + \rho_r \, r_{t-1}^n + \varepsilon_t^m \tag{2.55}$$

$$r_t^d = r_t^n - E_t[\pi_{t+1}] \tag{2.56}$$

2.3.5 Market Clearing

Market clearing conditions encompass the national income identity 2.57 and the banking capital identity 2.58. The housing stock is given by 2.59 and it evolves according 2.60.

$$y_t Y = c_t^h C^h + c_t^e C^e + \delta^k K k_t \tag{2.57}$$

$$l_t = d_t \,(1-\xi) + k_t \,\xi \tag{2.58}$$

$$H^h h_t^h + h_t^e H^e = h_t \tag{2.59}$$

$$h_t = H^n h_t^n + \left(1-\delta^h\right) h_{t-1} \tag{2.60}$$

2.3.6 External Shocks

$$a_t = \rho_z a_{t-1} + \varepsilon_t^a \tag{2.61}$$

$$s_t = \rho_x s_{t-1} + \varepsilon_t^x \tag{2.62}$$

2.4 Calibration and Impulse Responses

2.4.1 Calibration

Table 2.1: Calibrated Parameters

Parameter	Value	Description
σ	1	Consumption Elasticity
χ	1	Housing Consumption Elasticity
ϕ	1	Inverse of Frisch Elasticity
β	0.97	Borrower's Discount Rate
γ	0.99	Lender's Discount Rate
μ	0.12	Loan Monitoring Costs
δ^b	0.059	Bank Capital Depreciation
δ^h	0.005	Housing Stock Depreciation
ξ	0.09	Capital Adequacy Ratio
ν	0.3	Share of Housing in the Prod. Function
θ	11	Banking Capital Costs
Θ	0.75	Calvo Parameter
ς	0.20	Housing Preference
z	0.87	Labour Substitutability
Mortgage Parameters		
$F[\overline{\varpi}]$	0.03	Annual Default Probability
$\overline{\varpi}$	0.63	Shock Cut-Off
σ_ω	0.18	Shock Standard Deviation
μ_ω	0.016	Shock Mean
$\Phi[\overline{\varpi}]$	0.63	Steady State LTV
$\aleph$	14.5	Risk-weight elasticity

The calibration is standard. Elasticities are set to 1 for computational ease in solving analitically for the steady state and get log-linearised conditions. Banking sector parameters are taken from Gambacorta and Signoretti (2014). The Calvo parameter Θ is calibrated at 0.75 to signify a year average period of price adjustments.

The risky mortgage parameters are calibrated to be close to UK mortgage credit parameters and so I set a default rate of 3%, to which it corresponds a loan to value of 60% for $log(\omega) \sim N(-\frac{0.18^2}{2}, 0.18)$. These mortgage parameters are reasonably similar to the British credit market reported in Table 2.2.

Under this standard calibration, the derivative credit ratios are reasonably similar to UK series. The only parameter which is counterfactually different from national statistics is the consumption to housing value held by households. Even with a higher consumption elasticity σ, household consumption to housing is less than observed 20%.

Table 2.2: Steady State Ratios

Steady State Ratios	Calibrated	Matching	Source	Database
Default Rate	3%	1.5-2.5%	Fitch Ratings, Bank of England	
Loan-to-value	63%	60%-70%	Bank of England	ABS Portal
Debt-to-Income	581%	550%	ONS	National Statistics
New Housing to GDP	2.84%	2.15%	ONS	New Orders in the Construction Industry
Household Consumption to Housing Value	9%	20%	ONS	

2.4.1.1 Autoregressive Parameters and Shocks

Table 2.3: Autoregressive and Policy Parameters

Parameter	Value	Description
ρ^r	0.73	Monetary Policy Inertia
ρ^π	0.27	Inflation Weight
ρ^y	0	Output Weight
ρ^{rw}	0.1	Risk-Weight Inertia
ρ^a	0.9	Technology Shock Persistence

The Taylor rule parameters are calibrated as in Iacoviello (2005) as, for the baseline model, the weight on output is 0.

2.5 The Credit View of Monetary Transmission

In this section I show how the model presented above features a richer credit channel of transmission while displaying a tractable log-linearised structure in a simple two-agents setting. In particular, banking assets allow '*credit view*' to be reflected into the model along the two main channels of transmission: sticky prices and inter-temporal consumption.

Financial imperfections, as costly state verification, pertains to the '*(broad) credit channel*', as the wedge between internal and external finance premiums, and '*(bank) lending channel*' captures what happens on the supply side of loans, using the channels defined in Repullo and Suarez (2000). This friction intervenes on the dynamic profile of consumption, as households are collateral constrained, there is no consumption smoothing and future consumption depends on current housing investments.

A third, and quantitatively important channel, is the relative deflation of debts when prices rise (Fisher-effect). In the baseline model mortgages are indexed to inflation and therefore there is no inflation risk borne by borrowers and I abstract from the debt deflation channel. In the equations below, for illustrative purposes, I include in red the inflation rate π, to show that with a simple modification the nominal debt channel can feature in this model too.

Perhaps the most familiar item in the financial friction literature is the Borrowing Constraint, which dynamically links loans to asset prices. The key accelerating mechanism stems from the circularity between loans and house prices: if households post their proprietary housing value as collateral to get a mortgage loan, then falling real-estate prices tighten the borrowing constraint resulting in a lower housing investment in the following periods.

$$l_{t+1} = E_t[h^h_{t+1} + p^h_{t+1}] + E_t[\omega_{t+1}]\overline{\omega}\iota - r^l_t t + E_t[\pi_{t+1}] \qquad \text{(Borrowing Constraint)}$$

Here the key modification carried from Lambertini et al. (2017b) is the fact that LTV is *endogenous* and varies with the underlying credit conditions. This provides another source of transmission, which materialises in the $\overline{\omega}\omega_t\iota$ term of Borrowing Constraint (absent in Kiyotaki and Moore's style of models[14])

Since the LTV increases in $\hat{\omega}$, loans expand also in response of log-deviations from steady state of $\hat{\omega}$. Whilst LTVs aren't fixed and vary with credit conditions, this model delivers a counterfactual outcome, in the sense that the model-implied LTV *increases* in crises. LTV ratios usually increase in benign periods to reflect the more ready availability of credit. In this model LTV is synonymous with the lender's share, so that financiers require a higher share of the realised risky housing project to be persuaded to lend to such projects.

If loans are to be re-paid in nominal terms, then positive inflation is beneficial in the sense that it deflates the outstanding debt, easing the collateral constraint and allowing households to expand their real-estate ownership.

Another uncommon feature of this model is that households are net borrowers and entrepreneurs net lenders (through their deposit holdings). Usually the financial frictions literature emphasises the opposite case, with entrepreneurs projecting

[14] Among which the following ones are cited in this paper: Iacoviello (2005); Gerali et al. (2010); Gambacorta and Signoretti (2014)

their equity into risky projects and households smoothing their consumption with savings. Usually in the literature the presence of indebted households is modelled by introducing impatient households as a third and distinct agent.

In this model I wish to emphasise a household insolvency channel, therefore in the interest of simplicity and tractability I reduce the heterogeneity in agents retaining only credit constrained mortgagors. This modelling choice reflects an important and realistic feature of credit markets: a big proportion of households is of a wealthy hand-to-mouth type. i.e. holding no liquid assets but only illiquid ones (Kaplan et al., 2014) .[15] This is the case for this model: households' only (illiquid) asset is their housing, which yields them utility in terms of dwelling.

The collateral constraint prevents household consumption smoothing. Households are only able to lever their housing holdings to finance their current consumption and the fact that in the steady they are exactly hitting their collateral constraint makes them unable to smooth it throughout their infinite lifetime. Since the collateral constraint is binding, households' Euler equation is Eq. 2.11 links future consumption only to the period-to-period expansion of loans.

Financial frictions enter the housing problem by changing their consumption profile as households can substitute consumption of goods with housing value. The core net-worth channel stemming from risky mortgages lies in the following log-linearised condition:

$$\begin{aligned} y_t = c_t^h \frac{\zeta_1}{\Xi} - \frac{1}{\Xi}\left[(1-\delta^h)(h_t^h + p_t^h) - h_{t+1}^h - p_t^h - \mu G[\overline{\omega}](\eta\,\omega_t + h_t^h + p_t^h)\right] - \\ - \frac{(1-\nu)}{\Xi}\frac{H^n}{H^h}(h^n + p_t^h) + \frac{(1-\delta^h)}{\Xi}\frac{\Phi[\overline{\omega}]}{1/\gamma}\left[(l_t - \pi_t)\frac{1}{\gamma} + r_t^l\left(\frac{1-\gamma}{\gamma}\right) - l_{t+1}\right] \end{aligned} \tag{2.36, Aggregate Demand}$$

This equation captures the log-linear relationship between output and the other items in the households' budget constraint, approximately:

$$\Delta y_t = \Delta\text{consumption} - \Delta\text{housing worth} - \Delta\text{construction wage} + \Delta\text{loans} \tag{2.63}$$

[15] As opposed to poor hand-to-mouth, which hold neither liquid nor illiquid assets.

Output deviations are equal to consumption ones less deviations in housing value holdings.[16] Output is therefore increasing in consumption. The ratio of the steady state share of new housing captures the quantitatively small log-linear effect of labour income (and thus the state of production in the new housing sector). This equation can be conceived as the aggregate demand for the economy and shows the general equilibrium effects of the collateral channel and the insolvency channel, as in a credit expansion households can borrow more to finance their consumption, optimising at the same time their real-estate holdings which they can post as collateral. Again, the nominal debt channel can be built in this equation and it works exactly as above: it deflates loans, thereby increasing household leverage capacity. The insolvency channel has relatively minor effects on output in comparison to deviations from steady state loans (encapsulated in the last term).

Finally, $\mu G[\overline{\varpi}]$ captures the monitoring costs associated with default, i.e. a dead-weight loss eating into existing housing value. So a higher default threshold translates into a permanent destruction of housing value, impairing the capacity of households to borrow against it.

The Eq. 2.36 has the flavour of a dynamic aggregate demand function, where fluctuations stem from the portion of housing that borrowers can lever into loans (credit market) and the increase in value of accumulated housing (net-worth). Hence, such investment-savings equation ('IS') can be called a 'levered' IS, as opposed to the standard IS arising from textbook 3-equation models.[17] The elasticity of aggregate demand to house prices features prominently in the transmission mechanism. Since loans are the biggest driver for demand, a case where the policymakers commit to curb credit aggregates to control the demand might be compelling.

A built-in rule to control the expansion of credit is modelled by means of mechanistic risk-weighting presented above, which ensures a quick reversion to steady state lending whenever banking leverage and portfolio risk are too great with respect to their steady state values.

In this paper the aggregate demand is augmented with housing value and debt,

[16]Multiplying both sides by $P^h H^h$ makes evident that log-deviations are re-scaled by housing value owned by households.

[17]For comparison, in Appendix B.6 I derive a similar aggregate demand (IS) equation for a much simplified New-Keynesian model, finding that such IS is more sensitive to the interest rate that a IS curve form Walsh (2010) textbook 3-equations model.

hence it is a 'levered' IS schedule. Amplification does not only stem from fluctuations in output gap, but also from general equilibrium effect in the demand itself. Modifications to the interest rate do not only enter the demand by triggering a change in the inter-temporal evolution of consumption, but they are apt to restrict present consumption by constraining household leverage. Since wealthy hand-to-mouth agents do not have the option of consumption smoothing, the impact on output is greater, as a consequence of the high elasticity of such a levered IS to the interest rate[18]

On the supply side of credit, banks are constrained by capital adequacy ratio, forming a banking supply channel (as in (Gambacorta and Signoretti, 2014)). The key result is the lending spread equation, which endogenously arises from the banks' optimisation programme.

$$r_t^l - r_t^d = \frac{\theta \xi^3}{RR}(l_t - k_t + rw_t) \qquad \text{(Lending Spread)}$$

Banks can deviate from the risk-weighted CAR ratio $(l_t - k_t + rw_t)$ compensating the adjustment costs charging a higher spread.

As said above and consistently with the literature, I assume that the bank does not incur credit losses,[19] but banking losses stem only from balance-sheet management costs. Relaxing this modelling choice would provide a more realistic transmission of credit-negative shocks through the banking sector.

Envisioning a micro-founded risk-weight scaling the asset portfolio on the grounds of probability of default and recovery can circumvent that and represent an indirect way to factor credit risk into the banking choices. If assets are risk-weighted, loans and banking leverage will be more pro-cyclical.

The risk weight is the log-linear version of 2.32:

$$rw_t = \rho_{rw} \aleph \overline{\omega} \omega_t \qquad \text{(Risk Weight)}$$

[18] We can further substitute the 2.37 into 2.36, Aggregate Demand to better visualise how the interest rate enters the aggregate demand.

[19] Or, in other words the riskless rate is deemed to be equal to the ex-post realisation of the shock. See Benes and Kumhof (2015) for a relaxation of this condition holding in *in expectations* but not ex post.

As a preliminary result, we can see in Table 2.4 how the inclusion of defaults result in a heightened volatility of inflation as compared to a benchmark model without insolvencies.[20] There appears to be an output/inflation trade off associated with the standard Taylor rule. When insolvencies are modelled the variance of inflation is greater than in the case without defaults but the variance of output is lower. The volatility of inflation is larger with LATW rules, implicating that a deviation from inflation targeting works to keep output under control but at the cost of allowing more volatility in the price level.

Table 2.4: Variance of Output and Inflation with and without Insolvencies

The figure is obtained by simulating the base model for 1000 periods and calculating the variance of relevant variable under different monetary policy rules. The base model is considered against a benchmark model without default (presented in Appendix B.4). The base parametrization for the Taylor Rule in Eq. 2.35 is then augmented with positive weights to house prices and loans.

	Var(y)	Var(π)
Taylor Rule - Standard Calibration		
Base Model	2.42	0.84
w/out Defaults	2.25	0.51
LATW House Prices - $\rho_{ph} = 0.15$		
Base Model	2.00	0.85
w/out Defaults	1.91	0.79
LATW Loans - $\rho_l = 0.15$		
Base Model	2.28	0.83
w/out Default	1.91	0.78

2.5.1 Impulse Responses

The objective of this section is to analyse the effects of external shocks and the response of the dynamic system on impulse. My aim is therefore to show how the features described above characterise a model economy. In particular, the presence of a fleshed out credit channel implies a further conduit of shock transmission that reverberates on aggregate demand.

I test here three relevant shocks: a productivity shock, a negative interest rate shock and a cost-push shock. These shocks are recognised as the driving forces of business

[20] The simplified analytical model without defaults and a constant LTV ratio is presented in Appendix B.4. It is a version of Iacoviello (2005) with households borrowing and banks as financial intermediaries.

cycle moving variables away from their steady state. Since I am here concerned with the conduct of monetary policy, I am interested in plotting the nominal interest rate reaction to exogenous disturbances to gauge the time-path of variables post shock under a standard Taylor rule and different model configurations.

Although this Chapter strongly focuses on monetary policy, I include additional shocks in order to gain further insight into the workings of the model presented above with respect to a wider array of economic disturbances that affect the credit variables. This helps me to obtain a more informed analysis of the effects of aggregate shocks with regards to the different features of the model at hand.

I trace the monetary transmission mechanism given by the key equations showing the importance of the credit channel when the system is hit by exogenous disturbances. I then experiment a dynamic risk-weighting rule to investigate the interaction between monetary and macroprudential feedback stemming from default risk. Since the model is log-linearised, the charted impulse response functions are percentage deviations from steady state values.

Figure 2.2: Technology Shock

Percentage deviations from Steady State.

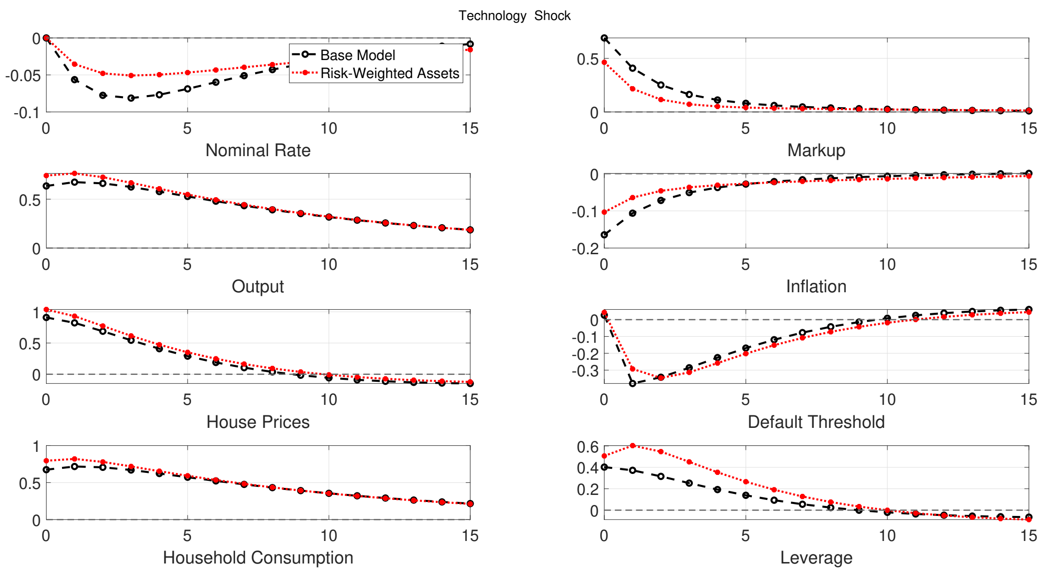

2.5.1.1 A Technology Shock

Wealth and Housing Demand Channel

In the baseline calibration, a 1% technology shock (shown in Fig. 2.2) affecting the goods production increases output and wages. Here the shock affects the aggregate supply first but it gets transmitted through the aggregate demand, as households end up consuming more goods and housing units. Over time, household consumption is persistently higher than its steady state. Households can scale up their real estate holdings leveraging real assets into loans. The increased demand for real estate impacts on prices, which display a hump shape tapering after 10 quarters - as seen in Fig. 2.2.

Imperfect substitution of labour inputs matter, as hours worked in the goods' sector depend positively on output and negatively on consumption and markup (shown in Eq. 2.69). With imperfect substitution wages in two sectors fail to equalise, hence the wage in the construction sector falls below steady state and the employment in the goods sector rises. This feeds into housing demand and supply consideration, contributing to higher housing prices.

Collateral Channel and Housing Market Acceleration

The technology shock causes a swift re-allocation of real-estate, from offices to dwellings. This makes households more able to project their housing ownership into new mortgage loans. The shift is facilitated by the fact that I do not model adjustment costs for the agents when they manage their housing stock.[21]

Concurrently, entrepreneurs produce and sell less housing units to households. Therefore the supply of real estate shrinks, driving up housing prices. The housing stock is non-stationary and therefore a shock produces a permanent re-allocation of housing units between entrepreneurs and borrowers.

Increasing housing prices ease the household collateral constraint (Eq. 2.37) driving up the lending expansion. Housing purchases from households sets in motion an economic accelerator as household can buy and accumulate more dwellings with which they can collateralise new loans. Households cannot smooth consumption so their only way to preserve wealth is by the accumulation of housing when the borrowing constraint releases.

On the back of the boom, the default threshold is lower, meaning that a greater idiosyncratic shock is needed to tip households into default. The technology shock

[21] Also Iacoviello (2005) sets adjustment costs to 0 after having verified that positive adjustment costs produce quantitatively small effects.

causes an output boom leading to an increase in house prices. This enables a higher valuation of existing housing, i.e. collateral and therefore makes the case of a default less likely.

Credit Supply and Lending Channel

In this model an economic boom corresponds to a credit boom. Since borrowers cannot smooth their consumption, the only way for them to access to more consumption in the future is to commit to more housing investment in the present. Thus, a technology shock increases the demand for mortgage loans. The profit-maximising banks accommodate the excess demand by scaling up their loan portfolio up to the point in which profits offset capital management costs for deviating from the target capital adequacy ratio.

As a result, mortgage loans depart significantly from their steady state as also do deposits. Entrepreneurs sell their housing stock converting their asset holdings from real-estate to deposits (Eq. 2.43) and they consume more. But aggregate loans grow faster than deposits, meaning that banks increase their leverage on the back of accumulating positive profits.

Aggregate Supply and Sticky Prices

The sluggish adjustment of inflation on impulse is due to the Calvo sticky price model. Monopolistic competition ensures that firms re-set prices following the evolution of their time-varying markup, i.e. the inverse marginal costs. A technology shock decreases the marginal costs, as it makes cheaper to produce goods, hence inflation reduces on impulse.

Nominal Rate Reaction

Since the Central Bank responds only to inflation deviation, it counters the shock easing monetary policy to bring back the price level to its steady state.

2.5.1.2 Monetary Policy Shock

Figure 2.3: Monetary Policy Shock

Percentage deviations from Steady State.

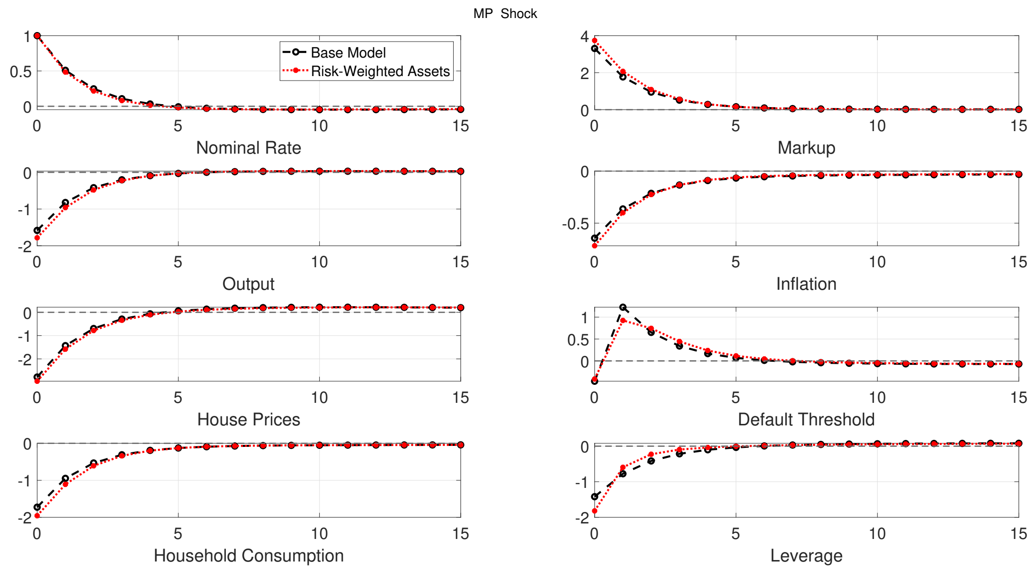

Wealth and Housing Demand

A negative monetary policy shock is identified by a 1% one-off increase in the nominal rate of interest. This transmits to entrepreneurs via their inter-temporal consumption and to households through the cost of lending. It alters the demand and supply in the housing market and entrepreneurs-households consumption patterns.

The key part of the transmission mechanism stems from wealth effects in the household and for firms. Since the interest rate is higher, households consume fewer goods and housing services.

On the aggregate supply side, the fall in house prices and housing demand from households prompts entrepreneurs to accumulate commercial real-estate and use it in the production process. This impacts firms' marginal costs and their markup falls below their steady state level. This feeds into their price setting protocol and inflation rises.

The Central Bank's modifies entrepreneurs' preferences via the Euler equation 2.44, as entrepreneurs increase their future consumption relative to the current, but this effect is negligible compared to the increase in its income due to the increased availability of real-estate. Simultaneously, housing consumption falls on the back of subdued housing investment since it is now more expensive to borrow. Total output is driven by a sharp fall in household consumption, lowering employment and wage in the goods sector.

Therefore, the exogenous tightening is accompanied by falling output and inflation.

Collateral Channel and Housing Market De-acceleration

The consumption channel translates into a credit crunch, as loans are falling quicker than deposits and house prices are also tightening household collateral constraint. Residential housing de-accumulation increases the default threshold, reducing borrowers' housing investments also in future periods.

Entrepreneurs, on the other hand, manage to increase their consumption as they accumulate real estate. At the same time, commercial real estate is an input to the production of new housing, there is thus an increase in the production of new housing also because since monitoring costs are proportional to housing value, an uptick in defaults causes a destruction of property in the model. Hence, hours worked and wages rise in the housing production sector. The increased production in the housing sector partially mitigates the downturn, sustaining house prices and increasing the housing supply.

Credit Supply and Lending Channel

The banking loan portfolio shrinks as well as deposits. Since loans are falling faster than deposits, the banking sector faces a collapse of the net interest margin, accruing losses. As the capital is impaired, banks need to slowly re-grow the loan portfolio to steady state while complying with the capital adequacy ratio. So the monetary policy shock is consistent with a credit crunch scenario that drags for some 5 periods, given the exogenous constraint on bank leverage.

Housing Production Increase

The low demand for goods depresses the wage in the goods sector, triggering a migration of labour to the housing production sector, where the wage has increased on the back of a production boom bolstered by greater entrepreneurial demand. The stock of real-estate grows but it remains more concentrated in the firms' hands, which can use it to produce more commercial units and hence continue investing.

Aggregate Supply and Sticky Prices

The striking feature of a monetary policy shock in this model is the level of persistence of inflation, and given the Taylor rule, of the interest rate.

This is attributable to the general equilibrium effects stemming from the heterogeneity of having two agents populating the model. In a single agent new-Keynesian model, the adjustment of a monetary policy tightening passes almost exclusively through the dynamic IS (which contains the Euler equation), making consumption now less appealing when compared to savings. Here this effect is present through the entrepreneurs Euler equation 2.44, but the effect that dominates is the income effect. More demand of housing from entrepreneurs fuels entrepreneurial activity and the higher investment at time 0 translates into higher real estate investments at later times. A monetary policy shocks reduces the price of housing and hence the marginal costs for entrepreneurs to produce new real estate using old real estate as input. The markup rises above steady state and inflation falls.

2.5.1.3 Risk-Weighting of Banking Assets

The risk-weight appended to banking assets provides another mechanism of shock propagation: assets are now made more procyclical as they are allowed to expand when the credit risk is moderated and contract when the default threshold endogenously rises.

This model amendment goes in the direction of making the lending channel more realistic: banks risk taking is normally reflected in the regulatory leverage ratio they are required to maintain.[22]

This modelling choice means that banking assets are more volatile than in the baseline model. In the case of a technology shock, the risk weight decrease allowing banks to gear up more, this reflects in an increase of house prices, household consumption and output through the mechanism described in the above paragraph.

The difference in aggregate demand from the base model makes changes markup and inflation dynamics. The Central Bank can keep the nominal interest rate higher for longer in the case of a technology shock coinciding with an economic boom.

In case of a monetary policy shock, these effects are much smaller due to the credit

[22] Risk weights were first established in the Basel Framework and embedded in financial regulation, as for instance the Capital Requirements Regulation in Europe.

crunch I described above. Risk weighted assets modulate the effective leverage the bank is allowed to take, therefore the supply of more or less loans impact the household borrowing constraint changing the default threshold.

2.5.1.4 Nominal Debt Channel

In Iacoviello (2005) the transmission channels belonging to the credit view are the collateral channel and the Fisherian channel of debt deflation. He shows how when both are 'activated' in the model, output falls more strongly upon a monetary policy shock. The emphasis on nominal debt is lacking in Gambacorta and Signoretti (2014), as they elect to focus more on the collateral channel, although they analyse a cost-push shock with indexed debt.

In Section 2.5, I show how to factor in nominal debt repayments in the base model.[23] In this alternative specification mortgages are no longer indexed to inflation and therefore there is an additional shock transmission channel at play. This modification carries implications for the credit equations block because in the base model, with indexed debt, inflation only affects the aggregate supply through the price setting friction. In the model with nominal debt, inflation enters the household credit decisions and modulates consumption/investment patterns for both households and entrepreneurs through time.

Higher inflation reduces the real value of pre-determined mortgage loans. So, a bit of inflation is credit positive as it improves the mortgagor's borrowing constraint. It also reduces the value of deposits, appearing in the entrepreneur's Euler equation and affecting inter-temporal consumption.

In this section I show the impulse response functions to the shocks described above, describing how the nominal debt channel changes the baseline model. Furthermore, I consider an additional inflation shock, i.e. an exogenous disturbance to the Phillips curve that reflects external inflationary pressure not originating from the endogenous staggered prices setting. In all cases, the debt deflation channel is quantitatively strong and changes the response of the default threshold on shocks.

Technology Shock

[23] The detailed derivations are in Appendix B.5.

Figure 2.4: Technology Shock

Percentage deviations from Steady State.

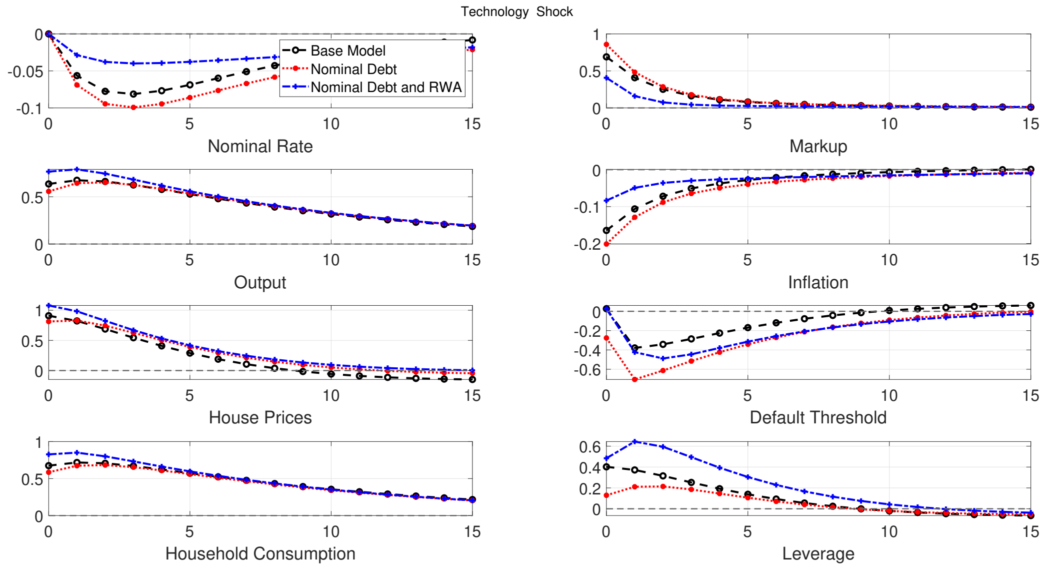

Following a technology shock (Fig. 2.4 the mechanism detailed in the previous section still holds, but now the default threshold falls more strongly on impact and peaks at a lower level on the back of inflation easing the collateral constraint during the economic boom. The marginal product of labour rises with wages, pushing up consumption and output. The economic boom is also a credit boom as households are able to invest more in real estate to secure more consumption in the future. Consequently households can take on more leverage, contributing to a house price increase in the first periods upon shocks.

On the credit supply side banks are also leveraging up to accommodate the fresh credit demand. When risk-weighted assets are modelled they enable banks to take up much more leverage on the back of the improved credit quality of borrowers. So the case of risk-weighted assets is the one with the strongest credit growth, as the banking portfolio is now less risky.

In turn, monetary policy is adjusted following the baseline backward looking Taylor rule. To nominal debt corresponds a better state of the aggregate demand and therefore the inflation falls by less than in the baseline model. Accordingly, the Central Bank can cut the nominal rate by less than it would in the base model.

Figure 2.5: Monetary Policy Shock

Percentage deviations from Steady State.

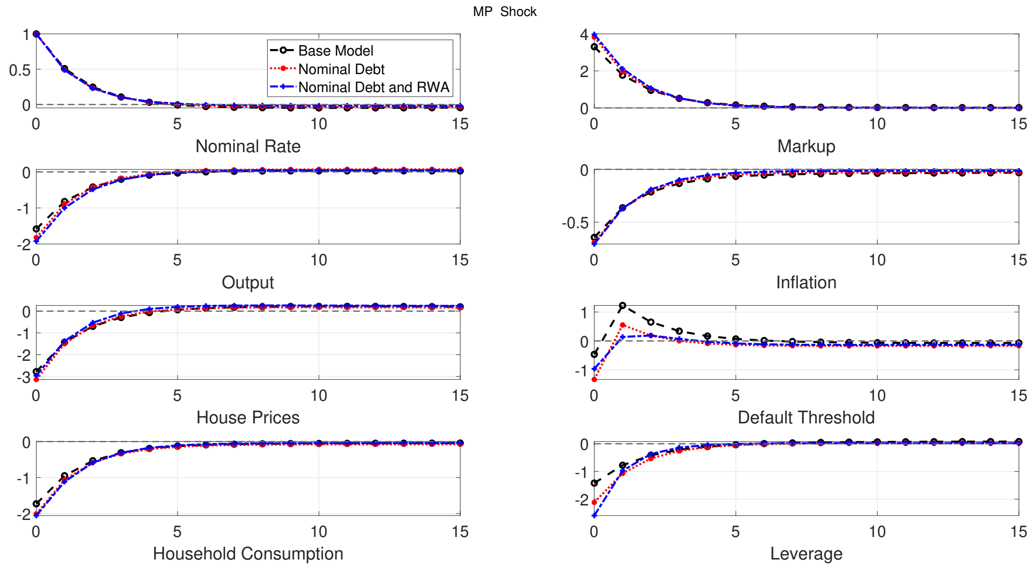

Monetary Policy Shock

In Fig. 2.5 I plot the response of key variables to a surprise 1% monetary policy shock. As stated in the Section 2.5 above, this is tantamount to model the Fisherian channel of debt deflation.

Here the time path of variables factoring in nominal debt is not drastically different from the baseline model. What changes is the credit block. Whilst high inflation transfers wealth from lenders to borrowers, a negative monetary policy shock does the opposite. As an higher interest rate keeps inflation below its steady state, borrowers reduce their holdings as it is more expensive to service debt in terms of higher interest rate but also in paying back it back in moneys that are worth more in real terms (due to debt deflation).

On the other side of the credit friction, entrepreneurs (lenders) get more real deposits and therefore are able to purchase more housing value to produce more goods and new real estate.

The debt deflation channel depresses the collateral constraint as the inflation rate enters with a negative sign in the Borrowing Constraint, but the general equilibrium effect on the default threshold makes it less responsive the shock than in the base model.

This effect is due to the presence of the banking sector, as a monetary policy shock improves the leverage ratio of intermediaries and it translates into a lower gross lending rate.

Figure 2.6: Cost Push Shock

Percentage deviations from Steady State.

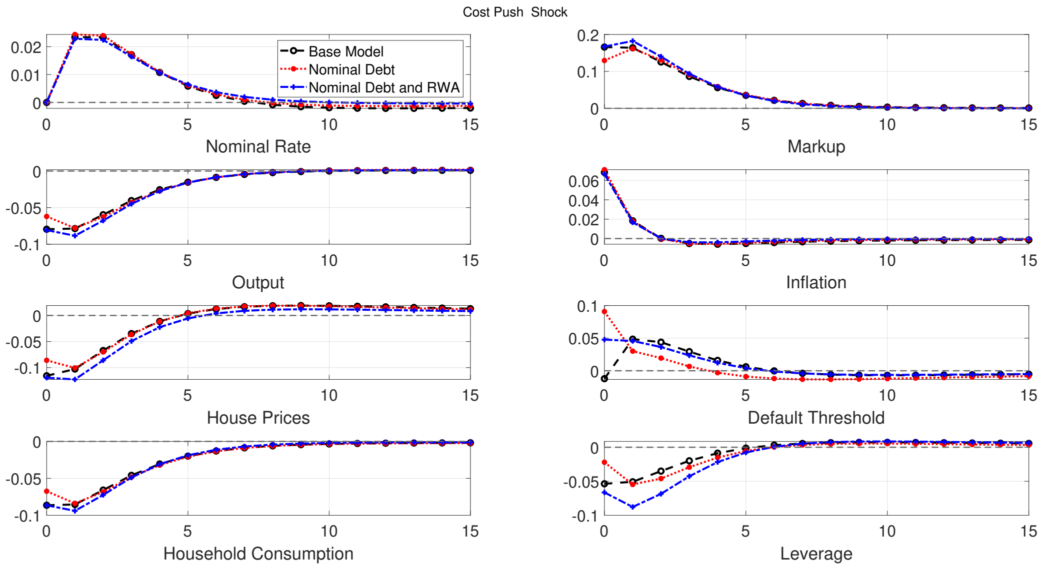

Cost Push Shock

The rationale behind the monetary policy reaction to a cost-push shock is to test the monetary policy rule response to a disturbance that affects aggregate supply.

A cost push shock is identified as an exogenous disturbance to the markup in the Phillips curve (Eq. 2.46),[24] with variance 1. It can be thought as an inflationary impulse originating outside the dynamic system, as e.g. a supply-side commodity shock.

$$s_t = \rho_x s_{t-1} + \varepsilon_t \tag{2.64}$$

A cost-push shock therefore entails a time-varying markup and, consequently, inflation and a persistence dictated by the parameter ρ_x in the equation above. I set

[24] As in Negro et al. (2020), the shock is scaled by the Phillips curve slope coefficient.

$\rho_x = 0.5$ as in Gambacorta and Signoretti (2014), meaning that a cost push shock is somewhat persistent but not as a technology shock.

The immediate result of a supply-sided shock is the spike in inflation, coupled with an instantaneous decline in output and household consumption. The variables hit on impulse converge back to steady state in less than 10 quarters due to the high coefficient given to inflation targeting in the Taylor rule.

This is because the main transmission happens on the back of the Central Bank response and formulation of monetary policy. Since the Central Bank reacts to inflation, it has to hike the nominal rate in the periods after the input (since the policy reaction is entirely backward looking).

The element of novelty year is the response of the default threshold under different model designs. Whether debt is indexed or not influences the magnitude of the fall in lending at time 0, which in turns tighten the borrowing constraint producing defaults.

If loans are risk-weighted, then portfolio credit risk is translated into a different banking leverage. The modulation of the credit supply due to the endogenous variation of the leverage ratio cools off the credit cycle, dampening defaults.

2.6 Policy experiments and Leaning Against the Wind

In this section I construct a policy exercises analogous to the one performed by Gambacorta and Signoretti (2014) to retrieve the efficient frontier for the output/inflation trade off. Consistently with the literature, I postulate a Central Bank's quadratic loss function of the type:

$$\text{Loss} = \text{Var}(\pi) + \alpha\,\text{Var}(y) \tag{2.65}$$

This ad hoc formulation of Central Bank's preference replaces a micro-founded social welfare function maximisation and assumes that central banking losses are a valid proxy for social losses. In this example the Central Bank cares about volatility of output and inflation, this assumption excludes housing and other macroeconomic variables in keeping with actual central banks official mandates, which concentrate on GDP and inflation stability.[25] As shown above, in the model hereby considered housing is an important role in the transmission mechanism, lying at the centre of the aggregate demand. I show in this Section that exactly for this reason targeting output is similar to targeting housing prices and in assessing the empirical performance of the model at hand, I find a Pearson's correlation coefficient of 0.8 between simulated series of house prices and output. This correlation, albeit milder, is also present in actual data (See 2.6.3).

The Central Bank is assumed to optimise the following backward-looking Taylor rule:[26]

$$r_t^n = (1-\rho_r)\left(\rho_l l_t + p_t^h \rho_{ph} + (1+\rho_\pi)\,\pi_t + \rho_y y_t\right) + \rho_r r_{t-1}^n \tag{2.66}$$

This rule is the same as the one presented above in Eq. 2.35 but with all contemporaneous variables. I consider it with and without interest rate smoothing, since Gambacorta and Signoretti (2014) have not factored any degree of interest rate inertia. I then construct a grid of parameters for ρ_{ph}, ρ_y and ρ_π and compute the loss for each combination in the grid. Finally, I retrieve the parameters under which the loss was minimised for the different weights on output α.

In the case of a standard Taylor rule, the grid spans across the following parameter space: $\rho_\pi \in [0\ 5]$, $\rho_y \in [0\ 2.5]$ and $\rho_{ph} = \rho_l = 0$. In the case of LATW, ρ_{ph} and ρ_l can vary in the interval [0 2.5]. The policy weight on output α is between 0 and 2. The points in the grid are 41 (including zero) along each dimension and equally spaced.

The results in Table 2.6 are very similar to those obtained by Gambacorta and Signoretti (2014), as I find that in cases of a standard Taylor Rule and LATW, the minimum loss is achieved with strict inflation targeting, which results in the maximum coefficient (5) assigned to inflation. In the case of a standard Taylor rule, the optimised output parameter is positive and gradually stronger when the central bank assigns a higher weight to output.

In the second case, with explicit housing prices targeting, the optimised output weight is always 0 or, in a few instances, close to nil. For stricter inflation targeting is not optimal to target house prices but augmenting the Taylor rule in that direction

[25] for a review of this approach see Benchimol and Fourçans (2019).

[26] I try also another rule with house price inflation $\pi_t^h = p_t^h - p_{t-1}^h$, obtaining similar results in the simulations below.

results in reduced losses for greater values of the output weight α. Incidentally, as the central bank deviates from the Taylor rule targeting asset prices, the output coefficient becomes irrelevant and according to the optimised rule, never greater than 0.06.

I share this result with Gambacorta and Signoretti (2014) even though their friction is a Iacoviello's exogenous LTV constraining entrepreneurs. Gambacorta and Signoretti (2014) speculate that in a New-Keynesian setting with credit, financial frictions lie at the heart of the transmission mechanism and therefore asset prices constitute a more direct target than output.

This might mean that targeting asset prices is substitute to output targeting rather than its complement. Targeting loan portfolio deviations from the steady state is similar to targeting housing prices because of the borrowing constraint relationship directly linking asset prices to loans:

$$l_{t+1} = E_t[h^h_{t+1} + p^h_{t+1}] + E_t[\omega_{t+1}]\overline{\omega}\iota - r^l_t t \qquad \text{(Borrowing Constraint)}$$

Table 2.5: Technology Shock

Optimised Taylor Rule Coefficients and Central Bank Losses

	Taylor Rule			House-prices Augmented			
α	Loss	ρ_π	ρ_y	Loss	ρ_π	ρ_y	ρ_{ph}
0	0.02	5.00	0.00	0.05	2.88	0.00	0.06
0.1	0.27	4.75	0.06	0.22	2.88	0.00	0.06
0.2	0.48	4.63	0.56	0.45	2.50	0.19	0.13
0.3	0.67	4.63	0.56	0.57	4.88	0.06	1.25
0.4	0.85	4.63	0.56	0.69	4.88	0.06	1.25
0.5	1.03	4.63	0.56	0.82	4.88	0.06	1.25
0.6	1.20	3.25	0.88	0.94	4.88	0.06	1.25
0.7	1.34	3.25	0.88	1.07	4.88	0.06	1.25
0.8	1.47	3.25	0.88	1.18	2.38	0.00	1.69
0.9	1.60	3.25	0.88	1.26	2.38	0.00	1.69
1	1.74	3.25	0.88	1.34	2.38	0.00	1.69
1.1	1.87	3.25	0.88	1.42	2.38	0.00	1.69
1.2	2.00	3.25	0.88	1.50	2.38	0.00	1.69
1.3	2.14	3.25	0.88	1.58	2.38	0.00	1.69
1.4	2.26	2.50	0.88	1.66	2.38	0.00	1.69
1.5	2.37	2.50	0.88	1.74	2.38	0.00	1.69
1.6	2.48	2.50	0.88	1.82	2.38	0.00	1.69
1.7	2.60	2.50	0.88	1.89	2.38	0.00	1.69
1.8	2.71	2.50	0.88	1.97	2.38	0.00	1.69
1.9	2.82	2.50	0.88	2.05	2.38	0.00	1.69
2	2.93	2.50	0.88	2.13	2.38	0.00	1.69

Table 2.6: Cost-push Shock

Optimised Taylor Rule Coefficients and Central Bank Losses

	Taylor Rule			House-prices Augmented			
α	Loss	ρ_π	ρ_y	Loss	ρ_π	ρ_y	ρ_{ph}
0	0.00	5.00	0.00	0.00	5.00	0.19	0.00
0.1	0.01	4.75	0.50	0.01	3.13	0.00	0.06
0.2	0.02	2.63	0.69	0.02	5.00	1.38	0.06
0.3	0.02	2.88	1.31	0.02	2.75	0.31	0.25
0.4	0.02	1.25	0.56	0.02	2.75	0.31	0.25
0.5	0.02	1.25	0.56	0.03	2.75	0.31	0.25
0.6	0.02	1.25	0.56	0.02	2.63	2.44	0.25
0.7	0.02	1.50	1.13	0.03	3.88	1.69	0.44
0.8	0.03	1.50	1.13	0.03	3.88	1.69	0.44
1	0.03	1.50	1.13	0.03	3.88	1.69	0.44
1.1	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.2	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.3	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.4	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.5	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.6	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.7	0.03	1.00	1.50	0.03	1.50	1.94	0.44
1.8	0.03	1.00	1.50	0.03	2.50	2.25	0.63
1.9	0.03	1.00	1.50	0.03	1.88	1.81	0.69
2	0.03	1.00	1.50	0.03	2.75	1.56	0.75

I also use Eq. 2.35 with a smoothing parameter of 0.9 to assess the effects of a gradual adjustment of the nominal interest rate. I find that interest rate smoothing shifts the efficient frontier mitigating the output-inflation trade-off. With inertia in the nominal interest rate, differences between the LATW and the standard Taylor rule are less significant than in the case without interest rate smoothing. This finding is not surprising and interest rate inertia often features as a desired characteristic for Taylor rules (Woodford, 2003a; Adolfson et al., 2011). Laureys et al. (2021) note that in the presence of financial frictions, interest rate smoothing contributes to attenuate the volatility of credit spreads. This might be the channel at play in the model above, as a gradual adjustment of the nominal rate of interest means that the adjustment of Bank's funding costs is more gentle and so is the squeeze in the supply of credit. In Figure 2.7 it is shown graphically how LATW compares to a standard Taylor rule. LATW gains are less clear-cut in the case of cost-push shocks.

Figure 2.7: Efficient Frontier

Inflation-Output Trade-off with and without interest rate smoothing.

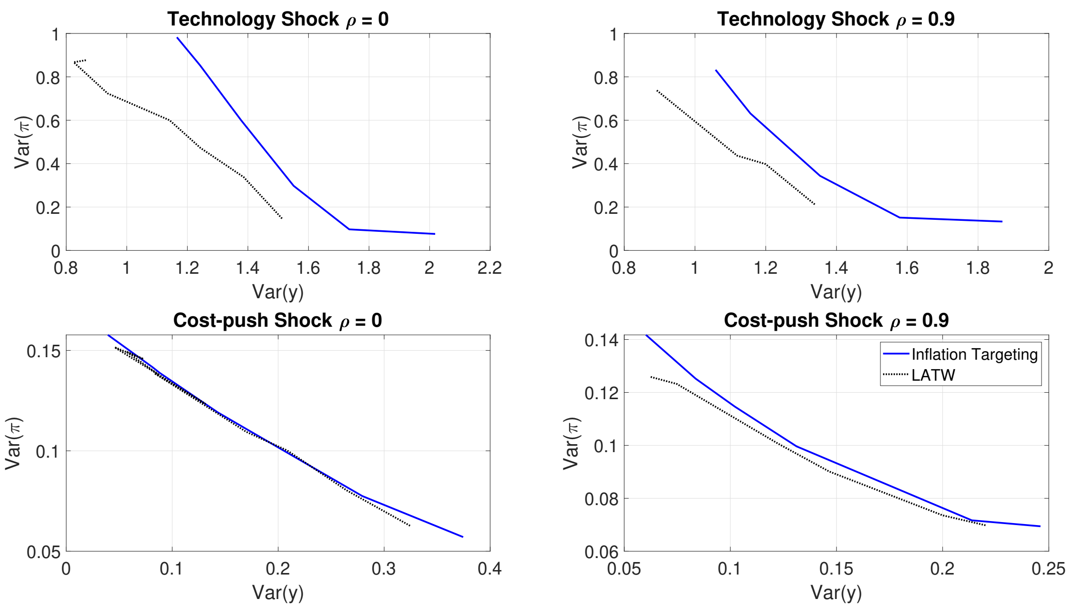

2.6.1 Why is asset price targeting similar to targeting output?

In this section I investigate why targeting house prices in the model outlined above is similar to targeting output. When I 'activate' the LATW rule in the grid search, the asset prices coefficient ρ_{ph} cannibalises output targeting, to which corresponds a constant 0 feedback coefficient in the optimised Taylor rule and a generally high inflation weight ρ_{ph}.[27]

In the Section 2.5 above, I showed how the aggregate demand compounds the net worth channel and the collateral channel of transmission, essentially feeding into the dynamic IS with house prices. This makes output very sensitive to the interest rate as far as it cools off the demand of housing when loans become more expensive and the default rate surges.

In this model, house prices are appearing in the borrowing constraint, as real-estate works as banking collateral, but real-estate prices are also informing and driving the production process. I conjecture that this vicinity of house prices to aggregate supply makes them all the more relevant to the monetary transmission mechanism.

On the demand side, when debt contract are not indexed to the price level, output

[27] As it happens in the Gambacorta and Signoretti's model.

depend positively on the present inflation rate. Inflation is credit positive as it expands the borrowing constraint and allows for more household leverage their household holding into loans thereby increasing output. As households lever up and the demand for housing units picks up, the price of real-estate increases.

On the side of aggregate supply, increasing real estate price means higher marginal product of producing new housing units, as real-estate is also necessary to carry out the production process. Entrepreneurs would ideally increase labour in the real estate sector but they cannot do so promptly since the labour inputs are imperfect substitutes.[28] The wage rate increases in the construction sector more than in the goods sector. Output increases thereby reducing the firms demanded markup with the price increase. There is a circularity of inflation, which first improves the balance-sheet of borrowers and then expands those of firms together with output.

If we eliminate n^c_t, w^c_t, n^h_t, w^h_t from the production functions for housing (Cobb-Douglas) and goods (linear), we can derive the following linearised conditions for output and new housing:[29]

$$y_t = 2a_t - c^h_t - x_t \tag{2.67}$$

$$x_t = a_t + h^n_t \frac{\nu + z}{(1-\nu)(1-z)} - h^e_{t-1}\nu \frac{1+z}{(1-\nu)(1-z)} - p^h_t \frac{1-\nu}{(1-\nu)(1-z)} \tag{2.68}$$

Consolidating the two conditions into a single output equation yields the following general production function for goods:

$$y_t = a_t + (p^h_t - c^h_t)\frac{1}{1-z} - h^n_t \frac{\nu + z}{(1-\nu)(1-z)} + h^e_{t-1}\frac{\nu(1+z)}{(1-\nu)(1-z)} \tag{2.69}$$

So that aggregate output depends positively on house prices with elasticity $\frac{1}{(1-z)}$, where z is labour substitution parameter between the two sectors. New real-estate units h^n_t are produced in a different sector, so in a construction boom the goods output is lower. Whereas if there is a positive deviation of commercial real-estate

[28] When labour inputs are perfect substitutes, the goods sector labour supply and wage will also increase in the goods sector, dampening the wage differential with the construction sector and reducing the impact of increased house prices on output.

[29] Resulting from a few algebraic manipulations of equations 2.40, 2.41, 2.47,2.48, 2.49 and 2.50 in Appendix 2.3.

h^e_{t-1}, then the output is increased, reflecting a better balance-sheet of entrepreneurs and higher wages in the housing construction sector. This shows in the table above, where the real estate augmented rule is associated to a 0 weight on output but a greater inflation parameter than it was the case under the Taylor rule.

In this setting, targeting house prices does affect the borrowing constraint, as detailed in the previous section, but it also does impact the production sectors. Targeting house prices is thus very similar to output targeting because of the way real-estate is used as a factor of production.

The two-agents structure of this model makes them compete for housing. When housing prices are down, firms have an edge in the market as it is cheaper for them to buy existing real estate, convert it into new real estate and harvest the profits thus boosting their future investments. This net transfer of housing wealth from households (borrowers) to entrepreneurs (firms) translates into a higher markup and keeps output sub-optimally low.

Vice versa, in an asset prices boom, household demand thrives and as enterprises are producing more to accommodate the financial acceleration. This interaction connecting demand sided development to supply lowers the markup and causes inflation. Hence there is a link between house prices and goods inflation through the production functions for the economy.

2.6.2 LATW and Nominal Debt

In the policy experiments above, I find that targeting house prices in the model with nominal debt yields dynamic indeterminacy for almost all values of ρ_{ph} in the grid (with the exception of lower ones, i.e. $\rho_{ph} \leq 0.125$). This constitutes a main difference from the Gambacorta and Signoretti (2014) paper, as they are able to tabulate the monetary policy frontier when the debt is nominal beside to when it is indexed to inflation.

Determinacy is an important property of linear rational expectations models and it is investigated analytically in tractable three equations models (for the textbook theory see Woodford (2001); Galí (2009); Walsh (2010)). Dynamic determinacy is associated with the economic intuition of the Taylor principle: the monetary policy rule has to respond more than one-to-one to inflation deviations from steady state for the system to have a stable solution. Lack of determinacy on the other hand means that there are multiple equilibria. [30]

Since my model is medium scale (is in 23 equations), it is not possible to study its roots analytically, so I resort to a numerical analysis of the policy parameters space as in Ascari and Sbordone (2014). I use the same grid of parameters as in the optimised monetary policy exercise laid out above and I check for which combination of ρ_π and ρ_y in the benchmark monetary policy rule Eq. 2.66 the system is determined. This procedure can be summarised graphically in a X-Y plot.

Figure 2.8: Determinacy regions of a standard Taylor rule

Base model and with the nominal debt channel. Eq. 2.66 with $\rho_{ph} = 0$

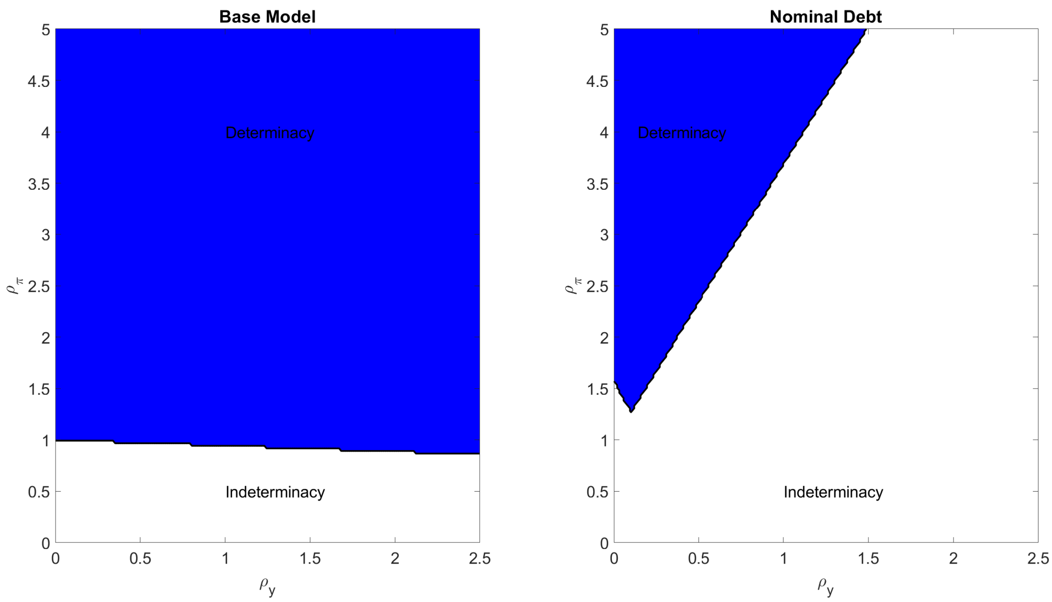

Fig. 2.8 shows how the backward looking benchmark Taylor rule in Eq. 2.66 produces similar results as in the textbook three equations model (Galí, 2009, Chapter 4). That is because it captures the Taylor principle, the Central Bank has to react aggressively to inflation and the weight on output does not matter that much as long as the weight on inflation is above one. This finding is in keeping with the literature on Taylor rules and the inflation/output. [31]

Nominal debt greatly reduces the region in which the model is stable. The Central

[30] In mainstream DSGE modelling determinacy is typically enforced by selecting a parametrisation which yields a stable equilibrium - e.g. one that satisfies the Blanchard and Kahn's conditions (Farmer, 2020).

[31] For a textbook analysis see: Woodford (2003b)

bank can deviate from strict inflation targeting and react to output as long as the coefficient assigned to inflation is very high, otherwise the system is indeterminate.

With nominal debt, the Central Bank faces a different inflation/output trade-off. Because of the inflation rate showing up in the 2.37, more inflation eases the collateral constraint contributing to more consumption fuelled by the ever higher leverage capacity.

In my model is not possible to find an optimised LATW policy rule in the nominal debt case due to a large region of indeterminacy (charted in Fig. 2.10). If debt is due in nominal terms, the space of policy parameters is significantly restricted. Importantly, this limits also the possibilities to 'lean against the wind', as targeting housing prices makes the model unstable.

I note that the determinacy region is greatly reduced with nominal debt under the two compared policy rules: with LATW and without. The key to understand the transmission mechanism in the case with nominal debt may lie within the aggregate demand: inflation deflates the debt stock allowing households to take up more leverage. On the other end, households with more debt are much more sensitive to interest rates developments and asset prices. The deleverage of borrowers causes a GDP destruction as they pay monitoring costs, hence in this instance may be difficult for the Central Bank to stabilise inflation and output at the same time.

Ultimately, it seems hard for monetary policy to stabilise house prices, output and inflation at the same time. House prices and output co-move as they both appear at the core of aggregate supply. The set-up of the production sector makes house prices and inflation inversely correlated. The degree of correlation depends also on monetary policy and the degree of interest rate smoothing.

Figure 2.9: Determinacy regions of an asset prices augmented rule

Base model and with the nominal debt channel. Eq. 2.66 with $\rho_{ph} = 0.05$

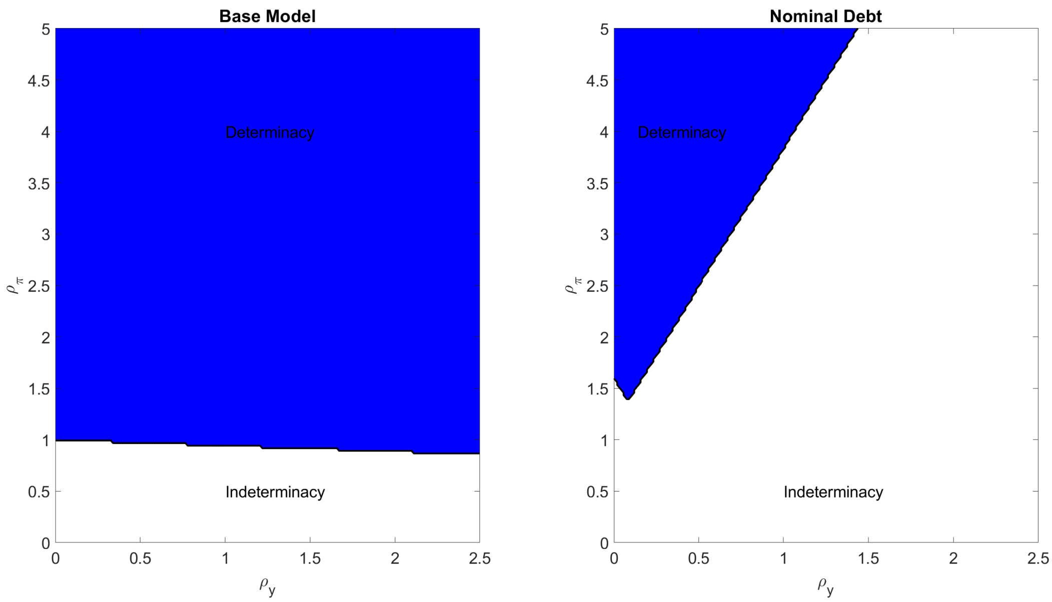

2.6.3 Model Performance: Interest Rate, House Prices and Defaults

Table 2.7: Standard Deviations of Key Variables

Data are HP-Filtered UK series for 1988-2018: π_t is CPI Inflation (ONS), p_t^h are average house prices (UK Gov) deflated with the implicit GDP deflator, and insolvencies (UK Insolvency Service) proxy the default threshold $\hat{\omega}_t$

Variable	Data	Base Model	Nominal Debt	Base Model and RW	Nominal Debt and RW	Without Insol-vencies
r_t^n	1.16	1.19	1.21	1.16	1.16	1.15
y_t	1.23	4.07	4.64	4.53	4.68	2.37
π_t	0.83	0.86	0.90	0.91	0.85	0.12
p_t^h	4.76	3.94	4.14	4.39	4.32	1.61
$\tilde{\omega}_t$	13.20	1.70	2.35	1.61	1.88	-

In this section I conduct a 'moment matching' exercise, with the aim of testing the empirical performance of the model presented above against the cyclicality of the UK credit market. For the sake of this empirical exercises, I postulate a Taylor rule in line with Eq. 2.35, with pure inflation targeting ($\rho_\pi = 0.27$ and $\rho_{rn} = 0$).

The model is successful in capturing the key statistics of the credit cycle. Table 2.7 tabulates the standard deviation of data against the model-implied ones, showing how a model with insolvencies can produce the persistence observed in actual data for the nominal interest rate, inflation and house prices. This provides an useful benchmark with respect to other similar models: Gelain et al. (2013) can match the volatility of US House Prices, Debt and Output by postulating a moving average forecast rule for expectation formation. The actual volatility of insolvencies is between 6 and 8 times greater than what the model simulates. Whilst alternatives to rational expectation can increase the volatility of defaults by adding uncertainty over the future realisation of house prices, it is doubtful that this modification would improve dramatically the currently low figure. The model shortcoming in this respect is more imputable to structural features.

Among the key variables, model-implied output is counter-factually more volatile than actual GDP, whereas insolvencies, represented by the default threshold $\tilde{\omega}$. This is to be expected given the simplistic nature of the goods market: the sum of entrepreneurial and household consumption stacks up to the model economy total output as in Eq. 2.57. In turn, consumption depends on the period-by-period configuration of the housing market, as households and entrepreneurs split their asset allocation between consumption and real estate. Hence, the volatility stems from the fact that mortgages and deposits are 1-period contracts with fully adjustable interest rate.

Insolvencies are much stickier in the model than they are in data, this also is a by-product of the selected friction. Insolvencies only materialise when households are in negative equity, meaning that they do not repay the mortgage loan if it is grater than the ex-post housing value that they receive. This is clearly not the case in reality, when defaults are not purely strategical but happen organically because of the impossibility to service a loan, this can happen on the back of a wider range of shocks, such as income or unemployment ones, which are not modelled in this paper.

In terms of cross-correlations, reported in Table 2.8, the model here presented continues to be consistent with data. Insolvencies are positively correlated with inflation and the house price index is negatively correlated with CPI inflation. My actual (hp-filtered) correlations are similar to those reported in Demary (2010), who also finds strong negative correlation between inflation and house prices in UK and in a panel with other countries. This stylised fact is captured, albeit in a weaker fashion, in the model implied cross-correlations.

Additional simulations show how the sign of the correlation between the inflation rate and output is determined by the interest smoothing parameter ρ_{rn}. If the Central Bank smooths interest rates, output and goods prices co-move, if the weight to interest rate smoothing is lower, then the model-implied cross-correlations for inflation rate, GDP and house prices are much closer to actual ones. If the Central Bank reacts to inflation more aggressively, it front-loads the output decline, whereas stickiness in the credit market provides for a more sluggish adjustment of house prices.

This may signify that a backward looking auto-correlated Taylor rule (as in Kapetanios et al. (2019)) is not a good representation of the Central Bank's response function in UK, for the reasons recalled in Section 2.2.6. Hence, the empirical correlations would support the Goodhart (2005) point of view that the high auto-correlation of the interest rate is in fact due to statistical mis-specification and the Bank of England is less concerned with interest rate smoothing than it appears from ex-post naive regressions.

Table 2.8: Credit Variables Model and Empirical Correlations

Model base calibration. UK data are for 1988-2018: π_t is CPI Inflation (ONS), p_t^h are average house prices (UK Gov) deflated with the implicit GDP deflator, and insolvencies (UK Insolvency Service) proxy the default threshold $\hat{\omega}_t$

Model

Base Model

$$
\begin{array}{c} \\ y_t \\ \pi_t \\ p_t^h \\ \tilde{\omega}_t \end{array}
\begin{array}{c}
\begin{array}{cccc} y_t & \pi_t & p_t^h & \tilde{\omega}_t \end{array} \\
\begin{bmatrix}
1 & -0.523 & 0.80 & -0.53 \\
- & 1 & -0.21 & -0.43 \\
- & - & 1 & -0.54 \\
- & - & 0 & 1
\end{bmatrix}
\end{array}
$$

Nominal Debt

$$
\begin{array}{c} \\ y_t \\ \pi_t \\ p_t^h \\ \tilde{\omega}_t \end{array}
\begin{array}{c}
\begin{array}{cccc} y_t & \pi_t & p_t^h & \tilde{\omega}_t \end{array} \\
\begin{bmatrix}
1 & -0.44 & 0.80 & -0.67 \\
- & 1 & -0.10 & 0.47 \\
- & - & 1 & -0.60 \\
- & - & 0 & 1
\end{bmatrix}
\end{array}
$$

Data

1-Sided HP-Filter

$$
\begin{array}{c} \\ y_t \\ \pi_t \\ p_t^h \\ \tilde{\omega}_t \end{array}
\begin{array}{c}
\begin{array}{cccc} y_t & \pi_t & p_t^h & \tilde{\omega}_t \end{array} \\
\begin{bmatrix}
1 & -0.52 & 0.69 & -0.35 \\
- & 1 & -0.55 & 0.50 \\
- & - & 1 & -0.19 \\
- & - & 0 & 1
\end{bmatrix}
\end{array}
$$

2-Sided HP-Filter

$$
\begin{array}{c} \\ y_t \\ \pi_t \\ p_t^h \\ \tilde{\omega}_t \end{array}
\begin{array}{c}
\begin{array}{cccc} y_t & \pi_t & p_t^h & \tilde{\omega}_t \end{array} \\
\begin{bmatrix}
1 & -0.29 & 0.66 & -0.37 \\
- & 1 & -0.36 & 0.47 \\
- & - & 1 & -0.38 \\
- & - & 0 & 1
\end{bmatrix}
\end{array}
$$

The points above hold when considering the response on impulse. In Fig. 2.10 I compare the IV-VAR impulse responses showed at the beginning of this paper with the theoretical ones produces by the baseline model. I find that the response on impulse of GDP, house prices and real debt are similar in magnitude and overlap to the VAR confidence bands. The actual response in insolvencies is much greater than the model implied one, this again is due to the fact that the DSGE models only strategic defaults and not unexpected ones. Moreover, in this paper mortgages figure as 1-period adjustable rate bonds (issued by the households), this is also very far from the reality of actual mortgages, which are fully-amortising contracts. Hence, in case of defaults, borrowers fail to re-pay the outstanding amount.

Figure 2.10: DSGE versus IV-VAR

Backward looking Taylor rule with $\rho_{rn} = 0.9$, $\rho_{\pi} = 1.5$ and $\rho_{y} = 0.1$ as in Kapetanios et al. (2019)

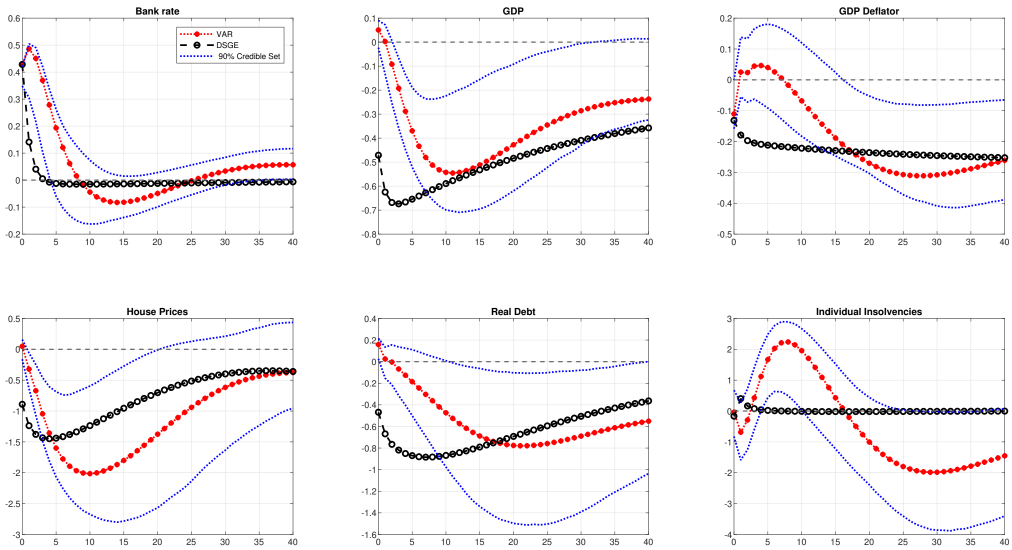

2.7 Directions for future research

This paper contributes to the literature on house prices targeting in New Keynesian models. It does so adopting a tractable New Keynesian framework. The focus on the extensive margin allows the model to capture the main stylised facts of the credit cycle, while falling short on others.

The workhorse model proposed above could fit reasonably well the empirical moments of the variables of interest but it failed to produce a sizeable response on impulse of the insolvency rate. I speculate that this might be due to the length of mortgage contracts, so one interesting expansions could be factoring in longer and risky mortgage contracts (for reference Bluwstein et al. (2018) has long household loans but no defaults). Whilst this modification might help better matching the data moment, does not change the fact that in the model presented above the transmission mechanism stems from a 'levered' aggregate demand, which makes mortgagors sensitive to interest rate movements.

Adam et al. (2011) show that house price dynamics and sentiment are difficult to reconcile with the rational expectation setting and indeed there is a developed literature of papers use the expectations channel to explain housing booms (few among many Lambertini et al. (2017a); Granziera and Kozicki (2015); Gelain et al. (2013)). Different formalisation of expectation-formation might provide a valuable environment to glean insights on the extent in which house price targeting helps the economy. In different modelling contexts, anchoring expectations around house prices using monetary policy instruments (besides say pure macroprudential policy) could represent a compelling hypothesis specifically because agents would have to infer the workings of the economy from its observed realisation.

Following this line of reasoning, one might also model Central Banks as not fully rational. For instance, the necessity for interest rate smoothing may naturally arise in a model in which the Central Bank learns about the economy and wants to avoid abrupt interest rate movements. Or, alternatively, reputational costs may be associated with a high volatility of the policy rate, hence providing a more compelling case for smoothing (Murray, 2012). Models with different form of expectations may offer a richer environment to test Taylor and LATW rules in a more meaningful way.

Even tough I calibrate my model on UK data, I find difficult to use an off-the-shelf Taylor rule for a baseline 'realistic' specification. In my first Chapter I dealt with the problem of endogeneity of monetary policy using an external instrument series, or trying to isolate the 'systematic component' of monetary policy with sign restrictions. In this Chapter I deal with monetary policy rules using them in a formulaic way, in accordance with a wider literature (recalled above) and operational DSGE example (the Swedish Riskbank Adolfson et al. (2011) use a rule similar to 2.35).

To overcome this problem I end up estimating an array of different rules and I cycle through different parameter combinations in my efficient frontier exercise. Finding the 'optimal' monetary policy rule is outside the scope of this chapter but, as a general point, I agree with Curdia et al. (2011) in saying that DSGE modelling often features too formulaic monetary policy rules.

This paper touches very slightly on welfare due to the computational complications associated with this exercise. I only compute a very simple monetary policy optimised rule without engaging in a complete analysis on optimal monetary policy with borrowing households and saving entrepreneurs. A proper optimal analysis is a useful extension but beyond the scope of this Chapter, as it entail setting up a quite complex simulation exercise. It could be interesting to see what the optimal level of default is. This would be particularly relevant if coupled with the extension above: how strongly the central bank can react to inflation by raising the interest rate if defaults are more volatile and sticky than they are in the present model?

Another point of attention is the nominal debt channel. Fisherian deflation of real debts appears to be quantitatively important as far as the response on impulse is concerned, improving borrower's credit quality represented by the default threshold. Nominal debt has also important implications for the stability of the system: with loans and deposits paid in nominal terms the policy space is greatly reduced and LATW is infeasible (it generates instability).

2.8 Conclusions

In this paper I model the credit view of monetary policy transmission in a New-Keynesian setting. There are two key agents: wealthy hand-to-mouth households (mortgagors) and firms (depositors), with a bank extending loans out of deposits wile setting aside regulatory capital. A costly state verification friction creates a modelling environment where credit is available up to an endogenous LTV ratio. This generates an acceleration mechanism, as aggregate demand and housing investments are sensitive to external disturbances that propagate them on impulse. This model captures also the lending channel, as $L \neq D$ and banks calibrate their leverage ratio as to offset their capital management costs with the net interest margin. So in periods of subdued loan demand, banks have to scale back their lending to mitigate losses.

A big advantage of the model here presented is its tractability: it is in 23 log-linearised equations and thus simple enough to allow for a tracing of the transmission mechanism. Form the frictions just described stem a collateral channel, a debt-deflation channel and a bank lending channel that enrich the dynamics of inter-temporal substitution and sticky prices common to workhorse New-Keynesian models. Through simple modelling devices, I am able to pin down the key stylised facts of the credit channel of monetary policy: the credit aggregates are very persistent and household finances provide a channel of amplification for the economy. Specifically, my model shifts the centre of the transmission mechanism from inter-temporal substitution to household leverage.

This model has policy implications: LATW monetary policy is consistent with lower volatility of a broadly defined central bank loss function when output has a positive weight, but is quantitatively similar to output targeting. This result may stem from the inner workings of the model: as house prices inform both aggregate demand and supply, they sit at the core of the transmission mechanism. Interest rate smoothing improves the efficient frontier, consistently with the wider DSGE literature.

Moreover, a credit friction that only envisions strategic default would tend to underestimate sudden insolvencies, which are linked to asset prices and output. Insolvencies are between 5.5 and 7 times more volatile of different declinations of the model tested above, and inversely correlated with house prices, so policies directed to house prices will impact the aggregate demand affecting the leverage of representative agents.

When household leverage depends also on inflation as agents repay the mortgage contract in nominal terms, debt deflation reduces the money value of the loans. This model configuration results in a reduced policy-parameters space, such as that the central bank has to react more vigorously to inflation in order to ensure dynamic determinacy. If mortgagors and depositors compete for the same real-estate, it becomes challenging for monetary policy to stabilise output, house prices and inflation at the same time.

Together, these results go in the direction of favouring a simpler monetary policy approach based on a regular Taylor rule.

Appendix B

Appendix to Chapter 2

B.1 Stylised Facts on the Credit Cycle

- Lower Frequency than the Business Cycle (Aikman et al., 2015);

Figure B.1: 1-sided HP-Filter UK GDP and Household Debt

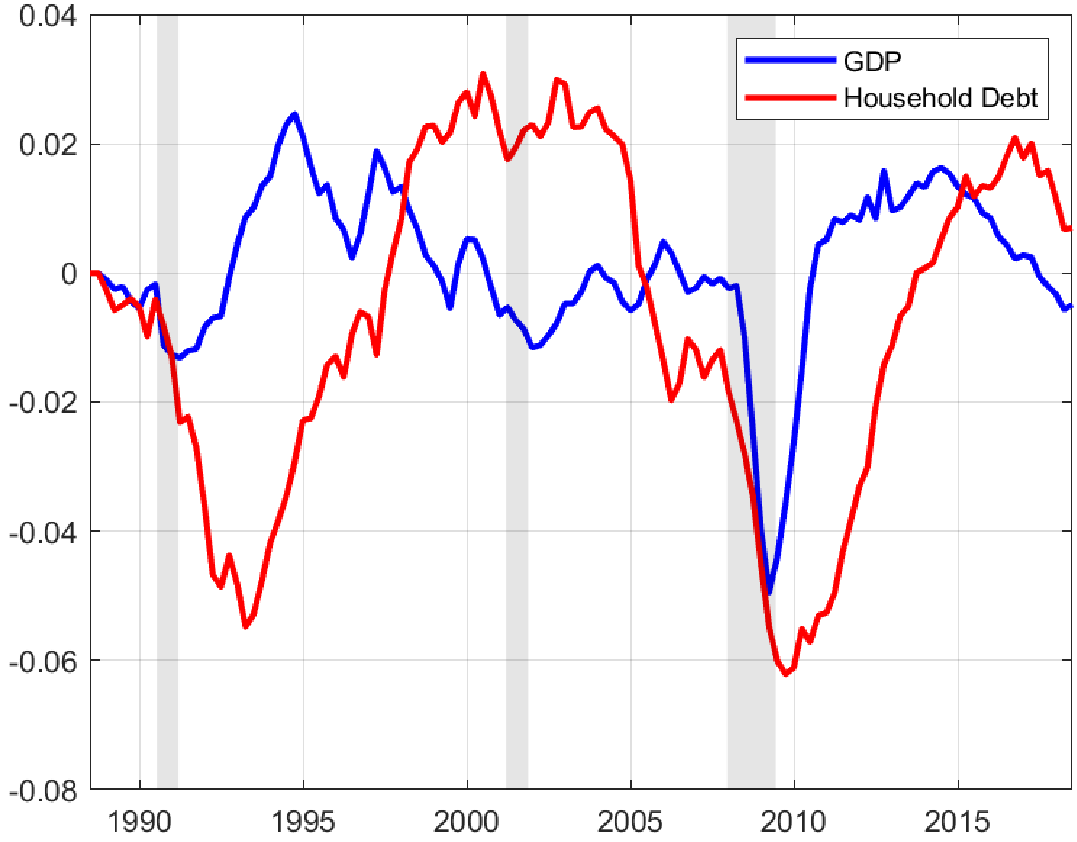

- Tightening of credit conditions accompanies recessions(Christiano et al., 2014);

Figure B.2: Mean above the Median LTV (BoE)

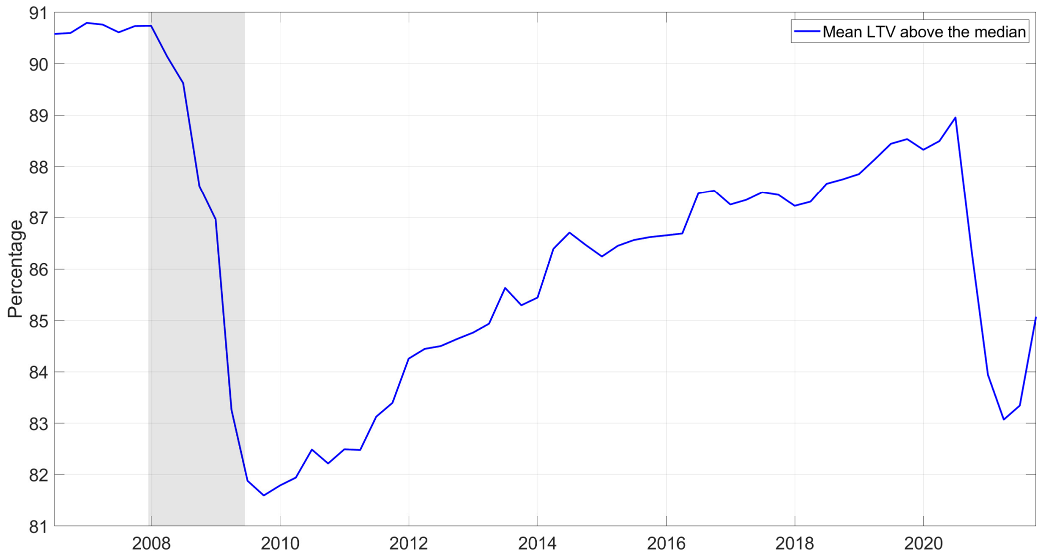

Figure B.3: UK Mortgage Approvals (LHS), High-Low LTV IR Spread (RHS)

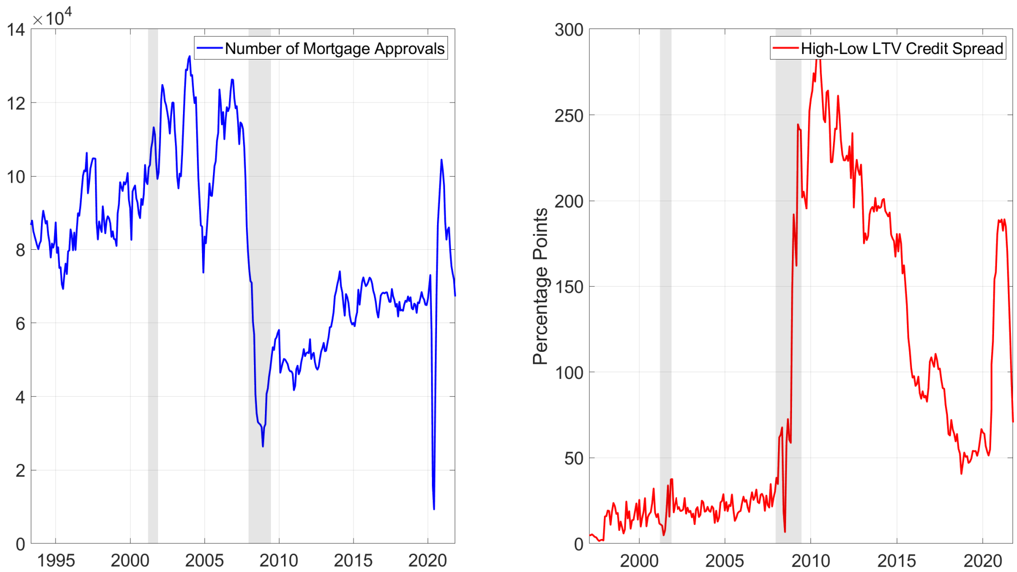

- Defaults are cyclical and associated with credit market conditions (Bernanke et al., 1999);

Figure B.4: Total Credit and Default Rates in US and UK

Total Credit comprises financing from all sources and is calculated by the Bank of International Settlements 'BIS' following the framework of the System of National Accounts 2008. Delinquent loans from US are 30-days past due disclosed by the Federal Reserve, and for UK the rate of insolvency per 100,000 inhabitants calculated by the Insolvency Service

.

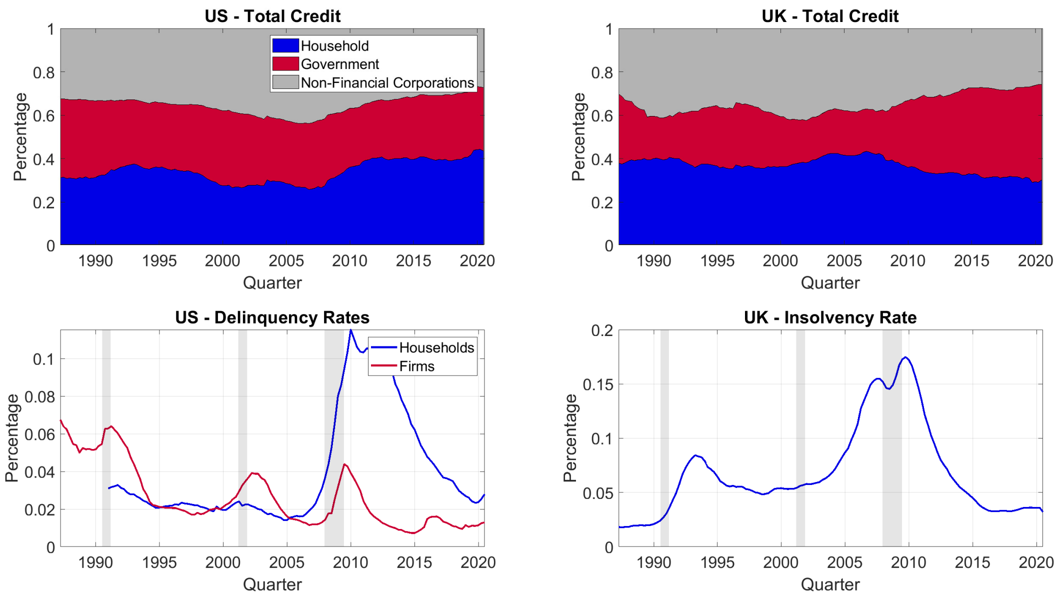

B.2 Steady State

To simplify the steady state I posit that the $\sigma = \phi = \chi = 1$, which is equivalent to assume a logarithmic utility function for both households and entrepreneurs.

$$P = 1 \tag{B.1}$$

$$1 + RR = 1/\gamma \tag{B.2}$$

$$H = H^e + H^h = 1 \tag{B.3}$$

$$\Lambda^c = \frac{\gamma - \beta}{C^h} \tag{B.4}$$

$$\frac{C^h}{P^h H^h} = \Big[1 - \beta(1 - \delta^h) + \beta\mu G[\overline{\omega}] - \Phi[\overline{\omega}](\gamma - \beta)\Big] / (\varsigma\beta) = \zeta_1 \tag{B.5}$$

$$\frac{C^e}{P^h H^e} = \frac{H^n}{H^e}\nu + \frac{H^h}{H^e}\frac{\Phi[\overline{\omega}]}{1 + RR}(1 - \delta^h)(1 - \xi) - \delta^h = \zeta_2 \tag{B.6}$$

$$\frac{H^e}{H^n} = \frac{\gamma\nu}{1 - \gamma + \gamma\delta^h} = \zeta_3 \tag{B.7}$$

$$L = (1 - \delta^h)\frac{\Phi(\overline{\omega})}{(1 + RR)} P^h H^h \tag{B.8}$$

$$\frac{D}{L} = 1 - \xi, \quad \frac{K}{L} = \xi \tag{B.9}$$

$$Y = C^h + C^e + \delta K \tag{B.10}$$

$$\frac{Y}{P^h H^h} = \frac{C^h}{P^h H^h} + \frac{C^e}{P^h H^h} + \delta\xi\frac{L}{P^h H^h} = \Xi \tag{B.11}$$

$$\frac{Y}{P^h H^h} = \zeta_1 + \zeta_2 + \delta \xi \Phi(\hat{\omega}) \frac{\gamma}{(\gamma + 1)} = \Xi \tag{B.12}$$

Figure B.5: Steady State Clearing

Clearing in the housing market (left pane) and in the loans market (right pane). Solid red lines show the steady state with banking capital requirements and dashed black lines without them

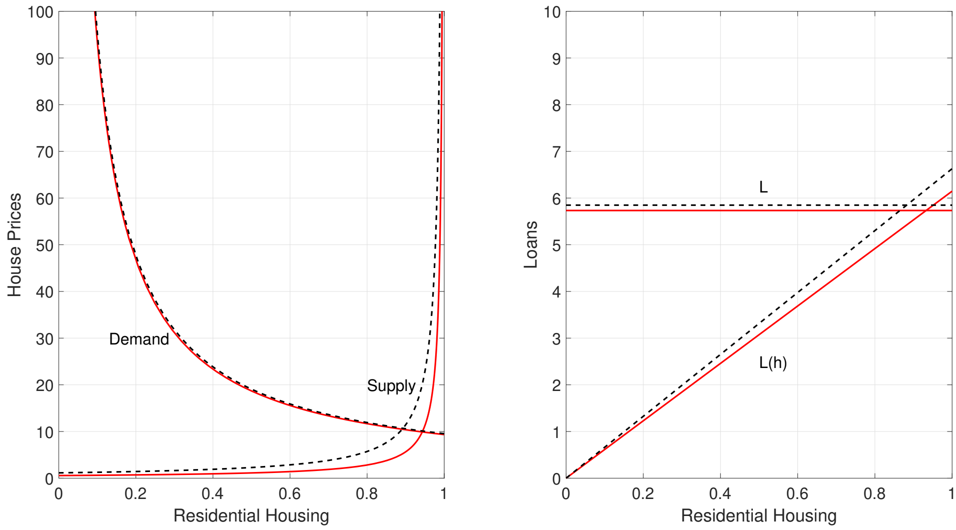

B.3 Log-Linearisation

B.3.1 Households Budget Constraint

The Household Budget Constraint is:

$$H^h_{t+1}P^h_t + C^h_t - L_{t+1} = -\Big(H^h_t P^h_t \mu G(\hat{\omega}_t) + L_t(R_{L,t}+1)\Big) + (1-\delta^h)H^h_t P^h_t + N^c_t W^c_t + N^h_t W^h_t + F_t \tag{B.13}$$

We can substitute in it the following definitions before log-linearising:

Labour demand curves:

$$N^h_t W^h_t = (1-\nu)H^n_t P^h_t \tag{B.14}$$

$$N^c_t W^c_t = Y_t / X_t \tag{B.15}$$

Households profit from retail goods producers:

$$F_t = Y_t - \frac{Y_t}{X_t} \tag{B.16}$$

$$H^h_{t+1}P^h_t + C^h_t - L_{t+1} = -\Big(H^h_t P^h_t \mu G(\hat{\omega}_t) + L_t(R_{L,t}+1)\Big) + (1-\delta^h)H^h_t P^h_t + \frac{Y_t}{X_t} + (1-\nu)H^n_t P^h_t + Y_t - \frac{Y_t}{X_t} \tag{B.17}$$

$$H^h_{t+1}P^h_t + C^h_t - L_{t+1} = -\Big(H^h_t P^h_t \mu G(\hat{\omega}_t) + L_t(R_{L,t}+1)\Big) + (1-\delta^h)H^h_t P^h_t + (1-\nu)H^n_t P^h_t + Y_t \tag{B.18}$$

To log-linearise the household flow of funds above we use the Uhlig's method.

$$H^hP^h(1+h^h_{t+1}+p^h_t)+C^h(1+c^h_t)-L(1+l_{t+1}) = -\Big[H^hP^h\mu G(\overline{\omega})(1+h^h_t+p^h_t+\eta\,\omega_t)+LRR(1+l_t+$$
$$+L(1+l_t)\big]+(1-\delta^h)H^hP^h(1+h^h_t+p^h_t)+(1-\nu)H^nP^h(1+h^n_t+p^h_t)+Y_t(1+y_t) \tag{B.19}$$

Subtracting the steady state values:

$$H^hP^h(h^h_{t+1}+p^h_t)+C^h(c^h_t)-L(l_{t+1}) = -\Big[H^hP^h\mu G(\overline{\omega})(h^h_t+p^h_t+\tilde{\eta}\,\omega_t)+LRR(l_t+r^l_t)$$
$$+L(l_t)\big]+(1-\delta^h)H^hP^h(h^h_t+p^h_t)+(1-\nu)H^nP^h(h^n_t+p^h_t)+Y_t(y_t) \tag{B.20}$$

Then collecting L, the resource constraint becomes:

$$C^h(c^h_t) = H^hP^h(h^h_{t+1}+p^h_t)+L\left[l_{t+1}-l_t-RR(l_t+r^l_t)\right]-\left[H^hP^h\mu G(\overline{\omega})(h^h_t+p^h_t+\tilde{\eta}\,\omega_t)\right]$$
$$+(1-\delta^h)H^hP^h(h^h_t+p^h_t)+(1-\nu)H^nP^h(h^n_t+p^h_t)+Y_t(y_t) \tag{B.21}$$

Dividing both sides by H^hP^h yields and using the definitions from L and ζ_1 in Section B.2:

$$\zeta_1C^hc^h_t = (h^h_{t+1}+p^h_t)+(1-\delta^h)\frac{\Phi(\overline{\omega})}{1+RR}\left[l_{t+1}-l_t-RR(l_t+r^l_t)\right]-\left[\mu G(\overline{\omega})(h^h_t+p^h_t+\tilde{\eta}\,\omega_t)\right]$$
$$+(1-\delta^h)(h^h_t+p^h_t)+(1-\nu)\frac{H^n}{H^h}(h^n_t+p^h_t)+\frac{Y_t}{H^hP^h}y_t \tag{B.22}$$

Re-arranging:

$$\zeta_1C^hc^h_t = \frac{Y_t}{H^hP^h}y_t+h^h_{t+1}+p^h_t-\mu G(\overline{\omega})(h^h_t+p^h_t+\tilde{\eta}\,\omega_t)+(1-\delta^h)(h^h_t+p^h_t)+$$
$$(1-\nu)\frac{H^n}{H^h}(h^n_t+p^h_t)+(1-\delta^h)\frac{\Phi(\overline{\omega})}{1+RR}\left[l_{t+1}-l_t-RR(l_t+r^l_t)\right] \tag{B.23}$$

B.3.2 Housing Demand

The Housing demand FOC is the following:

$$\varsigma\beta\frac{C_t^{h\,\sigma}}{H_t^{h\,\chi}P_t^h} = \frac{P_{t-1}^h}{P_t^h}\left(\frac{C_t^h}{C_{t-1}^h}\right)^{\sigma} - \beta(1-\delta^h) + \beta\mu G[\omega_t] - (1-\delta^h)C_t^{h\,\sigma}\lambda_{t-1}^c\Phi[\omega_t] \quad \text{(B.24)}$$

Using the Uhlig's method of approximation around the steady state:

$$\varsigma\beta\frac{C^{h\,\sigma}}{H^{h\,\chi}P^h}(1+c_t^h-\chi h_t^h-p_t^h) = [1+p_{t-1}^h-p_t^h+\sigma(c_t^h-c_{t-1}^h)] + \beta\mu G[\overline{\omega}](1+\eta\omega_t)$$
$$-(1-\delta^h)C^{h\,\sigma}\Lambda^c\Phi[\overline{\omega}](1+\sigma c_t^h+\lambda_t^c+\iota\omega_t) \quad \text{(B.25)}$$

Subtracting the steady state from both sides and setting $\sigma=\chi=1$ as per baseline calibration yields the following linearised equation:

$$\varsigma\beta\zeta_1(c_t^h-h_t^h-p_t^h) = (p_{t-1}^h-p_t^h+c_t^h-c_{t-1}^h)+\beta\mu G[\overline{\omega}](\eta\omega_t)-(1-\delta^h)C^h\Lambda^c\Phi[\overline{\omega}](c_t^h+\lambda_{t-1}^c+\iota\omega_t) \quad \text{(B.26)}$$

Using the same method, I derive the log-linear version of the budget constraint Lagrange multiplier is:

$$\lambda_{t-1}^c = -\frac{\gamma}{\gamma-\beta}(c_{t-1}^h+r_{t-1}^l) - \frac{\beta}{\gamma-\beta}(r_t^l-r_{t-1}^l-c_t^h) \quad \text{(B.27)}$$

B.3.3 Entrepreneurs Flow of Funds

The Entrepreneurs resource constraint is:

$$H_{t+1}^eP_t^h+C_t^e+N_t^cW_t^c+N_t^hW_t^h+D_{t+1} = (1-\delta^h)H_t^eP_t^h+\frac{Y_t}{X_t}+(1+R_{t-1}^d)D_t+H_t^nP_t^h \quad \text{(B.28)}$$

Substituting the labour demand FOCs inside it:

$$H^e_{t+1}P^h_t + C^e_t + \frac{Y_t}{X_t} + (1-\nu)H^n_t P^h_t + D_{t+1} = (1-\delta^h)H^e_t P^h_t + \frac{Y_t}{X_t} + (1+R^d_{t-1})D_t + H^n_t P^h_t \tag{B.29}$$

Simplifying:

$$H^e_{t+1}P^h_t + C^e_t - \nu H^n_t P^h_t + D_{t+1} = (1-\delta^h)H^e_t P^h_t + (1+R^d_{t-1})D_t \tag{B.30}$$

Linearising around the steady state yields:

$$\begin{aligned} H^e P^h(1+h^e_{t+1}+p^h_t) + C^e(1+c^e_t) - \nu H^n P^h(1+h^n_t+p^h_t) + D(1+d_{t+1}) = \\ (1-\delta^h)H^e P^h(1+h^e_t+p^h_t) + D(1+d_t) + RRD(1+r^d_{t-1}+d_t) \end{aligned} \tag{B.31}$$

Subtracting the steady state:

$$\begin{aligned} H^e P^h(h^e_{t+1}+p^h_t) + C^e(c^e_t) - \nu H^n P^h(h^n_t+p^h_t) + D(d_{t+1}) = \\ (1-\delta^h)H^e P^h(h^e_t+p^h_t) + D(d_t) + RRD(r^d_{t-1}+d_t) \end{aligned} \tag{B.32}$$

Dividing by $H^e P^h$, using the steady state definition of ζ_2 and re-organising the equation:

$$\begin{aligned} h^e_{t+1} + p^h_t + \zeta_2 c^e_t - \nu\frac{H^n}{H^e}(h^n_t+p^h_t) = \\ (1-\delta^h)(h^e_t+p^h_t) + \frac{D}{H^e P^h}[d_t - d_{t+1} + RR(r^d_{t-1}+d_t)] \end{aligned} \tag{B.33}$$

Substituting $(1-\xi)L$ for D and the steady state for L:

$$\zeta_2 c_t^e = \nu \frac{H^n}{H^e}(h_t^n + p_t^h) - h_{t+1}^e - p_t^h + (1-\delta^h)(h_t^e + p_t^h) + (1-\xi)(1-\delta^h)\frac{H^h}{H^e}\frac{\Phi(\overline{\omega})}{1+RR}[d_t - d_{t+1} + RR(r_{t-1}^d + d_t)] \tag{B.34}$$

B.4 A Model without Defaults

The following specification is along the lines of Iacoviello (2005), featuring an exogenous collateral constraint (hence the LTV Φ does not depend on $\hat{\omega}$).

B.4.1 Household Problem

The collateral constrain now embeds a fixed LTV ratio Φ:

$$\frac{L_{t+1}}{E_t[\Pi_{t+1}]} = (1-\delta^h)\Phi E_t\left[\frac{H_t P_{H,t}}{(1+R_t^l)}\right] \tag{B.35}$$

The demand for housing is:

$$\varsigma\beta H_t^{-\chi} = P_{t-1}^h \lambda_{t-1} - (1-\delta^h])P_t^h(\Phi\lambda_t^c + \beta\lambda_t) \tag{B.36}$$

B.4.2 Log-Linear

The borrowing constraint is:

$$l_{t+1} = E_t[h_{t+1}^h + p_{t+1}^h] - r_t^l \tag{B.37}$$

The log-linear housing demand is:

$$\varsigma\beta\zeta_1(c_t^h - h_t^h - p_t^h) = p_{t-1}^h - p_t^h + c_t^h - c_{t-1}^h + (1-\delta^h)(\gamma-\beta)(-\frac{\gamma}{\gamma-\beta}(c_t^h + r_t^l) - \frac{\beta}{\gamma-\beta}(r_{t+1}^l - r_t^l - c_{t+1}^h) - c_t^h) \tag{B.38}$$

where ζ_1 is:

$$\zeta_1 = \beta\left[\delta^h(1-\Phi) + \Phi - 1\right] + \gamma(\delta^h - 1)\Phi + 1 \tag{B.39}$$

B.5 Nominal Debt

In the case in which mortgage loans are not indexed to inflation, mortgagors pay back their 1-period ARM in money terms as well as depositors, who receive the nominal proceeds of their bank savings.

I show how the derivations of first order conditions change in this section, alongside with providing new log-linearised equations. This happens by introducing the budget constraints the gross inflation rate $\Pi_t = \frac{P_t}{P_{t-1}}$.

B.5.1 Households

The flow budget constraint is:

$$H_{t+1}^h P_t^h + C_t^h - L_{t+1} = -\left(H_t^h P_t^h \mu G(\hat{\omega}_t) + \frac{L_t}{\Pi_t}(R_{L,t}+1) \right) + (1-\delta^h) H_t^h P_t^h + N_t^c W_t^c + N_t^h W_t^h + F_t \tag{B.40}$$

The Collateral Constraint is:

$$\frac{L_{t+1}}{E_t[\Pi_{t+1}]} = (1-\delta^h) E_t \left[\Phi[\hat{\omega}_{t+1}] \frac{H_t P_{H,t}}{(1+R_t^l)} \right] \tag{B.41}$$

Housing Demand is the FOC with respect of H_t:

$$\varsigma \beta \frac{C_t^{h\,\sigma}}{H_t^{h\,\chi} P_t^h} = \frac{P_{t-1}^h}{P_t} \left(\frac{C^h{}_t}{C_{t-1}} \right)^{\sigma} - \beta(1-\delta^h) + \beta \mu G[\omega] - (1-\delta^h) C^{h\,\sigma}_t \lambda_{t-1}^c \Phi[\omega] \tag{B.42}$$

B.5.2 Entrepreneurs

$$H_{t+1}^e P_t^h + C_t^e + N_t^c W_t^c + N_t^h W_t^h + D_{t+1} = (1-\delta^h) H^e P_t^h + \frac{Y_t}{X_t} + (1+R_{t-1}^d)\frac{D_t}{\Pi_t} + H_t^n P_t^h \tag{B.43}$$

B.5.3 Log-Linear

The household's flow of funds becomes:

$$
\begin{aligned}
c_t^h \zeta_1 = y_t \frac{Y}{H^h P^h} + \Big[(1-\delta^h)(h_t^h + p_t^h) - h_{t+1}^h - p_t^h - \mu G[\overline{\omega}](\eta \omega_t + h_t^h + p_t^h)\Big] + \\
\frac{H^n}{H^h}(h^n + p_t^h)(1-\nu) + \frac{\Phi[\overline{\omega}]}{1+RR}(1-\delta^h)\Big[(l_{t+1} - l_t + \pi_{t+1} - (l_t + r_t^l - \pi_{t+1})RR\Big]
\end{aligned} \tag{B.44}
$$

The borrowing constraint is:

$$
l_{t+1} = E_t[h_{t+1}^h + p_{t+1}^h] + E_t[\omega_{t+1}]\overline{\omega}\iota - r_t^l t + E_t[\pi_{t+1}] \qquad \text{(Borrowing Constraint)}
$$

Housing demand:

$$
\begin{aligned}
&\varsigma \beta \zeta_1 \left(c_t^h - h_t^h - p_t^h\right) = c_t^h + p_{t-1}^h - p_t^h - c_{t-1}^h + \eta\, \omega_t\, G[\overline{\omega}]\mu\beta \\
&+ \left(1-\delta^h\right)(\gamma-\beta)\left[c_t^h - c_{t-1}^h + \omega_t \iota - \frac{\gamma}{\gamma-\beta}\left(r_{t-1}^l + c_{t-1}^h\right) - \frac{\beta}{\gamma-\beta}\left(r_t^l - r_{t-1}^l - c_t^h - \pi_{t+1}\right)\right]\Phi[\overline{\omega}]
\end{aligned} \tag{B.45}
$$

Entrepreneur Flow of Funds:

$$
\begin{aligned}
\zeta_2 c_t^e = \left(h_t^n + p_t^h\right)\frac{H^n}{H^e}\nu + \left[-h_t^e - p_t^h + \left(1-\delta^h\right)\left(p_t^h + h_{t-1}^e\right)\right] \\
+ \frac{H^h}{H^e}\frac{\Phi[\overline{\omega}]}{1+RR}(1-\delta^h)(1-\xi)\left[d_{t-1} - d_t - \pi_{t+1} + RR\left(d_{t-1} - \pi_{t+1} + r^d{}_{t-1}\right)\right]
\end{aligned} \tag{B.46}
$$

And its Euler Equation:

$$
r_t^n = c_{t+1}^e - c_t^e \tag{B.47}
$$

B.6 A Partial Equilibrium Model

In this section I state a partial equilibrium model. It is partial equilibrium because I only derive an IS equations, but the model can be expanded to three equations (IS, Phillips Curve and House Prices Phillips Curve), plus a monetary policy rule.

The utility function is:

$$U(C_t, N_t, H_t) = \frac{C_t^{(1-\sigma)}}{1-\sigma} + \varsigma ln(H_t^h) - \frac{N_t^\eta}{\eta} \tag{B.48}$$

The budget constraint is analogous to the one presented above, albeit simplified:

$$C_t + P_t^h H_t^h + R_{t-1}^l \frac{L_{t-1}}{\pi_t} = L_t + W_t N_t + P_t^h H_{t-1}^h \tag{B.49}$$

The representative household can borrow up to a certain fraction Φ (loan-to-value-ratio) of its housing value. The borrowing constraint is:

$$L_t = \Phi \frac{P_{t+1}^h H_t^h \pi_{t+1}}{R_t^l} \tag{B.50}$$

Substituting the borrowing constraint into the budget constraint yields:

$$C_t + P_t^h (H_t^h - H_{t-1}^h) + \Phi H_{t-1} P_t^h = \Phi \frac{P_{t+1}^h H_t^h \pi_{t+1}}{R_t^l} + W_t N_t \tag{B.51}$$

B.6.0.1 First Order Conditions

Maximising the utility function under the budget constraint yields this FOC for house prices demand.

$$\varsigma \frac{C_t^\sigma}{P_t^h H_t^h} = 1 - \Phi \pi_{t+1}^h \frac{\pi_{t+1}}{R_t^l} + \beta \left(\frac{C_t}{C_{t+1}} \right)^\sigma \pi_{t+1}^h (\Phi - 1) \tag{B.52}$$

which log-linearises in:

$$\tilde{c}_t = \tilde{c}_{t+1} - \frac{1}{\sigma} \frac{\Phi}{1-\Phi} (\tilde{r}_t - \tilde{\pi}_{t+1}) - \frac{1}{\sigma} \frac{1}{1-\Phi} \tilde{\pi}_{t+1}^h \tag{B.53}$$

We can reduce it even further by assuming a log-linear Cobb-Douglas production function with labour and real-estate held by entrepreneurs as its arguments.

$$\tilde{y}_t = \tilde{a}_t + (1-\alpha)\tilde{n}_t + \alpha\tilde{h}^e_{t-1} \tag{B.54}$$

Given that we postulate that consumption equals output $y_t = c_t$ and technology follows an AR1 process $a_t = \rho a_{t_1} + \varepsilon_t$, we can subtract from both sides the output implied under fixed prices and define output gap as $x_t = y_t - y^f_t$. The 'levered IS' is thus the following:

$$\tilde{x}_t = \tilde{x}_{t+1} - \frac{1}{\sigma}\frac{\Phi}{1-\Phi}(\tilde{r}_t - \tilde{\pi}_{t+1}) - \frac{1}{\sigma}\frac{1}{1-\Phi}\tilde{\pi}^h_{t+1} + (\rho-1)\tilde{a}_t \tag{B.55}$$

This dynamic IS is the analogue of the standard 3-equations model presented in Walsh (2010):

$$\tilde{x}_t = \tilde{x}_{t+1} - \frac{1}{\sigma}(\tilde{r}_t - \tilde{\pi}_{t+1}) + (\rho-1)\tilde{a}_t \tag{B.56}$$

The difference is that in the levered IS there is more amplification owning to the fact that households accumulate housing value, hence it is more elastic to the interest rate.

CHAPTER THREE

THE NASH WAGE ELASTICITY AND ITS BUSINESS CYCLE IMPLICATIONS

3.1 Introduction

The existing literature that seeks to estimate the degree of wage rigidity and to assess its importance in business cycles has so far failed to reach consensus, in part due to a fundamental obstacle to inference. The obstacle is that it has been difficult to argue that particular measures and estimation approaches correctly evaluate the rigidity of the marginal cost of labor except under strong and controversial assumptions about which theoretical model describes the labor market. As we discuss below, this obstacle has been a significant barrier not only for structurally estimated models featuring wage rigidity, but also for the reduced form literature on wage cyclicality. As such, without agreement on the preferred model of the labor market overall, it has been not been possible to agree on the level of wage rigidity supported by the data.

In an effort to surmount this obstacle, this paper develops and applies a new semi-structural method to estimate real wage rigidity and assess its business cycle implications. In doing so, we provide new evidence that wage rigidity is highly quantitatively important in US data, and plays a dominant role in explaining the volatility of unemployment over the business cycle. Our approach relies on an equation we derive relating the wage to other aggregate variables, which we show commonly holds across a very large class of different search and matching (SAM) models with Nash bargaining. By estimating specifications that regress actual wages on the Nash bargained wage implied by this equation, we are able to directly and straightforwardly compare US wage data to what would be implied by this large class of SAM models. The large class of models we study includes, for instance, models with many different shocks, rich firm and match heterogeneity, job-to-job transitions, nominal rigidity in goods markets and various other frictions in goods and financial markets. Our results indicate that models within this large class can only be made consistent with wage data (under conventional calibrations of other parameters) if the wages of both newly hired workers and job stayers are far more rigid than implied by Nash bargaining.

Our approach also allows us to infer the likely contribution of wage rigidity to the business cycle volatility of unemployment and to assess how far the data supports different models of rigid wages. For instance, in a simple SAM model with productivity shocks, we find that our estimated level of wage rigidity increases the volatility of unemployment more than sevenfold compared to what would occur under Nash bargaining, and can account for around half the unemployment volatility in the data. We show that our empirical estimates suggest that wages in the data are at least as rigid as in an alternating offer bargaining model based on Christiano et al. (2016).

Throughout this paper, we use the term 'wage rigidity' to represent the notion that wages do not vary with macroeconomic conditions to the extent that would be expected if they were set by Nash bargaining. We propose a new measure of wage rigidity, the Nash wage elasticity (NWE). The NWE represents the percentage increase in the cost of labor when the wage rate implied by the Nash bargaining solution increases by 1%.[1] By construction, if wages are indeed set by Nash bargaining, the NWE is equal to 1. On the other hand, if wages are very rigid compared to Nash bargaining, the NWE will be closer to zero.

Measuring wage rigidity against the benchmark of Nash bargaining is desirable, we

[1] Pissarides (2009) uses the phrase Nash wage elasticity on occasion to mean the elasticity of Nash wages with respect to productivity. To avoid confusion, we stress that we use this term to mean something completely different.

argue, because some flexible wage benchmark is required to meaningfully assess whether or not wages are rigid. That is, wages can only meaningfully be called rigid if their behavior deviates from what would be considered a flexible wage. We provide four reasons why Nash bargaining represents a logical flexible wage benchmark. First, Nash bargaining is perhaps the most common assumption used in the recent literature on unemployment over the business cycle, and so it is useful to know how far this is consistent with actual wage setting. Second, Nash bargaining is constrained efficient in the labor market in important cases, and so the NWE provides a useful yardstick of how flexible or rigid wages are likely to be compared to what would be constrained efficient. Third, as we discuss below, we show that the NWE is a strong predictor of the effects of wage rigidity on the cyclical volatility of unemployment, regardless of whether or not wages are actually set by Nash bargaining. Fourth, we show that different models without Nash bargaining, such as various rigid wage models, imply significantly different values for the NWE, and the implied NWE is less sensitive to other model assumptions aside from those about wage setting. Therefore, the NWE allows us to adjudicate which of these models are more consistent with wage data.

Our paper begins by developing and applying a semi-structural approach to estimate the NWE in US data. We derive a common wage equation that holds across a large class of models with Nash bargaining. We use this equation to impute a time series for the Nash wage from US data, without needing to adjudicate over which model in this large class corresponds to the true data generating process. We obtain estimates of the NWE by regressing measures of the actual cost of labor on the Nash wage. Across 180 regressions using various series for the cost of labor, various (or no) instruments for the Nash wage and various assumptions about the opportunity cost of employment and hiring costs, we mainly obtain NWE estimates between 0 and 0.1, indicating that wages of both job stayers and new hires are highly rigid in comparison to Nash bargaining. The data consistently favors an NWE below 0.65, except in the most extreme specifications, which use *both* the most procyclical wage series we consider (the 'user cost of labor' from the NLSY) *and* also assume high values of the opportunity cost of employment and/or very large fixed hiring costs. Intuitively, our consistently small estimates of the NWE are a consequence of the fact that the Nash wage is much more procyclical than measures of actual wages, across these many different specifications.

Next, we provide novel analytical and simulation results to show that the NWE is a strong predictor of the cyclical volatility of unemployment, across a large class of models with shocks affecting the marginal revenue product of labor (e.g. productivity or markup shocks). We derive a tight mathematical relationship between the NWE and the Fundamental Surplus, which Ljungqvist and Sargent (2017) have shown is a valuable predictor of the cyclical volatility of unemployment in many search models. When the NWE is as low as most of our empirical estimates, we show that wage rigidity amplifies unemployment fluctuations in a simple SAM model with productivity shocks more than sevenfold compared to the case of Nash bargaining, and that such a model can easily account for around half of the empirical volatility of unemployment over the business cycle.

Lastly, we investigate how far our results are consistent with various other models of wage setting in the literature. We first examine models in which the wage is consistent with constrained efficiency in the labor market, such as many models of directed search. We find that these models would imply values of the NWE equal to or greater than 1, which is inconsistent with our empirical findings. We then examine a model of staggered wage bargaining similar to Gertler and Trigari (2009) and a model of alternating offer bargaining similar to Christiano et al. (2016). We find that wages in the data are perhaps less rigid than implied by the calibrated staggered wage bargaining model but are more rigid than implied by the alternating offer bargaining model.

Overall, this paper makes five main contributions relative to the literature. These contributions are, first, to show that a large class of search and matching models imply a common equation for the Nash bargained wage. Second, to develop the concept of the Nash wage elasticity, which can be estimated for this large class of models without having to specify which model in this class corresponds to the true data generating process. Third, to provide a range of empirical estimates of the Nash wage elasticity, which overwhelmingly imply extremely rigid wages. Fourth, to show that, across a class of models with or without wage rigidity (or Nash bargaining), the Nash wage elasticity is a strong predictor of the contribution of wage rigidity to the volatility of unemployment over the business cycle. Fifth, to use our Nash wage elasticity estimates to make inferences about how far different models of non-Nash wage setting are consistent with wage data, such as models with constrained efficient wages or rigid wages.

Finally, while our approach is motivated by a desire to estimate wage rigidity and its business cycle implications, we anticipate that the general methodology that we develop could be useful in other contexts. For instance, it may be possible to use similar approaches to estimate price rigidity in goods markets; to estimate the elasticity of asset prices to fundamentals, or to estimate the elasticity of nominal exchange rates to differences in relative goods prices across countries.

The remainder of the paper is structured as follows. Section 3.1.1 discusses the related literature and compares our method and findings to this literature. Section 3.2 discusses the intuition for our approach, develops the modeling framework, derives equations to calculate the Nash wage and formally defines the Nash wage elasticity. Section 3.3 outlines the data sources and calibration used to calculate the Nash wage. Section 3.4 presents our empirical results and discusses the intuition behind our findings. Section 3.5 discusses the implications of our NWE estimates for business cycles and for models with non-Nash wage setting. Section 3.6 concludes.

3.1.1 Related Literature

In this section, we compare our approach to the large existing literature that seeks to estimate the level of wage rigidity and to infer its importance for business cycles. We also outline why our findings differ substantially from some of the work in this literature. While the literature has contributed greatly to our understanding of wage dynamics and business cycle propagation mechanisms, it is, as of yet, still far from consensus on the key questions of how far wages are rigid, and how far wage rigidity matters for business cycles.[2] The literature has been dominated by two broad approaches which differ quite significantly from ours: fully structural models and reduced form estimation.

The first of these two approaches taken by the literature has been to build structural SAM or other DSGE models and either calibrate them and compare to data or structurally estimate them against the data.[3] Since the wage setting process

[2]For instance, see Christiano et al. (2021), Dupraz et al. (2019), Gertler et al. (2020), Pissarides (2009), Basu and House (2016), and Bellou and Kaymak (2021) for recent contrasting views.

[3]Examples from this literature using SAM models include Christiano et al. (2016), Gertler and Trigari (2009), Hagedorn and Manovskii (2008), Hagedorn and Manovskii (2013), Hall and Milgrom (2008) and Pissarides (2009). A number of studies in this literature, such as Merkl and Stüber (2017), Hagedorn and Manovskii (2013), and Gertler et al. (2020) provide new reduced-form estimates of wage cyclicality using an approach along the lines of the reduced form literature discussed below, and use the results of these regressions to calibrate or empirically evaluate a structural model. Studies that assess the implications of wage rigidity in New Keynesian DSGE models with search and matching frictions include Krause and Lubik (2007), Blanchard and Galí (2007) and Christoffel and Linzert (2010). Prominent recent examples from the large literature on the importance of wage rigidity in New Keynesian models include Auclert et al. (2021) and Broer et al. (2020).

is modelled explicitly, the resulting parameters that are found to best fit the data are directly informative about the presence and/or nature of wage rigidity. Nevertheless, the conclusions drawn about the degree and importance of wage rigidity differ substantially across the different studies in this literature. A major reason is that the level of wage rigidity implied by any such calibrated or structurally estimated model may be very model specific, in that a different model may fit the same data equally well with a very different level of wage rigidity. As such, it is unsurprising that the findings of this literature regarding wage rigidity differ across models, and it is not usually clear how far the findings from any particular model are robust to model misspecification.

The principal difference between our approach and this literature is that we rely on a wage equation for the Nash bargained wage that we show holds commonly across a very large class of SAM models. Therefore, our conclusions about wage rigidity do not depend on which model in this class is the correct one, and so are arguably less sensitive to model misspecification.

Our findings also differ substantially from many (although certainly not all) of the SAM models in this literature in that we find a high level of wage rigidity. A key reason for this difference is that our approach is based on a Nash wage equation that applies under models with many different shocks. Therefore, our NWE estimates may be valid even if multiple shocks are important influences on labor demand. In contrast, many models in this literature assume that the only shocks driving fluctuations in labor demand are productivity shocks.[4] Under this assumption, Nash bargaining implies a tight link between wages and productivity, and so much of this literature considers the elasticity of wages with respect to productivity in the data to be very informative about wage rigidity, and empirically evaluates models accordingly.[5] However, given the lack of a strong correlation between unemployment and productivity in the data, it seems implausible that productivity shocks are the only driver of unemployment fluctuations. It is not clear then whether the elasticity of wages with respect to productivity is very informative about rigidity once we allow for other shocks. On the contrary, in our framework, which is consistent with multiple shocks, we find that the Nash wage is practically

[4] Examples of this include Hagedorn and Manovskii (2008), Hall and Milgrom (2008), Pissarides (2009) and Malcomson and Mavroeidis (2017).

[5] Thus, Hagedorn and Manovskii (2008), for instance, calibrate their model to match this elasticity, while Pissarides (2009) suggests that Nash bargaining is supported in the data if the elasticity of new hire wages with respect to productivity is close to 1.

uncorrelated with productivity in the data, and so the elasticity of wages with respect to productivity is not informative about the NWE.

The second approach taken by the prior literature to estimate wage rigidity has been to estimate the cyclicality of real wages via reduced-form regressions.[6] Following Bils (1985), this literature has typically regressed a measure of wages (at the individual or aggregate level) on a cyclical indicator, such as the unemployment rate or productivity.[7] An advantage of this approach, relative to the structural approach above, is that, to estimate the cyclicality of a particular wage measure, it is not necessary to write down a structural model of the labor market. This might seem to avoid the concerns of model misspecification inherent in the structural approach.

However, in practice, the reduced form literature has faced the same barrier to inference as the structural approach – the interpretability of its conclusions often rely on strong assumptions about the underlying theoretical model of the labor market. This is because estimates of wage cyclicality found by the reduced form literature are often highly sensitive to the choices of wage measure (e.g. the wage of all workers, new hires, or new hires out of unemployment) and of cyclical indicator (e.g. unemployment or productivity). Which of these choices seems most justified depends on the theoretical model that the researcher has in mind. For instance, as discussed above, models in which productivity is the only shock suggest that productivity is the natural cyclical indicator on which to regress wages, but this conclusion does not necessarily follow if there are other shocks. Equally, whether the wage measure that best captures the marginal cost of labor is the average wage of all workers, the wage of newly hired workers, the 'user cost of labor' developed by Kudlyak (2014), or none of these varies across theoretical models depending on whether a worker's current wages in the model are influenced by conditions when they were hired, and depending on whether average match quality may vary over time (Kudlyak, 2014; Gertler et al., 2020). Moreover, for a given wage series and cyclical indicator it is regressed on, it is impossible to infer whether an estimated wage cyclicality of e.g. 2% signifies a sticky or flexible wage without knowing how a flexible wage should behave. This is hard to ascertain without a theoretical model.

[6] This approach originates with Dunlop (1938) and Tarshis (1939).

[7] Recent examples of work in this vein includes Haefke et al. (2013), Martins et al. (2012), Carneiro et al. (2012), Kudlyak (2014), Basu and House (2016), Gertler et al. (2020) Grigsby et al. (2021), Hazell and Taska (2020) and Schaefer and Singleton (2021). Much of the earlier literature that studies the wage cyclicality of job stayers and new hires in this way is surveyed by Pissarides (2009)

Consequently, the literature has overwhelmingly interpreted the results of reduced form regressions using specific calibrated models; conclusions from this literature regarding whether the data supports a flexible or rigid wage then depend on the particular theoretical model and calibration strategy chosen.

Our approach differs from this literature in that it delivers an estimate of the Nash wage elasticity that can immediately be interpreted theoretically, as evidence in favor or against Nash bargaining for example, without the need to commit to a particular theoretical model. Furthermore, our theoretical derivation of the Nash wage equation makes clear that the ideal measure of actual wages to estimate the NWE is the user cost of labor based on workers newly hired out of unemployment (without any additional adjustments for match quality), and the cyclical indicator it should be regressed on is the Nash wage. This is true across a large class of models, since the same Nash wage equation holds across a large class of models. Therefore, we are able to answer questions of which wage measure and which cyclical regressor should be used, without needing to commit to a particular theoretical model. Finally, we find that estimates of the Nash wage elasticity are ultimately far below one in many specifications for *all* the measures of actual wages we consider, because all these wage measures are much less cyclical than the Nash wage. Therefore it turns out that the question of how to measure actual wages is relatively less important for estimating the NWE.

In addition to the literature reviewed above, our approach relates closely to recent work by Malcomson and Mavroeidis (2017), Bils et al. (2018), Koenig et al. (2021) and Ljungqvist and Sargent (2017). Malcomson and Mavroeidis (2017), like us, seek to estimate a wage-setting equation while imposing weak assumptions on the data generating process. Unlike us, they find that the data is consistent with new hire wages being set by Nash bargaining. We conjecture that the difference in results arises because they do not allow for markups and implicitly assume that fluctuations in labor demand are driven entirely by productivity shocks. This could lead to a bias against finding wage rigidity for the reasons discussed above on page 136. Koenig et al. (2021) show that a canonical Diamond-Mortensen-Pissarides model implies a wage elasticity to unemployment far higher than the data, consistent with our finding of a low NWE. They also suggest that the elasticity of wages with respect to unemployment is informative across a class of models, and develop a model of reference-dependent wages to account for rigidity. Bils et al. (2018) study the cyclicality of the labor market wedge under search models, finding that much of this cyclicality can be accounted for by the product market wedge. Their findings provide additional evidence that labor market dynamics are affected by time-varying markups, as allowed for in our approach. Ljungqvist and Sargent (2017) show that in many different search and matching models, the determinants of unemployment fluctuations is driven by a common factor they call the 'Fundamental Surplus'.

More broadly, our work relates to the literature that studies the implications of search and matching models with wage rigidity for business cycle fluctuations. Hall (2005), Hall and Milgrom (2008), Christiano et al. (2016) and Gertler and Trigari (2009), among others, develop models of rigid wages and show that these can help explain the volatility of unemployment over the business cycle. Dupraz et al. (2019) find that downward wage rigidity can help account for business cycle asymmetries.

Finally, our approach of developing a measure of wage rigidity, the NWE, that is useful across different models has significant similarities to the literature on estimable sufficient statistics originating with Chetty (2009). Analogous to this literature, the Nash wage elasticity is a rough sufficient statistic that is highly informative about, for instance, the contribution of wage rigidity to business cycle fluctuations, across many different models.

3.2 Modeling Framework

In this section, we formally derive an equation for the Nash bargained wage that holds across a large class of search and matching models, incorporating rich firm and match heterogeneity, a wide variety of different shocks, possible frictions in goods and financial markets, job-to-job transitions and varying labor force participation.

We proceed in stages. First, to provide intuition for how it can be possible to derive a wage equation that holds under such broad conditions, we briefly discuss the case of perfectly competitive labor markets in Section 3.2.1. In Section 3.2.2 we derive the equation for the Nash wage in a framework featuring no firm or match heterogeneity. In Section 3.2.3 we expand our approach to show that virtually the same equation for the average Nash wage arises in a model which is much more general on a number of dimensions, including (but not limited to) firm and match heterogeneity, job-to-job transitions and time-varying labor force participation. Our aim is to derive an equation for the Nash wage which holds in as broad a class of

SAM models as possible. We discuss the cases nested by our modeling framework in Section 3.2.4. Finally, in Section 3.2.5, we discuss how our Nash wage equation can be used to estimate the Nash Wage Elasticity using data on wages and labor market flows. Since our Nash wage equation holds across a very large class of models, it is possible to estimate the NWE without needing to make assumptions about which model in this class accurately describes the data generating process.

3.2.1 Intuition From Perfectly Competitive Labor Markets

Assume that identical households supply labor in a single perfectly competitive spot labor market. As is well known, the resultant equilibrium wage rate must be on the household's labor supply curve, which means that it must equal the household's marginal rate of substitution (MRS) between consumption and leisure.

The essence of our approach is to note that the wage rate will equal the household's MRS under a competitive spot labor market regardless of the determinants of labor demand. For instance, if firms have sticky prices in goods markets, or their ability to hire is affected by working capital constraints, or their capital investment is constrained by financial frictions, all of these things will affect their labor demand and affect equilibrium labor hours and wages but the wage will continue to equal the MRS in all these cases. Likewise, if firms have heterogeneous productivity levels or markups, this will affect the aggregate demand for labor but the wage will continue to equal the MRS.

Then, in a spot labor market, a natural metric of wage rigidity is the elasticity of the observed wage rate with respect to the MRS. Since the perfectly competitive wage will equal the MRS under a wide variety of different assumptions about labor demand, the elasticity of observed wages with respect to the MRS provides a measure of how far observed wages are rigid, compared to competitive wages, and this measure remains equally valid and useful under a wide variety of different assumptions about labor demand. Of course, to measure wage rigidity in this way, it is necessary to have a time series for the MRS. The literature on the cyclicality of the labor wedge Chari et al. (2007) has shown that it is straightforward to calculate a value of the MRS from aggregate data under standard assumptions about preferences.

Our approach differs from simply measuring wage rigidity in terms of the elasticity of wages with respect to the MRS because we allow for search frictions in labor markets. With search frictions, there is no longer any reason to expect that a flexible wage would equal the MRS. Instead, we derive a similarly general expression for the Nash wage which holds under a wide variety of different assumptions about firm and match heterogeneity and about the determinants of labor demand. Just as different assumptions about labor demand affect equilibrium labor hours and wages but do not affect the basic equality between wages and the MRS in the competitive case, so different assumptions about labor demand also affect employment and wages but do not affect the Nash wage equation in the search theoretic case. In effect, the Nash wage equation we will derive is a search theoretic analogue to the labor supply curve in a competitive market. That is, the Nash wage equation defines a locus of points that the wage rate should satisfy, conditional on labor market stocks and flows, and this locus is unaffected by the determinants of labor demand, just as the labor supply curve is unaffected by the determinants of labor demand in the competitive case.

Therefore, we define the Nash wage elasticity as the elasticity of observed wages with respect to the Nash wage derived from our Nash wage equation. We now derive the Nash wage equation formally.

3.2.2 The Nash wage without heterogeneity

In this section, we derive an equation for the Nash wage in a broad framework which nests a substantial number of different SAM models but does not allow for firm or match heterogeneity. We extend the results to a substantially more general setting in the next section.

We first outline the assumptions of our framework with no firm or match heterogeneity. Time is discrete. The economy consists of measure 1 of households and some measure of firms. Households live in large families, made up of employed and unemployed agents. Each large family shares consumption among its members. Unemployed agents match at the start of each period with vacancies v_t posted by firms in period t, according to the constant returns to scale continuously differentiable matching function $M_t = \overline{m}_t \cdot M(u_{t-1}, v_t)$, where $\overline{m}_t$ is a possible shock to the efficiency of the matching function, u_{t-1} is the number of unemployed at the end of the period $t-1$ and start of period t, and v_t is the number of vacancies posted in period t. The unemployed therefore find jobs at the job finding rate $f_t = \frac{M_t}{u_{t-1}}$. There is no on-the-job search.

At the start of period t, fraction s_t of employed agents separate from jobs. We allow that s_t may evolve stochastically over time in response to shocks. The measure of households who are unemployed, u_t, evolves over time according to the following law of motion:

$$u_t = (1 - f_t)u_{t-1} + s_t(1 - u_{t-1}) \tag{3.1}$$

3.2.2.1 Preferences

The members of each large family act to maximize the expected value of :

$$U = \sum_{t=0}^{\infty} (1-\rho)^t u(c_t),$$

where c_t is the consumption of the family and $u(\cdot)$ is strictly increasing and concave. Employed agents earn wage rate w_t in period t. Unemployed agents engage in home production. Each unemployed agent produces z_t units of consumption each period, where z_t changes over time according to some stochastic process.

It is not realistic to interpret z_t as literally representing home production alone. We instead view z_t as a black box for the opportunity cost of employment, which, in a more general model, would include the utility value of the time that an unemployed person does not need to spend working, adjustments to reflect that the unemployed face different tax rates to wage earners, the various cash and in-kind benefits an unemployed person is entitled to, the possibility that these benefits may expire after a certain period of unemployment, and the utility cost of applying for these benefits. In a richer model that incorporates all these features, Chodorow-Reich and Karabarbounis (2016) show that it is possible to derive time series for z_t from aggregate and survey data under various assumptions about preferences. They point out that, as far as SAM models are concerned, what matters for aggregate wage and employment dynamics is the behavior of z_t, rather than the various components of z_t. As such, for simplicity, we do not model the components of z_t, and instead simply treat z_t as home production. In our empirical analysis we calculate the Nash wage using estimated series of z_t from Chodorow-Reich and Karabarbounis (2016), so that our conclusions depend on behavior of z_t that they argue fits the data.

Let $\mathscr{W}_t$ and $\mathscr{U}_t$ denote, respectively, the marginal present value to the household of having an extra employed and unemployed agent. These evolve according to the following Bellman equations:

$$\mathscr{W}_t = \qquad w_t + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-s_{t+1})\mathscr{W}_{t+1} + s_{t+1}\mathscr{U}_{t+1}]\right],$$

$$\mathscr{U}_t = \qquad z_t + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-f_{t+1})\mathscr{U}_{t+1} + f_{t+1}\mathscr{W}_{t+1}]\right],$$

where $\frac{u'_{c_{t+1}}}{u'_{c_t}}$ is the household's stochastic discount factor.

3.2.2.2 Firms

Firms post vacancies, which each cost κ_1 per period. Fraction q_t of vacancies are assumed to match with workers each period. If a vacancy matches with a worker, the firm hires the worker at additional hiring cost κ_0.

The total number of new matches M_t each period must satisfy:

$$M_t = q_t v_t = f_t u_{t-1}.$$

It follows that $q_t = \frac{u_{t-1} f_t}{v_t}$.

We assume, for notational convenience, that vacancies match with unemployed agents before production takes place, and so newly hired workers are productive in the period that they are hired.

Employed agents provide a gross flow value to the firm of r_t in period t, which we allow to evolve over time according to some stochastic process. We are completely agnostic about the determinants of r_t. The term r_t can be interpreted as the marginal revenue product of a worker – so it might depend on the markup as well as on the labor productivity and on the number of hours that a worker works.[8] More generally, if workers provide other useful services to a firm apart from producing output, such as research and development or training of other workers, then these may also enter r_t. Since we are agnostic about the determinants of r_t or about how it varies over time, our framework can nest any friction in goods or financial markets that maps into aggregate quantities via its effect on productivity or markups (and,

[8]It might seem that allowing for time-varying hours should require the disutility of working a particular number of hours to enter the value function of an employed worker. However, we may instead follow Chodorow-Reich and Karabarbounis (2016) and normalize so that the disutility of working the current mandated number of hours features as part of the value of being unemployed, and so this is incorporated into z_t.

therefore, on r_t). This might include, for instance, the effect of sticky prices in goods markets as in Ravenna and Walsh (2008), or working capital constraints, as in Christiano et al. (2015).

Let $\mathscr{J}_t$ and $\mathscr{V}_t$ denote, respectively, the marginal present value of an extra worker and an extra vacancy to a firm. These satisfy the following Bellman equations:

$$\mathscr{J}_t = r_t - w_t + E_t\left[\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-s_{t+1})\mathscr{J}_{t+1} + s_{t+1}\mathscr{V}_{t+1}]\right], \tag{3.2}$$

$$\mathscr{V}_t = -\kappa_1 + (1-q_t)E_t\left[\frac{u'_{c_{t+1}}}{u'_{c_t}}\mathscr{V}_{t+1}\right] + q_t(\mathscr{J}_t - \mathscr{V}_t - \kappa_0). \tag{3.3}$$

Here, $\mathscr{J}_t$ appears on the right-hand side of the Bellman equation for the vacancy in period t, because vacancies that are filled in period t already become productive that period.

Firms are able to create vacancies for free, so, in equilibrium, vacancy posting satisfies the free entry condition $\mathscr{V}_t = 0$. Substituting this into (3.3) and rearranging, we obtain:

$$\tilde{\mathscr{J}}_t = \kappa_0 + \frac{\kappa_1}{q_t} = \kappa_0 + \frac{\kappa_1 v_t}{u_{t-1}f_t} = h_t, \tag{3.4}$$

where h_t denotes the expected hiring cost.

3.2.2.3 Worker Share of Match Surplus

We now derive an equation that defines the share of match surplus that accrues to workers in our framework. Below, we use this to derive a formula for the Nash wage. We define the worker share of match surplus, β_t, as the worker match surplus, divided by the total surplus. That is:

$$\beta_t = \frac{\mathscr{W}_t - \mathscr{U}_t}{\mathscr{J}_t - \mathscr{V}_t + \mathscr{W}_t - \mathscr{U}_t}.$$

Using that $\mathscr{V}_t = 0$ and $\mathscr{J}_t = h_t = \kappa_0 + \frac{\kappa_1 v_t}{u_{t-1}f_t}$, this can be written as:

$$\mathscr{W}_t - \mathscr{U}_t = \beta_t\left[\mathscr{W}_t - \mathscr{U}_t + \kappa_0 + \frac{\kappa_1 v_t}{u_{t-1}f_t}\right],$$

which rearranges to,

$$\mathscr{W}_t - \mathscr{U}_t = \frac{\beta_t}{1-\beta_t}\left[\kappa_0 + \frac{\kappa_1 v_t}{u_{t-1}f_t}\right]. \tag{3.5}$$

Subtracting the Bellman equation for an unemployed agent from the Bellman equation for an employed agent, and substituting in (3.5) to eliminate $\mathscr{W}_t$ and $\mathscr{U}_t$ terms, we obtain the following dynamic equation for β_t:

$$\begin{aligned}\frac{\beta_t}{1-\beta_t}\left(\kappa_0+\frac{\kappa_1 v_t}{u_{t-1}f_t}\right) &= w_t - z_t \\ &+ E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}(1-s_{t+1}-f_{t+1})\frac{\beta_{t+1}}{1-\beta_{t+1}}\left(\kappa_0+\frac{\kappa_1 v_{t+1}}{u_t f_{t+1}}\right)\right].\end{aligned} \tag{3.6}$$

Evidently, the share of match surplus that goes to workers depends on the wage w_t, as is intuitive. Note that we have made no assumptions about how wages are actually set – the dynamic equation (3.6) characterizes the implied share of match surplus that is going to workers, for *any* well-behaved stochastic process governing w_t.

We assume that the economy fluctuates around a steady state. In the steady state, equation (3.6) implies that:

$$\frac{\beta}{1-\beta} = \frac{(w-z)}{[1-(1-f-s)(1-\rho)]h}, \tag{3.7}$$

where, abusing notation, we simply omit the time t subscript to denote the steady state value of a variable. Here we used that $h = \kappa_0 + \frac{\kappa_1 v}{uf}$.

3.2.2.4 The Nash Wage

We define the Nash wage, w_t^N, as the value that the wage w_t would have to take each period in order for the worker surplus share β_t to remain constant over time at its steady state value β, where β_t is calculated according to equation (3.6). Then, it follows that w_t^N satisfies:

$$\begin{aligned}\frac{\beta}{1-\beta}\left(\kappa_0+\frac{\kappa_1 v_t}{u_{t-1}f_t}\right) &= w_t^N - z_t \\ &+ E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}(1-s_{t+1}-f_{t+1})\frac{\beta}{1-\beta}\left(\kappa_0+\frac{\kappa_1 v_{t+1}}{u_t f_{t+1}}\right)\right],\end{aligned} \tag{3.8}$$

where $\frac{\beta}{1-\beta}$ is given by equation (3.7). In our empirical analysis, we assume that the steady state values of variables are equal to their average in the sample period.

Thus, w_t^N is the value that w_t would need to take each period in order for $\frac{\beta_t}{1-\beta_t}$ to remain forever equal to its average value over the sample period.

We refer to w_t^N defined in this way as the Nash wage, since, under the Nash sharing rule, the worker share of match surplus β_t is given by the bargaining strength of workers. The standard assumption in the SAM literature is that this is constant over time so that $\beta_t = \beta$, in which case equations (3.6) and (3.8) imply that $w_t = w_t^N$.

Two further remarks are in order regarding the relationship between w_t^N and the concept of Nash bargaining in SAM models. First, it is more precise to say that w_t^N is defined according to the Nash sharing rule rather than the Nash bargaining solution. In the basic Diamond-Mortensen-Pissarides model, the two coincide: Nash bargaining implies that the worker gets a constant fraction β of the match surplus. However, in some models with frictions, such as Schoefer (2021), the Nash bargaining solution does not imply that the worker's share of match surplus remains constant over time, even though the bargaining weight of workers remains constant. The wage-setting arrangements in such models are therefore not consistent with our notion of the Nash wage, although we conjecture that the Nash bargaining solution will nevertheless deliver something close to a constant share of the match surplus going to workers in many such models in practice, in which case our Nash wage will continue to provide a good approximation to the outcome of Nash bargaining.

Second, since our Nash wage implies a constant share of surplus going to workers, it rules out the possibility of shocks to worker bargaining power, as considered by, for instance, Shimer (2005). If, in reality, there are shocks to worker bargaining power, this could lead to systematic errors in the time series we derive for the Nash wage and bias our estimates of the Nash wage elasticity. We discuss in Section 3.3 how we use other structural shocks as instruments for the Nash wage in our estimation strategy to avoid these biases.

3.2.3 Nash Wage Equation in a More General Environment

The framework used in the previous section to derive the Nash wage equation (3.8) was relatively general on some dimensions. We made almost no assumptions about the stochastic process underlying the opportunity cost of labor z_t, wages w_t, the marginal revenue product of labor r_t, the separation rate s_t or possible shocks to the matching function. Thus, the framework in the previous nests many different models with different assumptions about how these variables are determined.

Nevertheless, the framework in the previous section was special in that, for instance, it did not allow for firm or match heterogeneity. We now generalize our framework to allow for rich heterogeneity in firms and matches, including heterogeneity in firms' discount factors (for instance, due to financial frictions), as well as allowing for endogenous separations, job-to-job transitions and time-varying labor force participation. Our approach is to derive an equation for the Nash wage while making as few assumptions as possible, in order to create a framework which nests as many different SAM models as possible. We show that, even in this much more general case, the equation describing the Nash wage looks very similar to the equation derived in the previous section.

We now outline the assumptions of our more general framework. We maintain the same assumptions as in Section 3.2.2 except where noted.

3.2.3.1 Labor Market Flows

As before, there are measure 1 of households who live in large families. We allow for the possibility that some members of a family are economically inactive, i.e. not in the labor force. Let i_t denote the measure of agents who are economically inactive in period t.

All unemployed agents are identical. We let f_t denote the fraction of unemployed agents who find a job in period t by successfully matching with vacancies. We do not require that all matches are accepted – since we allow for the possibility that poor quality matches are rejected. Thus f_t denotes the probability that an unemployed agent finds a match in period t that she accepts.

It is assumed that agents may shift between being unemployed or economically inactive, but economically inactive agents cannot go straight into employment without first becoming an unemployed agent who looks for a job in some period t, and potentially finding a job in period $t+1$.[9]

Then, the law of motion of the measure of unemployed agents at the end of period t, u_t, is:

$$u_t = (1-f_t)u_{t-1} + s_t(1-i_{t-1}-u_{t-1}) - (i_t - i_{t-1}),$$

where s_t denotes the average separation rate of all employed agents. For simplicity,

[9]It is well known that, in the data, economically inactive individuals do find jobs without previously being registered as unemployed. We interpret such individuals as people who were, in truth, looking for work and therefore unemployed, but were mismeasured as economically inactive.

we do not model agents' choices of whether to be unemployed or economically inactive, we merely assume that i_t evolves over time according to some endogenous stochastic process. Explicitly modeling this motivation would not affect our conclusions provided the flow value of unemployment, z_t, is taken as given.

3.2.3.2 Household Bellman Equations

We allow that workers in different matches and/or at different firms may earn different wages, and that the wages of a worker in a match may evolve idiosyncratically and endogenously over time due to, for instance, match-specific human capital accumulation, or long-term wage contracting as in Rudanko (2009). We also allow the separation rate into unemployment to vary over time and across matches and firms. This could be due to e.g. endogenous separations, where low productivity matches have a higher probability of separating. Let $w_t^{i,k}$ and $s_t^{i,k}$ denote the wage and separation rate in match k at firm i at time t.

We also allow for possible job-to-job transitions, which occur at the start of each period, simultaneously with separations into unemployment. Let $\lambda_t^{i,k}$ denote the probability of a worker in match k at firm i transferring to a new job at the start of period t. This may vary over time and across matches, since workers may be more likely to look for other jobs if their match quality is low.[10]

The Bellman equations for an unemployed agent, and for an employed agent in the match k, are as follows:

$$\mathscr{U}_t = z_t + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-f_{t+1})\mathscr{U}_{t+1} + f_{t+1}\tilde{\mathscr{W}}_{t+1,t+1}]\right], \quad (3.9)$$

$$\mathscr{W}_t^{i,k} = w_t^{i,k} + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-s_{t+1}^{i,k}-\lambda_{t+1}^{i,k})\mathscr{W}_{t+1}^{i,k} + \lambda_{t+1}^{i,k}\tilde{\mathscr{W}}_{t+1}^{i,k,T} + s_{t+1}^{i,k}\mathscr{U}_{t+1}]\right], \quad (3.10)$$

where $\mathscr{W}_t^{i,k}$ denotes the value of a worker in match k at firm i at time t and $\tilde{W}_{t+1}^{i,k,T}$ denotes the expected value that the worker in match k at firm i at the start of $t+1$ expects to have, if she transitions directly to a new job in period $t+1$. $\tilde{\mathscr{W}}_{t,\tau}$ denotes the average value among workers at time t, who were most recently unemployed at the start of period τ, and found a job during period τ. By the usual abuse of the law

[10] Of course, our setting also nests models with no job-to-job transitions, which amounts to fixing $\lambda_t^{i,k} = 0$ for all i,k,t.

of large numbers, we assume therefore that an unemployed agent who finds a job in period $t+1$ has expected value $\tilde{W}_{t+1,t+1}$ in that period.

3.2.3.3 Firms

Firms, i, and matches, k, are heterogeneous in terms of the marginal flow value to the firm of the match, which we denote by $r_t^{i,k}$. This may be due to heterogeneous productivity or markups across firms. If firms have concave production functions or downward-sloping demand curves then $r_t^{i,k}$ will also depend on the number of workers employed by a firm.

Firms may hire out of unemployment, or may hire already employed agents, which precipitates a job-to-job transition. For mathematical tractability, we assume that the unemployed and the already employed match with firms in different submarkets, with potentially more than one submarket for matching with the already employed, as occurs in models of directed search with job-to-job transitions. Then, when a firm posts a vacancy, it decides whether to target the unemployed or the already employed.

As before, we assume a vacancy posting cost of κ_1 and a fixed hiring cost of κ_0 if the vacancy successfully yields a new hire. Let $q_t^{i,u}$ denote the probability that a vacancy posted by firm i that targets the unemployed successfully turns into a match (and an employment relationship) in period t. We allow that $q_t^{i,u}$ varies across firms because, in a setting with firm and match heterogeneity, it is possible that the matches at lower productivity firms are less likely to be accepted, and so less likely to turn into employment relationships.[11]

Then, the Bellman equations for a firm with a match $\mathscr{J}_t^{i,k}$ and a vacancy at firm i that targets the unemployed $\mathscr{V}_t^{i,u}$ are as follows:

$$\mathscr{J}_t^{i,k} = \quad r_t^{i,k} - w_t^{i,k} + E_t\left[m_{t+1}^i(1 - s_{t+1}^{i,k} - \lambda_{t+1}^{i,k})\mathscr{J}_{t+1}^{i,k}\right], \tag{3.11}$$

$$\mathscr{V}_t^{i,u} = \quad -\kappa_1 + q_t^{i,u}(\tilde{\mathscr{J}}_{t,t}^i - \mathscr{V}_t^{i,u} - \kappa_0). \tag{3.12}$$

Here $\tilde{\mathscr{J}}_{t,\tau}^i$ denotes the expected value of a match at firm i at the start of t if the worker was hired out of unemployment in period τ. m_{t+1}^i is the firm i's stochastic

[11] For instance, if a worker and firm observe the idiosyncratic match productivity of a match before deciding whether to go ahead with the match, then it may be that matches at a low productivity firm will only be accepted if the idiosyncratic match-specific component of productivity is particularly good, which may be a low probability event.

discount factor, which we allow to potentially depend on the firm i – this might occur if, for instance, some firms value their cashflow today more relative to the future due to short-term financing constraints.

The Bellman equations above already incorporate that a firm can freely create and dispose of new vacancies. This means that it must be the case that $\mathscr{V}_t^{i,u} \leq 0$, with equality if the firm is maintaining at least one vacancy in period t.

Then, equation (3.12) implies that, if firm i hires in period t, then,[12]

$$\tilde{\mathscr{J}}^i_{t,t} = \kappa_0 + \frac{\kappa_1}{q_t^{i,u}}. \tag{3.13}$$

3.2.3.4 Worker Share of Match Surplus

We now derive an expression for the worker's share of match surplus in this more general framework. We define β_t as the average share of match surplus at time t that is earned by workers who are newly hired out of unemployment in that period (where the average is across new matches of such workers in period t). As will be seen below, it is this measure of worker surplus share for which there exists a mathematical formulation in terms of labor market flows that is almost identical to equation (3.6) above. Thus, we define β_t as:

$$\beta_t = \frac{\tilde{\mathscr{W}}_{t,t} - \mathscr{U}_t}{\tilde{\mathscr{W}}_{t,t} - \mathscr{U}_t + \frac{\int_i q_t^{i,u} v_t^{i,u} \tilde{\mathscr{J}}^i_{t,t} \mathrm{d}i}{\int_i q_t^{i,u} v_t^{i,u} \mathrm{d}i}}.$$

Here, the term $\tilde{\mathscr{W}}_{t,t} - \mathscr{U}_t$ represents the average match surplus of workers newly hired out of unemployment at time t, as defined above. The term $\tilde{\mathscr{J}}_{t,t}$ represents the expected surplus each firm i gets from hiring such workers, and the integral reflects that this should be averaged across firms i in proportion to their share of total hiring out of unemployment, with the hiring of firm i out of unemployment given, in expectation, by $q_t^{i,u} v_t^{i,u}$.[13]

[12]Note that, in many models, it is possible for equation (3.13) to hold in equilibrium for multiple firms with different values of $q_t^{i,u}$ and $\tilde{\mathscr{J}}^i_{t,t}$. In particular, in a model with large firms facing downward-sloping demand curves or concave production functions, firms with a higher $q_t^{i,u}$ will, all else equal, hire more and will see the marginal value of $\tilde{\mathscr{J}}^i_{t,t}$ fall for these firms until (3.13) holds with equality.

[13]Here we are making the usual abuse of the law of large numbers by taking the expectations of the values of firms and hiring costs when integrating across firms.

Rearranging this, and using that (3.13) holds for all firms that set $v_t^{i,u} > 0$, we obtain:

$$\tilde{\mathscr{W}}_{t,t} - \mathscr{U}_t = \left(\frac{\beta_t}{1-\beta_t}\right) \frac{\int_i q_t^{i,u} v_t^{i,u} \tilde{\mathscr{J}}_{t,t}^i \mathrm{d}i}{\int_i q_t^{i,u} v_t^{i,u} \mathrm{d}i} = \left(\frac{\beta_t}{1-\beta_t}\right) \left(\kappa_0 + \frac{\kappa_1 v_t^u}{\int_i q_t^{i,u} v_t^{i,u}}\right),$$

where $v_t^u = \int_i v_t^{i,u} \mathrm{d}i$ denotes the total number of vacancies targeted at the unemployed.

The total number of unemployed agents that find jobs must equal the total number of such vacancies that match, so that $\int_i q_t^{i,u} v_t^{i,u} \mathrm{d}i = u_{t-1} f_t$. Then, we can rewrite the equation above as:

$$\tilde{\mathscr{W}}_{t,t} - \mathscr{U}_t = \left(\frac{\beta_t}{1-\beta_t}\right) \left(\kappa_0 + \frac{\kappa_1 v_t^u}{u_{t-1} f_t}\right). \tag{3.14}$$

In order to use this to derive a dynamic equation for β_t similar to (3.6), it is necessary to characterize $\tilde{\mathscr{W}}_{t,t}$. We do this by gradually integrating the Bellman equation for $\mathscr{W}_t^{i,k}$ across workers and employment histories. This requires some additional notation. Let $P(j,m,t|i,k,\tau)$ denote a probability measure representing the probability that a worker who left unemployment at time τ, obtaining the match k in firm i will subsequently find themselves in period t in match m in firm j, without having spent a spell of unemployment in between (where these probability measures are based on the information available at the start of period τ). Let $f_\tau^{i,k}$ denote a probability measure representing the probability that an unemployed agent at the start of period τ successfully forms the match k with firm i in period τ (where these probability measures are based on the information available at the start of period τ). Let

$$\overline{\mathscr{W}}_{t,\tau}^{i,k} = \frac{\int_{j,m} \mathscr{W}_t^{j,m} \mathrm{d}P(j,m,t|i,k,\tau)}{\int_{j,m} \mathrm{d}P(j,m,t|i,k,\tau)}, \qquad \overline{w}_{t,\tau}^{i,k} = \frac{\int_{j,m} w_t^{j,m} \mathrm{d}P(j,m,t|i,k,\tau)}{\int_{j,m} \mathrm{d}P(j,m,t|i,k,\tau)},$$

$$\overline{s}_{t,\tau}^{i,k} = \frac{\int_{j,m} s_t^{j,m} \mathrm{d}P(j,m,t|i,k,\tau)}{\int_{j,m} \mathrm{d}P(j,m,t|i,k,\tau)}, \qquad \tilde{s}_{t,\tau} = \frac{\int_{i,k} \overline{s}_{t,\tau}^{i,k} (\overline{\mathscr{W}}_{t,\tau}^{i,k} - \mathscr{U}_t) \mathrm{d}f_\tau^{i,k}}{\int_{i,k} (\overline{\mathscr{W}}_{t,\tau}^{i,k} - \mathscr{U}_t) \mathrm{d}f_\tau^{i,k}}.$$

That is, $\overline{\mathscr{W}}_{t,\tau}^{i,k}$, $\overline{w}_{t,\tau}^{i,k}$ and $\overline{s}_{t,\tau}^{i,k}$ denote the average values of $\mathscr{W}$, w and s that a worker expects to obtain at time t, if the worker is hired in match k by firm i at time τ, and remains continuously employed (without a spell of unemployment) between time τ and time t. $\tilde{s}_{t+1,\tau}$ is the average separation rate at time $t+1$ of workers who were hired out of unemployment in period τ (and have not since become unemployed) where the average is weighted across matches in proportion to the surplus of those matches.

Then, it follows from equation (3.10) that $\overline{\mathscr{W}}^{i,k}t,\tau$ evolves according to:[14]

$$\overline{\mathscr{W}}^{i,k}_{t,\tau} = \quad \overline{w}^{i,k}_{t,\tau} + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-\overline{s}^{i,k}_{t+1,\tau})\overline{\mathscr{W}}^{i,k}_{t+1,\tau} + \overline{s}^{i,k}_{t+1,\tau}\mathscr{U}_{t+1}]\right] \tag{3.15}$$

Integrating across all new workers hired out of unemployment at time τ, we obtain, after some rearrangement:

$$\tilde{\mathscr{W}}_{t,\tau} = \quad \tilde{w}_{t,\tau} + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}[(1-\tilde{s}_{t+1,\tau})\tilde{\mathscr{W}}_{t+1,\tau} + \tilde{s}_{t+1,\tau}\mathscr{U}_{t+1}]\right].$$

Repeatedly recursively substituting (3.14) and (3.9) into this equation to eliminate $\mathscr{W}$ and $\mathscr{U}$ terms, we obtain the dynamic equation that describes β_t:

$$\begin{aligned}\frac{\beta_t}{1-\beta_t}\left(\kappa_0 + \frac{\kappa_1 v^u_t}{u_{t-1}f_t}\right) &= w^{UC}_t - \Phi_t - z_t \\ &\quad + E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}(1-\tilde{s}_{t+1,t} - f_{t+1})\frac{\beta_{t+1}}{1-\beta_{t+1}}\left(\kappa_0 + \frac{\kappa_1 v^u_{t+1}}{u_t f_{t+1}}\right)\right],\end{aligned} \tag{3.16}$$

[14]To prove formally that (3.10) implies (3.15), first note that (3.10) implies (3.15) when $t = \tau$. This follows from integrating (3.10) across all matches that a worker could move to as part of an on-the-job transition. Then, by a symmetrical argument, note that if (3.10) implies (3.15) for some $t = n$ and τ, then (3.10) implies (3.15) for $t = n + 1$ and τ. The result then follows by induction.

where,

$$w_t^{UC} = \tilde{w}_{t,t} + E_t \sum_{j=1}^{\infty} (1-\rho)^j \frac{u'_{c_{t+j}}}{u'_{c_t}} \Bigg[\tilde{w}_{t+j,t} \left(\prod_{k=1}^{j} (1-\tilde{s}_{t+k,t}) \right) - \left(\frac{1-\tilde{s}_{t+1,t}}{1-\tilde{s}_{t+1+j,t+1}} \right) \tilde{w}_{t+j,t+1} \left(\prod_{k=1}^{j} (1-\tilde{s}_{t+1+k,t+1}) \right) \Bigg], \tag{3.17}$$

$$\Phi_t = E_t \sum_{j=2}^{\infty} (1-\rho)^j \left(\frac{u'_{c_{t+j}}}{u'_{c_t}} \right) \mathscr{U}_{t+j} \Bigg[- \left(\frac{\tilde{s}_{t+j,t}}{1-\tilde{s}_{t+j,t}} \right) \left(\prod_{k=1}^{j} (1-\tilde{s}_{t+k,t}) \right) + \left(\frac{\tilde{s}_{t+j,t+1}}{1-\tilde{s}_{t+j,t+1}} \right) \left(\frac{1-\tilde{s}_{t+1,t}}{1-\tilde{s}_{t+1+j,t+1}} \right) \left(\prod_{k=1}^{j} (1-\tilde{s}_{t+1+k,t+1}) \right) \Bigg], \tag{3.18}$$

$$\mathscr{U}_t = z_t + E_t \sum_{j=1}^{\infty} (1-\rho)^j \frac{u'_{c_{t+j}}}{u'_{c_t}} \left[z_{t+j} + f_{t+j} \left(\frac{\beta_{t+j}}{1-\beta_{t+j}} \right) h_{t+j} \right]. \tag{3.19}$$

Equation (3.16) is identical to equation (3.6), which described β_t in the model with no firm or match heterogeneity, except for the following differences:

1. The vacancy rate v_t in equation (3.6) is replaced by v_t^U: the number of vacancies targeted at the unemployed.

2. The separation rate s_t in equation (3.6) is replaced by $\tilde{s}_{t+1,t}$: the average next period separation rate of workers newly hired out of unemployment, where the average is weighted by the match surplus.

3. The wage rate w_t is replaced in equation (3.16) with the wage component of the *user cost of labor*, w_t^{UC}, based on workers hired directly out of unemployment. The user cost of labor is a concept first developed by Kudlyak (2014), who observed that in models where wages in a match continue depend on conditions under which the worker was first hired, the macroeconomically relevant measure of the cost of labor is not the current wage rate, but instead depends on the wage of newly hired workers and also the future wage changes that these newly hired workers expect in future.
   Our expression for w_t^{UC} is the same as the expression for the wage component of the user cost of labor in Kudlyak (2014) except for two key differences. First, we allow for a time varying stochastic discount factor and a separation rate that varies across matches and over time, which complicates the user cost equation.

Second, our derivation of w_t^{UC} makes clear that the correct measure of the user cost of labor for our purposes depends on the expected present and future wages of workers *hired directly out of unemployment*. For instance, the key first term in our user cost equation, $\tilde{w}_{t,t}$, is the average wage of workers newly hired out of unemployment. In contrast, Kudlyak measures the user cost using the wages of all newly hired workers – many of whom are workers transitioning from one job to another. As argued by Gertler et al. (2020), the behavior of the wages of workers transitioning from one job to another can give a very misleading impression of the cost of labor for firms, in cases where workers are more likely to move to higher quality matches in booms. At the same time, contrary to what Gertler et al. (2020) and Bellou and Kaymak (2021) have suggested, it is possible to accurately estimate the relevant notion of user cost in the models encompassed by our framework without needing to measure, control or make assumptions about match quality, since there are no measures of match quality in the equation (3.17).

4. There is an additional term Φ_t, which is non-zero if (and only if) the probability of a worker losing their job and entering unemployment depends on the number of periods that the worker has been employed.[15] The term Φ_t enters equation (3.16) because if, for instance, holding a job for longer reduces the likely future separation rate, then a worker who is hired at t and retains their job into $t+1$ is less likely to be unemployed at, e.g. time $t+10$ than a worker who is hired at $t+1$. The consequence is that the value of being unemployed at time $t+10$ therefore affects the value of accepting a job today, and so future unemployment values $\mathscr{U}_{t+j}$ enter into the expression for Φ_t.

3.2.3.5 The Nash Wage

In the more general case, we define the Nash wage w_t^N as the time series that the user cost of labor w_t^{UC} would take each period according to equation (3.16), if it were the case that the share of match surplus going to each worker newly hired from unemployment β_t, remained forever constant at some steady state value β (which we take, in our empirical analysis, to be the average observed value of β_t in

[15] This will be typically true in, for instance, models with endogenous separations – where a longer period in which a worker has been employed may e.g. indicate a likely higher match quality and therefore a lower chance of future separation.

the sample period). Fixing $\beta_t = \beta$ in (3.16), this implies the Nash wage equation:

$$w_t^N = \Phi_t^N + z_t + \frac{\beta}{1-\beta}\left(\kappa_0 + \frac{\kappa_1 v_t^u}{u_{t-1} f_t}\right) - E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}(1-\tilde{s}_{t+1,t} - f_{t+1})\frac{\beta}{1-\beta}\left(\kappa_0 + \frac{\kappa_1 v_{t+1}^u}{u_t f_{t+1}}\right)\right], \tag{3.20}$$

where β is given by the steady state value

$$\frac{\beta}{1-\beta} = \frac{(w - z - \Phi^N)}{[1-(1-f-s)(1-\rho)]h},, \tag{3.21}$$

and where Φ_t^N and Φ^N are calculated by substituting $\beta_t = \beta$ into equation (3.18) and (3.19).

3.2.4 Cases Nested by Our Modeling Framework

Our modeling framework in Section 3.2.3 deliberately imposed minimal structure in order to encapsulate a wide variety of models and mechanisms discussed in the literature. Since we can derive the equation for the Nash wage without imposing more structure, it follows that our equation for the Nash wage holds across a very large range of different models. The consequence of this is that the estimates for the Nash wage elasticity in Section 3.3 below are valid across a large class of models.

Here we briefly outline some of the many cases nested by the modeling framework of Section 3.2.3. Of course, an important special case of our framework is a discrete time version of the canonical Diamond-Mortensen-Pissarides model in Shimer (2005). Our framework captures this special case if all matches are assumed to be homogeneous, r_t^k is the same across matches and equal to aggregate productivity, separation rates are constant over time, and stochastic discount factors are the same across agents and over time.

Equally, our model also covers significant departures from the framework of Shimer (2005). First, the model allows for a wide variety of frictions outside the labor market. Any friction that affects the labor market via the flow value of labor r_t^k is covered. This may occur, for instance, if there are distortions to product markets, such as product market price rigidities. Working capital constraints as in Christiano et al. (2015) are also covered. Likewise, we allow for firms' discount rates to vary over time and differ from household discount rates, which covers cases where financial frictions make firms behave as if they are relatively impatient, as in Schoefer (2021).

Second, the model allows for rich heterogeneity across firms and matches and over time. Since r_t^k can vary across matches over time, our framework allows for the possibilities that workers improve at their jobs over time, that matches persistently vary in quality, or that there are match-specific productivity shocks. Equally, allowing r_t^k and firms' discount rates to vary across firms and over time allows the model to capture cases where the effects of goods and financial market frictions differ across firms and over time. Since we allow that there can be variation across firms and time in the probability that vacancies match with workers and match specific time varying separation rates, our model also covers cases where the probability of matching and separating depend endogenously on e.g. match specific productivity.

Third, our framework allows for time varying labor force participation and job-to-job transitions.

Fourth, our framework allows for essentially any wage-setting protocol. By allowing the wage to depend upon both the current period t and the period at which the match began, we allow for history dependence in wages, as in e.g. Rudanko (2009). While our Nash wage equation by definition captures the wage set according to the Nash sharing rule, we nowhere assume that actual wages are set in this way, since there is no assumption that they equal Nash wages.

Fifth, our framework allows for the possibility of a wide range of macroeconomic shocks. Any shock that operates via r_t^k is covered, such as productivity shocks, markup shocks or shocks that increase the costs of working capital. Since we did not specify a matching function, we also allow for the possibility of variation in matching efficiency over time. We also allow for the possibility of shocks to the separation rate or to firms' discount rates, where the latter could be a consequence of financial shocks.

It is worth emphasizing that we are not making the absurd assumption that none of these many features matter for labor markets. Our framework is fully consistent with these various types of friction and heterogeneity mattering for both equilibrium unemployment and wage determination. However, our common Nash wage equation reveals that, under the Nash sharing rule, the wage has to be set according to the same equation across these many cases, just as workers have to be on their supply curve equation in a competitive labor market, regardless of what is assumed about labor demand.

3.2.5 The Nash Wage Elasticity

In Section 3.3, we use the equations (3.16) and (3.20) for the Nash wage w_t^N and the worker bargaining share β_t from the more general model described in Section 3.2.3 to compute time series for these two variables using data on wages, labor market flows, and series for the opportunity cost of employment z_t from Chodorow-Reich and Karabarbounis (2016). We use the resulting time series for the Nash wage to estimate the Nash wage elasticity, as we now discuss.

3.2.5.1 Computing Series for w_t^N and β_t

To compute series for w_t^N and β_t, we make several assumptions that simplify equations (3.16) and (3.20), due to data limitations.

First, we impose $\Phi_t = \Phi_t^N = 0$ for all t. It is not possible to evaluate Φ_t or Φ_t^N without either data or assumptions about how the separation rate varies across jobs and over time due to job tenure effects and varying match surplus across matches. Neither the sign nor the cyclicality of Φ_t are straightforward to determine, since a separation rate that decreases with job tenure implies both $\left(\frac{\tilde{s}_{t+j,t}}{1-\tilde{s}_{t+j,t}}\right) < \left(\frac{\tilde{s}_{t+j,t+1}}{1-\tilde{s}_{t+j,t+1}}\right)$ and $\left(\frac{1-\tilde{s}_{t+1,t}}{1-\tilde{s}_{t+1+j,t+1}}\right) < 1$. In most models in the literature, the probability of a worker separating from a job and entering unemployment is either exactly or approximately unrelated to the length of time the unemployed individual has been in a job. In such cases, Φ_t is exactly or approximately equal to zero. For this reason, we judge that $\Phi_t = 0$ may be a plausible approximation. Future work could investigate how far allowing for time variation in Φ_t affects estimates of the Nash wage elasticity.

Second, we set $\tilde{s}_{t+1,t} = s_{t+1}$. Recall that the former is a weighted average separation rate, while the latter is simply the economy-wide average separation rate. $\tilde{s}_{t+1,t}$ will tend to be greater than s_{t+1} insofar as workers who have been in a job less time are more likely to separate, but will tend to be less than s_t insofar as matches with a higher surplus are less likely to separate. In practice, the $\tilde{s}_{t+1,t}$ terms are of very small quantitative significance in equations (3.16) and (3.20), and, consequently, we find that our empirical NWE estimates are essentially unaffected by different assumptions about $\tilde{s}_{t+1,t}$, provided $\tilde{s}_{t+1,t}$ is of the same order of magnitude as s_{t+1}.

Third, we assume that v_t^u is proportional to v_t. Specifically, we fix $v_t^u = v_t$, where setting the constant of proportionality to equal 1 is a normalization, since doubling κ_1 and halving v_t^u leaves equations (3.16) and (3.20) unchanged. We make this assumption because we do not know of data available for our long time period regarding what fraction of vacancies is targeted primarily at the unemployed. Setting $v_t^u = v_t$ makes our Nash wage equation consistent with the case of no job-to-job transitions, which is a common, if counterfactual, assumption in models in the literature. Furthermore, making v_t^u proportional to v_t may roughly approximate reality because the job finding rate of the unemployed and frequency of job-to-job transitions appear to have a roughly similar cyclicality in US data (Mukoyama, 2014).

For the empirical analysis, it is convenient to work with log-linearized forms of the equations (3.16) and (3.20). Log-linearizing (3.20) around the steady state using (3.21), setting $\Phi_t = 0$, $\tilde{s}_{t,t+1} = s_{t+1}$ and $v_t^u = v_t$ and using hat variables to denote log deviations from the steady state, we obtain

$$\begin{aligned}\frac{w^{UC}\hat{w}_t^N - z\hat{z}_t}{w^{UC} - z} &= \left(\frac{\kappa_1 v}{huf}\right)(\hat{v}_t - \hat{u}_t - \hat{f}_t) \\ &\quad + \left(f\hat{f}_t + s\hat{s}_t\right)\frac{(1-\rho)}{1-(1-f-s)(1-\rho)} + E_t[\hat{A}_{t+1} - \hat{A}_t], \end{aligned} \tag{3.22}$$

where

$$\hat{A}_t = \frac{(1-\rho)(1-f-s)}{1-(1-f-s)(1-\rho)}\left[\frac{f\hat{f}_t + s\hat{s}_t}{1-f-s} - \left(\frac{\kappa_1 v}{huf}\right)(\hat{v}_t - \hat{u}_t - \hat{f}_t) + \sigma\hat{c}_t\right]. \tag{3.23}$$

Log-linearizing equation (3.16) for β_t, imposing $\Phi_t = 0$ and recursively substituting in (3.22), we obtain:

$$\frac{\hat{\beta}_t}{1-\beta} = \left(\frac{w^{UC}}{w^{UC}-z}\right)E_t\sum_{j=0}^{\infty}\left[1-(1-f-s)(1-\rho)\right](1-\rho)^j(1-f-s)^j(\hat{w}_{t+j}^{UC} - \hat{w}_{t+j}^N). \tag{3.24}$$

This equation confirms that the deviation of the worker's share of the match surplus from its steady state value is proportional to a discounted sum of expected future deviations of the user cost of labor from Nash wages.

We use equations (3.22) and (3.24) to infer the level of the deviation of the worker surplus share, $\hat{\beta}_t$, and the Nash wage, $\hat{w}_t^N$, from empirical data. Given values of these, we propose a measure of aggregate wage rigidity which we term the Nash Wage Elasticity (NWE). The NWE represents the percentage change in the actual wage rate, w_t, when the Nash wage, w_t^N, changes by 1%. An NWE of 1 would imply that Nash bargaining provides an accurate model of wage fluctuations. On the other hand, if the NWE is positive but close to 0, this would imply a wage rate which is relatively insensitive to the macroeconomic factors that influence the Nash wage.

Specifically, we assume a relationship of the form

$$\hat{w}_t^{UC} = \gamma \hat{w}_t^N + \varepsilon_t \tag{3.25}$$

where γ is the NWE and ε_t is a disturbance term. We estimate this equation by OLS and using various instruments for the Nash wage to address concerns of possible measurement error.

The next section discusses the data series and details of the specifications used to estimate the Nash wage. The Nash wage elasticity is closely related to the procyclicality of wages. As we discuss in the next section, our derived value of the Nash wage turns out to be highly procyclical. Consequently, higher values of γ suggest more procyclical wages. Generally, we find γ far below 1, and, accordingly, our measures of the user cost of labor are less procyclical than the Nash wage. As such, we find that the Nash wage tends to be lower than the user cost of labor in recessions. Furthermore, since we showed above that $\hat{\beta}_t$ depends on the difference between the Nash wage and the actual user cost of labor, we find that $\hat{\beta}_t$ is strongly countercyclical.

3.3 Data and Calibration

We compute time series for the model-implied Nash Wage $\hat{w}_t^N$ and worker surplus share ($\hat{\beta}_t$) in US data using the equations (3.22) and (3.24) using various measures of the user cost of labor (w_t^{UC}) various measures of the opportunity cost of employment (z_t), calibrated parameters time series for the job finding rate, separation rate, unemployment and consumption derived from US data. The national statistical sources used for these series are reported in Appendix C.1). We use the headline series of the level of unemployed workers (u^l), number employed workers (e) and number of unemployed workers for less than 5 weeks (u^s) monthly released by the

U.S. Bureau of Labor Statistics (BLS) to derive f and s as in Shimer (2005):

$$f_t = 1 - \frac{u_t^l - u_{t-1}^s}{u_t^l} \tag{3.26}$$

$$s_t = \frac{u_{t+1}^s}{e_t(1 - \frac{1}{2}f_t)} \tag{3.27}$$

The job vacancy rate (v) is taken from Petrosky-Nadeau and Zhang (2020), which they construct by combining the BLS's Job Openings and Labor Turnover Survey (JOLTS) with a range of earlier sources, in order to cover periods where JOLTS data is unavailable.

For the cost of labor, we consider multiple empirical measures. In the previous section, it was shown that the relevant measure of the cost of labor that should ideally be used to calculate the Nash wage elasticity is the user cost of labor based on newly hired workers out of unemployment. The user cost of labor incorporates not only the wage earned by a newly hired worker today, but also the possibility that being hired today rather than tomorrow affects a worker's likely wages in future periods, thereby making it more or less expensive for a firm to hire a worker today versus tomorrow. Unfortunately, no empirical measure of this user cost has been constructed and it is not straightforward to construct a reliable such measure given the available data. In particular, to calculate the user cost accurately, a relatively long individual panel is needed to incorporate the possibility that being employed today affects a worker's wages some distance into the future, as implied by models of implicit contracts such as Rudanko (2009). One such panel is the National Longitudinal Survey of Youth (NLSY), however the NLSY does not ask whether an individual has been hired out of unemployment, and so the wages of new hires out of unemployment cannot be distinguished from those of job switchers. Kudlyak (2014) and Basu and House (2016) have constructed measures of the user cost of labor from the NLSY, based on lumping together all new hires, whether they are job switchers or new hires out of unemployment. This is problematic because Gertler et al. (2020) have found that the wages of job switchers and wages of new hires out of unemployment have quite different cyclical properties, with the latter being substantially more cyclical. If this is the case, the user cost series of Kudlyak (2014) and Basu and House (2016) may have quite a substantial procyclical bias, which would substantially bias upwards estimates of the NWE based on this series (since the Nash wage is very procyclical, as we discuss below).

For this reason, we consider five different empirical proxies for the user cost of labor: the BLS average weekly earnings wage series, two series from Basu and House (2016) and two series from Haefke et al. (2013). We consider all these series because of the substantial debate in the literature regarding the true cyclicality of the cost of labor (Gertler et al., 2020; Bellou and Kaymak, 2021). As we discuss, we believe that some of the series are likely to lead to downward biased estimates of the NWE and other series are likely to lead to upward biased estimates of the NWE. Therefore, by considering the range of measures, we hope to provide a range of estimates of the NWE, with the true value likely to fall somewhere within this estimated range.

We call the first series the CES wage as it is based on the monthly Current Employment Survey (CES) and captures the U.S. employed population average hourly earnings. It is the most widely used measure of an actual wage as it gauges the effective salary across the whole pool of continuously employed workers at a given time. There are two reasons why this headline series may have a countercyclical bias, thus failing to be a representative 'user cost' series: a) since the CES wage is aggregate, this measure suffers from likely countercyclical composition bias if the workers hired in booms have on average lower quality characteristics than the workers hired in recessions and b) an aggregate wage will have a countercyclical bias if the wages of newly hired workers are more influenced by the current state of the labor market than are those of job stayers (Basu and House, 2016; Kudlyak, 2014). Given this likely countercyclical bias, we anticipate that the CES series will tend to understate the Nash wage elasticity.

As a second proxy for the user cost of labor, we use a measure of the new hire wage computed from NLSY data by Basu and House (2016), which specifically corrects for composition bias by controlling for individual fixed effects. As discussed above, this new hire wage includes both job changers and those of newly hired workers out of unemployment, whose wages might have quite different cyclical properties. Another concern with the NLSY series is that the NLSY consists of only a single cohort, which might not be representative of the wider population.

The third series we use is the 'user cost of labor' calculated by Basu and House (2016) from NLSY data (this is similar to the series calculated by Kudlyak (2014)). We refer to this series as the NLSY user cost, to distinguish it from the true user cost of labor in our theoretical framework. This series also adjusts for individual fixed effects to address composition bias. As discussed above, this series gets close conceptually to the relevant concept of the user cost of labor, but may overstate the Nash wage elasticity, possibly by a significant amount, since it lumps together job switchers with newly hired workers out of unemployment.

The remaining two wage series are calculated by Haefke et al. (2013) on the basis of Current Population Survey (CPS) hourly earnings. The first such wage series, which we call Haefke New Hire, is the wage of newly hired workers out of unemployment, since the CPS has information on recent past unemployment status and so can distinguish this group from job switchers. The second series, which we call Haefke All, is a wage of all workers. Both series adjust for composition bias using controls for education, demographic characteristics and experience. The composition bias adjustment means that the 'Haefke All' series is less at risk of countercyclical bias than the CES series discussed above. Nevertheless, since neither of these series is based on panel data, neither of them can fully capture the possible dynamic effects of being hired today on future wages, which should be taken into account as part of the user cost of labor. Thus, neither of these measures fully capture the user cost. Since these two indices have some missing values due to the unavailability of more granular information in the third and forth quarters of 1985 and 1995, we resort to linear interpolation to produce a continuous series.

Next, we multiply our wage indices, which are hourly, by average weekly hours worked (according to the CES) in order to enumerate our series for the cost of labor in terms of the per-worker cost. This transformation is necessary to make the derivative wage indicators consistent with the SAM model stated above, where the wage corresponds to the pay for being employed rather than a hourly rate.

We obtain measures of the opportunity cost of employment, z_t, from Chodorow-Reich and Karabarbounis (2016). As discussed above on page 142, these measures represent the combined advantages of being unemployed relative to the mean marginal product of labor in terms of benefits in cash and in kind, taxes, and more free time. The reader is referred to Chodorow-Reich and Karabarbounis (2016) for details of how these series for z_t are constructed. Using different assumptions on preferences, Chodorow-Reich and Karabarbounis derive four different time series for z_t, computed on the basis of a) separable utility in hours and consumption (SEP); b) Constant Frisch Elasticity (CFE) and two different Cobb-Douglas parametrizations (CD1 and CD2). For completeness, we derive our results using all four series for z_t. As we show in Table 3.1 below, these different series imply very different levels for the average value of z_t over time, but Chodorow-Reich and Karabarbounis robustly find z_t to be highly procyclical, contra the assumption of a constant z commonly considered in the SAM literature. The intuition for this is that the marginal rate of substitution between consumption and leisure is highly cyclical (i.e. workers value consumption relatively more in recessions), which dominates the countercyclicality of unemployment benefits.

Chodorow-Reich and Karabarbounis's different assumptions about preferences imply different assumptions about σ. To maximize consistency with their approach, we use the value of σ associated with each z_t series in Chodorow-Reich and Karabarbounis (2016)

3.3.1 Steady State Variable Values and Calibrated Parameters

Having obtained time series for s_t, f_t, u_t, z_t and so on, our log-linearization approach requires that we calculate the steady state values of these variables as well as their deviations from the steady state. For each variable, we assume that the HP-filtered log value from the data isolates the log deviation of the variable from its steady state value. For almost all variables, we consider the longer term (whole sample) average as the steady state, value. These values are shown in Table 3.1.

Inferring the steady state value of the user cost of labor is less straightforward, since many of our series for this control for individual fixed effects or characteristics, so it is not clear what these series imply for the average user cost of labor over time. Instead, we calibrate the steady state value of the user cost of labor based on the hiring first order condition for firms, in order to maximize consistency with the literature.

This approach requires first that we calibrate the hiring costs κ_0 and κ_1, representing the fixed hiring cost and vacancy posting cost. As a baseline, we set the fixed hiring cost to zero (the traditional assumption in the literature) and set the vacancy posting cost to 0.44. This second parameter deserves more discussion, since there is no consensus in the literature on the total value of hiring costs (relative to the steady state marginal product of labor). Both Hagedorn and Manovskii (2008) and Michaillat and Saez (2021) calibrate the costs based on microdata, both assuming fully variable costs ($\kappa_0 = 0$). We consider different combinations of fixed/variable hiring costs consistent with a steady state hiring cost of roughly 0.58, halfway between the value used by Hagedorn and Manovskii (2008) and Michaillat and Saez (2021). Our steady state costs are stated in the following equation:

$$h = \kappa_0 + \kappa_1 \left(\frac{v}{uf} \right) \simeq 0.58 \tag{3.28}$$

At the same time, since there is also no consensus on the likely split between fixed and variable hiring costs, we explore multiple calibrations of κ_0 and κ_1 below, keeping the steady state hiring cost at 0.58, but varying the fraction of fixed hiring costs in the total from 0 to 90%. This implies values of κ_0 ranging from 0 to 0.52, and values of κ_1 ranging from 0.044 to 0.44.

For the user cost of labor, we assume that steady state w^{UC} is the equilibrium steady state wage that would result if all workers are homogeneous and paid the same wage and there are no goods market frictions in the steady state. Hence, we calculate it according to (3.12):

$$\mathscr{J} = r - w + (1-\rho)(1-s)\mathscr{J} \tag{3.29}$$

The first order condition is $\mathscr{J} = h$. Normalizing $r = 1$, we obtain that in the steady state:

$$w = 1 - (1 - (1-\rho)(1-s))h \simeq 0.98 \tag{3.30}$$

We note that 0.98 is consistent with the SAM literature (e.g. Pissarides (2009) assumes a steady state wage of 0.98).

We summarize the calibration in the Table 3.1 below, which shows the range for κ_0 and κ_1 consistent with the steady state h of 0.58 and the σ parameter values used by Chodorow-Reich and Karabarbounis (2016) for each of the four z series. In the analyses below, we do not show results for other values of σ, ρ or h (i.e. we only vary the ratio $\frac{\kappa_0}{\kappa_1}$). This is because we have found that plausible alternatives to these parameter assumptions make truly negligible difference to our empirical results – affecting NWE estimates by less than 1%.

Table 3.1: Steady States and Calibration

Variables in Steady State		Description
u	0.064	Unemployment Rate
f	0.37	Finding Rate
w	0.98	Wage
z	0.47, 0.76, 0.96	Opportunity cost of employment
s	0.03	Separation Rate
v	0.03	Vacancy rate
β	0.68	Bargaining Share
h	0.58	Hiring Costs
Calibrated Parameters		
k_0	0, 0.29, 0.52	Fixed Hiring Costs
k_1	0.44, 0.22, 0.044	Proportional Hiring Costs
σ	1, 1.52, 1.25, 1.19	Risk Aversion Coefficient
ρ	0.012	Discount Rate

3.4 Results

Figure 3.1: Worker Surplus Share (β_t) based on the NLSY user cost of labor

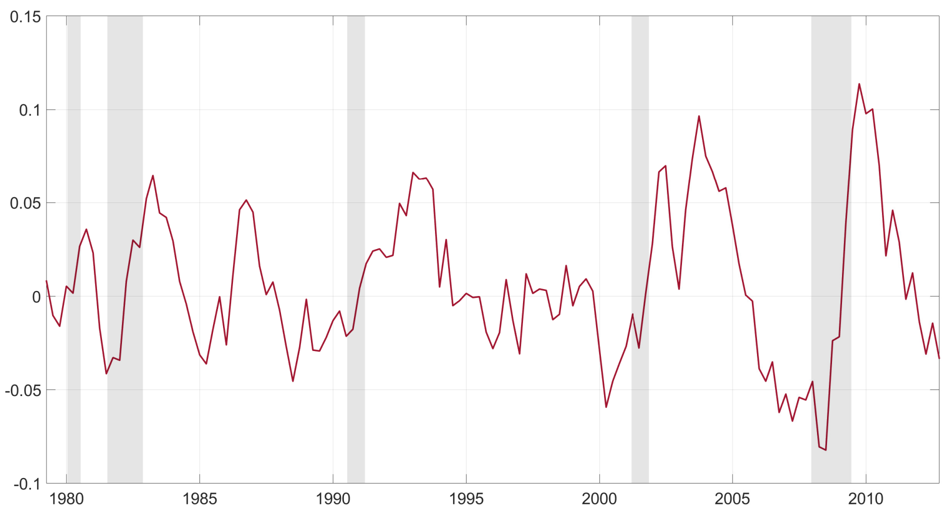

Before estimating the NWE, we graph values of the worker surplus share and Nash wage computed according to the log-linearized expressions for these in Section 3.2.5. In Figure 3.1 we graph the worker surplus share. The baseline series is computed using the input series described in the previous section. The left-hand side only requires arithmetic operations whereas on the right-hand side the terms in the expectations operator have been forecast with a reduced-form VAR with y_t = [1 year interest rate; Real GDP; GDP Deflator; f_t; s_t; w_t; z_t; v_t; u_t; c_t] up to horizon $j = 50$, although the geometrical decay of the summation term makes it effectively nil after the 20^{th} period.

In these equations, the expected forward value $\hat{w}^N_{t+j}$ appears in expectations and since we stated in Eq. (3.22) the Nash wage as a combination of known parameters j periods ahead, given the VAR forecasts, we recursively iterate the VAR to calculate the expectations $E_t\hat{w}^N_{t+j}$ which we use in equation (3.24) to calculate the bargaining share and Nash wage.

In Figure 3.2 we chart the empirical Nash Wage, calculated as in (3.22). Comparing Figures 3.1 and 3.2, it is evident that the measured worker share of match surplus is strongly countercyclical and the Nash wage is strongly procyclical. The countercyclical worker surplus share suggests that, in recessions, workers are earning more than they would if their share of the surplus was constant (as in the Nash sharing rule). This is consistent with the evidence for wage rigidity we document below.

Figure 3.2: Nash Wage $\hat{w}^N_t$

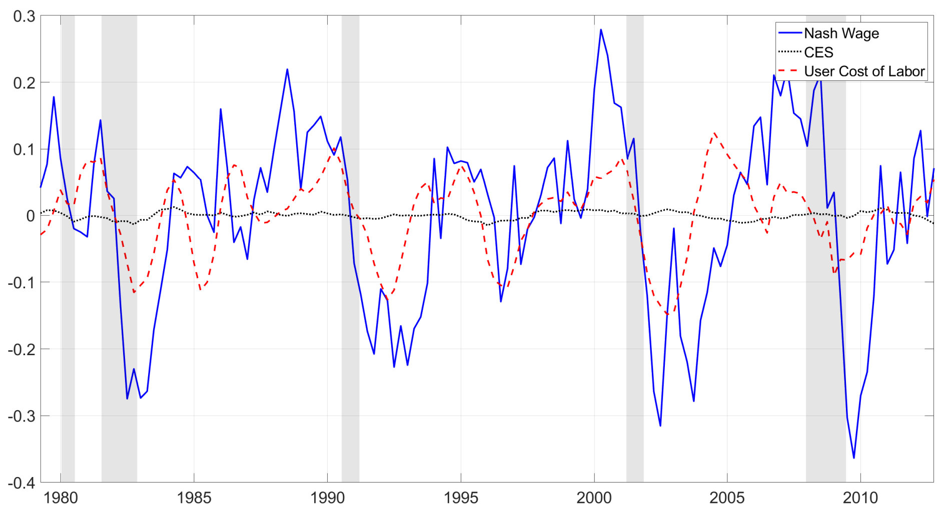

3.4.1 Estimating the Nash Wage Elasticity

In this section, we estimate the NWE. To do so, we estimate 60 linear regressions for which the NWE is the slope coefficient of a regression where the Nash wage is the independent variable and the dependent variable is a measure of the cost of labor. The Nash Wage Elasticity is represented by the coefficient γ in the following ordinary least squares (OLS) regression:

$$\hat{w}_t^{UC} = \gamma \hat{w}_t^{N} + \varepsilon_t \tag{3.31}$$

We estimate 60 different such regressions by varying the cost of labor measure $\hat{w}_t^{UC}$ by alternating between the five proxies for the user cost of labor discussed above, and by using 12 different combinations of the four series for the opportunity cost of labor z with three different combinations of fixed and variable hiring costs. In our equation for the Nash wage, what matters is the size of variable hiring costs relative to fixed hiring costs. To capture the likely range of such variation, we consider specifications with only variable costs ($\kappa_0 = 0, \kappa_1 = 0.43$), as is standard in the literature, specifications where the fixed hiring cost is roughly half of total hiring costs ($\kappa_0 = 0.29$, $\kappa_1 = 0.22$) and specifications where the fixed hiring cost is 90% of total hiring costs ($\kappa_0 = 0.52, \kappa_1 = 0.04$).

In Table 3.2 we find estimates of the NWE ranging from 0.01 to 1.56. However, the vast majority of estimates are below 0.1 (and usually statistically indistinguishable from zero), and only a few estimates are above 0.6. This small number of estimates above 0.6 all use the NLSY user cost of labor (our most procyclical labor cost series) *and* use either almost entirely fixed hiring costs, or use the CD2 series of z. This series of z has an average value of 0.96 – close to the calibration of Hagedorn and Manovskii (2008), which is viewed by most of the subsequent literature as implausibly high (Chodorow-Reich and Karabarbounis, 2016; Christiano et al., 2021). Given the extreme assumptions needed to find an NWE above 0.6, we interpret our estimates as clearly supporting an NWE below 0.6, and favoring an NWE of 0.1 or below.

Table 3.2: Results of Regression (3.31). We calculate the NWE for 5 wage indexes and 12 distinct calibrations of z and hiring costs. Newey–West Standard Errors in brackets.

Z Series	Steady state z	κ_0	κ_1	CES	NLSY New Hire	NLSY User Cost	Haefke All	Haefke New Hire
SEP	0.47	0.00	0.44	0.01	0.02	0.24	0.01	0.02
				(0.01)	(0.04)	(0.04)	(0.01)	(0.02)
SEP	0.47	0.29	0.22	0.01	0.02	0.39	0.03	0.03
				(0.01)	(0.07)	(0.07)	(0.01)	(0.03)
SEP	0.47	0.52	0.04	0.02	0.02	0.65	0.06	0.04
				(0.02)	(0.12)	(0.15)	(0.02)	(0.07)
CFE	0.47	0.00	0.44	0.01	0.01	0.24	0.01	0.02
				(0.01)	(0.04)	(0.04)	(0.01)	(0.02)
CFE	0.47	0.29	0.22	0.01	0.02	0.39	0.03	0.03
				(0.01)	(0.07)	(0.07)	(0.01)	(0.03)
CFE	0.47	0.52	0.04	0.02	0.02	0.67	0.06	0.05
				(0.02)	(0.13)	(0.15)	(0.02)	(0.07)
CD1	0.76	0.00	0.44	0.02	0.04	0.53	0.03	0.04
				(0.01)	(0.09)	(0.10)	(0.01)	(0.05)
CD1	0.76	0.29	0.22	0.03	0.06	0.81	0.06	0.06
				(0.02)	(0.15)	(0.16)	(0.02)	(0.08)
CD1	0.76	0.52	0.04	0.06	0.07	1.16	0.13	0.09
				(0.04)	(0.24)	(0.26)	(0.04)	(0.15)
CD2	0.96	0.00	0.44	0.19	0.38	1.56	0.26	0.13
				(0.06)	(0.40)	(0.54)	(0.06)	(0.30)
CD2	0.96	0.29	0.22	0.21	0.40	1.39	0.28	0.12
				(0.06)	(0.40)	(0.57)	(0.07)	(0.32)
CD2	0.96	0.52	0.04	0.22	0.40	1.18	0.28	0.10
				(0.06)	(0.39)	(0.59)	(0.07)	(0.33)

Next, we estimate the NWE using the (hp-filtered) unemployment rate as an instrument for the Nash wage. To do so, we estimate an IV regression with the same second stage equation as Eq. (3.31), but where the hp-filtered unemployment rate instruments for the Nash wage in the first stage.

$$w_t = \theta w^n + \xi_t \tag{3.32}$$

This approach has two advantages. First, it reduces concerns of measurement error in the Nash Wage. Second, it corresponds to dividing the elasticity of actual wages with respect to the unemployment rate by the elasticity of Nash wages with respect to the unemployment rate.[16] Since the reduced-form literature on wage cyclicality

[16] That is, we choose this specification as analog of the ratio θ_1/θ_2 in the following OLS regressions:

$$w_t = \theta_1 y_t + \xi_{t,1}$$

$$w_t^n = \theta_2 y_t + \xi_{t,2}$$

commonly computes the elasticity of wages with respect to the unemployment rate, this approach has the virtue of easy comparison to that literature.

We report the results in Table 3.3. Throughout, IV estimates are relatively similar to OLS estimates. This is because the Nash Wage is strongly correlated with the unemployment rate (correlation $\simeq 0.8$) in most specifications. Again, in most of cases NWE is close or indistinguishable from zero. Intuitively, as we discuss in Section 3.4.3, this is because the elasticity of the Nash wage with respect to the unemployment rate is much higher than the elasticity of the actual cost of labor with respect to the unemployment rate, for most measures of the cost of labor. The NWE based on NLSY User Cost is markedly positive and in some few instances also greater than 1, but with the caveat of being high under a relative non-standard calibration – using the CD1 and CD2 specification entailing a high value of z, and/or a high value of fixed hiring costs κ_0.

Table 3.3: Results of Regression (3.32). We calculate the NWE for 5 wage indexes and 12 distinct calibrations of z and hiring costs. Newey–West Standard Errors in parentheses.

Z Series	Steady state z	κ_0	κ_1	CES	NLSY New Hire	NLSY User Cost	Haefke All	Haefke New Hire
SEP	0.47	0.00	0.44	0.01	- 0.01	0.27	0.02	0.01
				(0.01)	(0.05)	(0.07)	(0.01)	(0.02)
SEP	0.47	0.29	0.22	0.01	- 0.01	0.41	0.03	0.02
				(0.01)	(0.08)	(0.10)	(0.01)	(0.04)
SEP	0.47	0.52	0.04	0.01	- 0.02	0.68	0.05	0.03
				(0.02)	(0.13)	(0.16)	(0.02)	(0.07)
CFE	0.47	0.00	0.44	0.01	- 0.01	0.28	0.02	0.01
				(0.01)	(0.05)	(0.07)	(0.01)	(0.02)
CFE	0.47	0.29	0.22	0.01	- 0.01	0.42	0.03	0.02
				(0.01)	(0.08)	(0.10)	(0.01)	(0.04)
CFE	0.47	0.52	0.04	0.02	- 0.02	0.70	0.06	0.03
				(0.02)	(0.14)	(0.17)	(0.02)	(0.07)
CD1	0.76	0.00	0.44	0.01	- 0.02	0.62	0.04	0.02
				(0.02)	(0.12)	(0.15)	(0.02)	(0.05)
CD1	0.76	0.29	0.22	0.02	- 0.03	0.91	0.07	0.03
				(0.03)	(0.18)	(0.22)	(0.03)	(0.08)
CD1	0.76	0.52	0.04	0.03	- 0.04	1.46	0.12	0.06
				(0.05)	(0.28)	(0.35)	(0.04)	(0.14)
CD2	0.96	0.00	0.44	0.09	- 0.12	4.11	0.29	0.15
				(0.13)	(0.79)	(0.98)	(0.11)	(0.35)
CD2	0.96	0.29	0.22	0.11	- 0.14	4.85	0.35	0.18
				(0.15)	(0.94)	(1.15)	(0.13)	(0.42)
CD2	0.96	0.52	0.04	0.12	- 0.16	5.66	0.42	0.22
				(0.18)	(1.09)	(1.35)	(0.16)	(0.50)

This definition of an IV estimate is also useful for our approach to estimating the NWE conditional on monetary shocks below.

3.4.2 NWE Conditional on Monetary Shocks

To estimate the NWE conditional on monetary shocks we borrow from the dynamic fiscal multiplier literature. The multiplier is defined as the cumulative change in GDP relative to government spending on an exogenous impulse. It is often approximated as the ratio of the integral of GDP and government spending impulse response functions (IRFs) at an arbitrary horizon h. Ramey (2016a). Nekarda and Ramey (2021) have extended this framework to analyze the conditional response of markup to monetary policy, government spending, productivity and investment specific technology shocks.

A handy way to retrieve IRFs and their ratio is by setting up a vector auto-regression (VAR), a methodology that is simple to implement and provides an intuitive identification when structural restrictions are made explicit. Ramey (2016a, 2018); Barnichon et al. (2021) use a more direct and assumptions-free way to calculate the dynamic fiscal multiplier by comparing the impulse response functions of a h-step ahead local projection (LP). This methodology is appealing because no identifying assumptions are needed and LPs can be computed in a single equation rather than stating a full system. Control variables are usually included to avoid serial correlation and improve the regression fit instead of the cross lags of variables.

In this regard, the literature offers two alternative testing routes. They yield identical results, but it is useful to recall them in order to compare such methodologies with the Regressions stated for the unconditional NWE and provide further intuition. The first method is based on the calculation of LP impulse response functions (IRFs) h periods ahead as:

$$w_{t+j} = a_h + \gamma_j C_t + \psi_h x_t + u_{t+h} \tag{3.33}$$

Where the dependent variable is the variable subject to an identified shock x_{t+j} and C_t is a matrix of control variables. ψ_h is the response of y_t on impulse. In this setting the multiplier can be calculated as the ratio of cumulative IRFs without having to specify a full-blown structural model.

The second method consolidates the IRF ratio in a single equation, providing a direct point-estimate of the multiplier by means of a direct instrumental local projection (IV-LP). As in footnote 16, this IV-LP consolidates the LP in equation (3.33) in a single equation by using the identified shock as an excluded instrument for the Nash wage.

$$\sum_{j=0}^{h} w_{t+j} = m_h \sum_{j=0}^{h} w^n_{t+j} + \gamma_j C_t + u_{t+j} \tag{3.34}$$

Adapting this equation to our case, $\sum_{j=0}^{h} w_{t+j}$ is the cumulative measure of wage, $\sum_{j=0}^{h} w^n_{t+j}$ is the Nash Wage and C_t is a matrix of control variables. The identified shock the is used as an external instrument for w^n. The coefficient m_h is then the estimate NWE multiplier at time horizon h. The IV-LP approach has the advantage to be conducive to the calculation of standard error [17] and weak instrument first stage statistics.

We use the IV-LP to retrieve the NWE measures conditioning on externally identified shocks akin to Nekarda and Ramey (2021). We consider an exogenous monetary shock. Hence we instrument the Nash Wage using the Romer and Romer narrative series for monetary policy Romer and Romer (2004b) as updated in Wieland and Yang (2020).

Our specification is exactly Eq. (3.34) where the matrix of controls contains 1 lag of the shock, log GDP, Nash Wage and wage proxy. The multiplier is calculated at 4 quarters after the monetary policy shock. We summarize the results in Table 3.4.

[17] which have to be heteroscedasticity and autocorrelation robust (HAC), since the IV-LP error term is MA

Table 3.4: Results of Regression (3.34). We calculate the NWE for 5 wage indexes and 12 distinct calibrations of z and hiring costs. Newey–West Standard Errors in parentheses.

Z Series	Steady state z	κ_0	κ_1	CES	NLSY New Hire	NLSY User Cost	Haefke All	Haefke New Hire
SEP	0.47	0.00	0.44	0.01	0.04	0.37	0.02	- 0.01
				(0.01)	(0.06)	(0.18)	(0.01)	(0.04)
SEP	0.47	0.29	0.22	0.02	0.07	0.60	0.04	- 0.02
				(0.02)	(0.09)	(0.30)	(0.02)	(0.07)
SEP	0.47	0.52	0.04	0.03	0.18	1.25	0.09	- 0.05
				(0.04)	(0.21)	(0.63)	(0.06)	(0.16)
CFE	0.47	0.00	0.44	0.01	0.04	0.37	0.02	- 0.01
				(0.13)	(0.14)	(0.13)	(0.01)	(0.04)
CFE	0.47	0.29	0.22	0.02	0.07	0.61	0.04	- 0.02
				(0.13)	(0.15)	(0.13)	(0.02)	(0.07)
CFE	0.47	0.52	0.04	0.03	0.19	1.30	0.09	- 0.05
				(0.13)	(0.17)	(0.15)	(0.06)	(0.16)
CD1	0.76	0.00	0.44	0.02	0.09	0.88	0.02	- 0.01
				(0.13)	(0.15)	(0.14)	(0.01)	(0.04)
CD1	0.76	0.29	0.22	0.03	0.17	1.44	0.04	- 0.02
				(0.13)	(0.17)	(0.15)	(0.02)	(0.07)
CD1	0.76	0.52	0.04	0.01	0.20	2.85	0.09	- 0.05
				(0.11)	(0.19)	(0.17)	(0.06)	(0.16)
CD2	0.96	0.00	0.44	0.02	0.09	8.27	0.02	- 0.01
				(0.11)	(0.19)	(0.23)	(0.01)	(0.04)
CD2	0.96	0.29	0.22	0.04	0.06	11.06	0.04	- 0.02
				(0.11)	(0.20)	(0.24)	(0.02)	(0.07)
CD2	0.96	0.52	0.04	0.08	0.05	15.26	0.09	- 0.05
				(0.11)	(0.20)	(0.25)	(0.06)	(0.16)

The NWE conditioning on the monetary policy shock is similar to the OLS and the IV specifications. The NWE is generally very small or nil one year after a monetary policy shock. This suggests that the actual wage is not sensitive to moves in the NWE also contingent on shocks.

In Figure 3.3 we summarize all the 180 estimates for the NWE derived so far in a box plot. The 12 boxes represent the 12 alternative calibrations, using the four different series for z and three different calibrations of hiring costs. The outliers are all estimates involving the NLSY user cost series. Estimates with a higher steady state z and CD2 are more dispersed and present more outliers. The rest of estimates is very concentrated around 0. Taken together, this evidence is strongly suggestive that the NWE is positive but much closer to 0 than to 1.

Figure 3.3: Box Plot of NWE across all regressions.

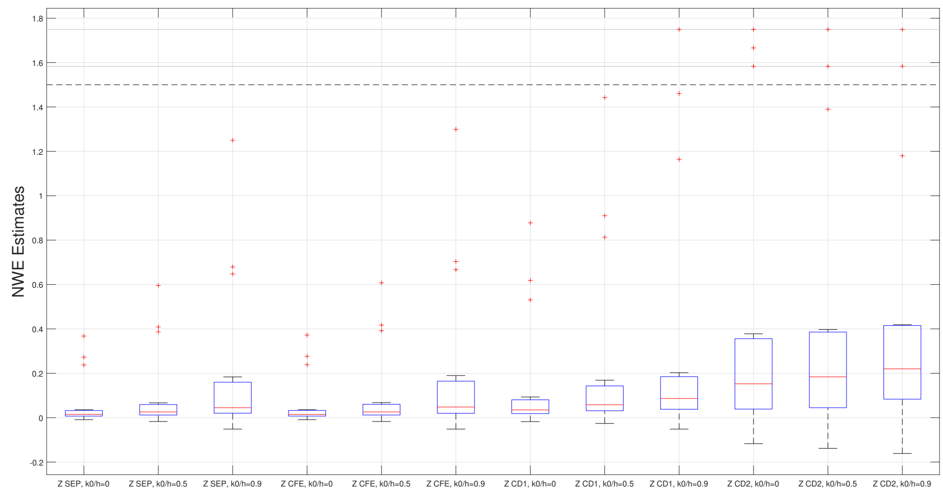

3.4.3 What is driving our results?

Across many different measures of wages and different empirical specifications, we have estimated values of the NWE that are positive, but substantially below 1. Perhaps surprisingly, this is true in many specifications even when we use the highly procyclical NLSY user cost of labor from Basu and House (2016). Key to understanding our results is that our measured series for the Nash wage are highly procyclical, more procyclical even than the NLSY user cost of labor.[18] Since the actual cost of labor is significantly less procyclical than the Nash wage, it follows that the NWE is substantially below 1.

In this section, we discuss why we find that the Nash wage is so procyclical. We first provide some informal intuition, before making the discussion more precise.

Informally, the key to our results is that firm match surpluses appear to be procyclical, whereas worker match surpluses appear highly countercyclical, across different measures of the cost of labor. The reason that the firm match surplus appears to be procyclical is that the firm's hiring decision implies that the firm match surplus must equal the hiring cost, and hiring costs are procyclical since vacancies take longer to be filled when unemployment is low. On the other hand, the worker match surplus is strongly countercyclical, as the value of unemployment is much lower in recessions due to the longer time required to find a job. With a

[18] The procyclicality of the Nash wage and actual wage were shown in Table 3.3 above.

procyclical firm surplus and a strongly countercyclical worker surplus, it follows that the worker's share of the total surplus is strongly countercyclical, and so the Nash wage must be substantially more procyclical than the actual wage. Then, the NWE must be substantially below 1. In principle, if wages were procyclical enough then this would be consistent with a procyclical worker match surplus and an NWE of 1, but it turns out that this would require wages to be a lot more procyclical than any of our empirical measures of labor cost.

To develop this argument formally, assume, for simplicity, that all variables follow approximate random walks, so that for each variable $E_t[x_{t+1}] \simeq x_t$. Assume also that the separation rate is roughly constant, so that $\hat{s}_t = 0$. Then, equation (3.22) gives the following approximate formula for the Nash wage.

$$\frac{w^{UC}\hat{w}_t^N - z\hat{z}_t}{w^{UC} - z} - f\hat{f}_t\left(\frac{(1-\rho)}{1-(1-f-s)(1-\rho)}\right) = \left(\frac{\kappa_1 v}{huf}\right)(\hat{v}_t - \hat{u}_t - \hat{f}_t)$$

Furthermore, log linearizing equation (3.1), setting $u_{t+1} = u_t$ and $\hat{s}_t = 0$ yields:

$$(f+s)\hat{u}_t = -f\hat{f}_t.$$

Substituting this into the equation above, we obtain:

$$\underbrace{\frac{w^{UC}\hat{w}_t^N - z\hat{z}_t}{w^{UC} - z} + (f+s)\hat{u}_t\left(\frac{(1-\rho)}{1-(1-f-s)(1-\rho)}\right)}_{\text{Deviation of Worker Match Surplus}} = \underbrace{\left(\frac{\kappa_1 v}{huf}\right)\left(\hat{v}_t - \frac{s}{f+s}\hat{u}_t\right)}_{\text{Deviation of Firm Surplus}} .$$

The right-hand side of this equation is the deviation from steady state of the hiring cost, which equals the deviation from steady state of the firm's match surplus. Since vacancies are procyclical and unemployment is countercyclical, the firm's surplus is procyclical.

The equation states that this has to equal the deviation from steady state of the worker surplus. The reason that the left-hand side of the equation has to equal the right-hand side is that the Nash wage $\hat{w}_t^N$ is defined as the wage that keeps the worker share of surplus constant, requiring that the log deviation of the firm and worker surplus are equal.

The worker surplus term on the left-hand side has two components. The first term is the deviation of the worker surplus is the change in the wage minus opportunity cost of working. The second term, proportional to $\hat{u}_t$, represents that the worker's match surplus relative to unemployment depends on how long a worker would expect to be unemployed if they were to quit the job. Since the duration of unemployment increases as the unemployment rate increases, this term relates positively to $\hat{u}_t$.

If the wage were too acyclical, so that $\hat{w}_t^N = 0$, then the left-hand side of this equation would be countercyclical, since our data series for $\hat{z}_t$ is procyclical (Chodorow-Reich and Karabarbounis, 2016), and the unemployment rate is countercyclical. Since the right-hand side is procyclical, $\hat{w}_t^N$ has to be procyclical for the left-hand side to equal the right-hand side. That is, the firm's match surplus is procyclical, and everything apart from wages makes the worker surplus countercyclical. As such, the wage rate would have to be quite procyclical in order for the worker's surplus to be as procyclical as the firm's, which is required by Nash bargaining.

How procyclical then does $\hat{w}_t^N$ have to be? To answer this question, rewrite the left-hand side of the equation above as:

$$\left[\frac{w^{UC}\left(\frac{\hat{w}_t^N}{\hat{u}_t}\right) - z\left(\frac{\hat{z}_t}{\hat{u}_t}\right)}{w^{UC} - z} + \left(\frac{(1-\rho)(f+s)}{1-(1-f-s)(1-\rho)}\right)\right]\hat{u}_t.$$

For this to be procyclical, it must be negatively related to $\hat{u}_t$. That is, we need that:

$$\frac{w^{UC}\frac{\hat{w}_t^N}{\hat{u}_t} - z\frac{\hat{z}_t}{\hat{u}_t}}{w^{UC} - z} + \left(\frac{(1-\rho)(f+s)}{1-(1-f-s)(1-\rho)}\right) < 0$$

Using that $\frac{\hat{z}_t}{\hat{u}_t} < 0$, this can be rearranged to:

$$\left|\frac{\hat{w}_t^N}{\hat{u}_t}\right| > \left(1 - \frac{z}{w^{UC}}\right)\left(\frac{(f+s)(1-\rho)}{\rho+(f+s)(1-\rho)}\right) + \frac{z}{w^{UC}}\left|\frac{\hat{z}_t}{\hat{u}_t}\right|$$

The right-hand side term in round brackets will be very close to 1 in an empirically plausible calibration. Since the last term is also positive, this requires in practice that:

$$\left|\frac{\hat{w}_t^N}{\hat{u}_t}\right| > 1 - \frac{z}{w^{UC}}.$$

This condition implies that the Nash wage must be extremely procyclical. In particular, for our baseline series of z, the right-hand side is equal to 0.53. The left-hand side is the percentage change in Nash wages when the number unemployed increases by 1%. If this exceeds 0.53, and the average unemployment rate is 6.3%, then a 1 percentage point increase in the unemployment rate (i.e. from 6.3% to 7.3%) must reduce the Nash wage by more than 0.53/0.063 = 8.4 percent. Note that this is merely the minimum level of cyclicality of the Nash wage required for the worker match surplus to be at all procyclical. If the firm surplus is highly procyclical (which will be the case if κ_0 is substantially below 1) then Nash bargaining requires for the worker match surplus also to be highly procyclical, implying that a 1% decrease in the unemployment rate must increase the Nash wage by substantially more than 8.4 percent. As such, the Nash wage is highly procyclical. The implication, then, is that the measured wage would have to be roughly this procyclical for us to find an NWE close to 1. Since none of our series for the cost of labor are anything like this procyclical (see Table 3.3), we find an NWE considerably below 1.

3.5 Business Cycle Implications of Wage Rigidity

In this section, we assess the business cycle implications of our estimated level of aggregate wage rigidity. For simplicity, we restrict attention to the environment with homogeneous firms and matches and no on-the-job search outline in Section 3.2.2. Furthermore, as in much of the theoretical literature, we restrict attention to an economy which experiences only one shock, to the marginal revenue product of labor r_t, which follows an exogenous stochastic process. Shocks to r_t could be interpreted as, for instance, productivity shocks or markup shocks, or aggregate demand shocks in a model with sticky prices in goods markets. In addition, we assume for simplicity that the matching function and separation rate are time invariant so that $M_t = M(u_{t-1}, v_t)$ and $s_t = s$. We investigate the consequences of different assumptions about wage setting in this environment.

First, we show analytically that, if all variables follow approximate random walks, then the NWE is approximately a sufficient statistic for the contribution of wage rigidity to the cyclical volatility of unemployment in such a model. This is because we show that there is a tight mathematical relationship between the NWE and the Fundamental Surplus, which Ljungqvist and Sargent (2017) have shown is a useful predictor of the cyclical volatility of unemployment in many search models.

Next, we show via simulations that the NWE does indeed closely predict the volatility of unemployment in a simple SAM model with shocks to r_t, just as the link with the Fundamental Surplus would lead to us to expect. To do this, we calibrate a very simple log-linearized business cycle model based on our search and matching framework. When the NWE is as low as most of our empirical estimates, we show that wage rigidity amplifies unemployment fluctuations in the model more than sevenfold compared to the case of Nash bargaining, and that such a model can easily account for around half of the empirical volatility of unemployment over the business cycle if the only shocks to r_t are productivity shocks.

Lastly, we investigate how far our results are consistent with various other models of wage setting in the literature, including models with constrained efficient wages, such as many directed search models, and rigid wage models based on Hall (2005), Gertler and Trigari (2009) and Christiano et al. (2016).

3.5.1 The NWE and the Fundamental Surplus

We now revisit, in our framework, recent results of Ljungqvist and Sargent (2017), who show in the context of a number of SAM models, including some with sticky wages, that the elasticity of the unemployment rate with respect to productivity shocks depends closely on a term that they call the 'Fundamental Surplus'. We extend their results to the class of models studied in this section.[19] Furthermore, we show that the fundamental surplus depends closely on the Nash wage elasticity, and that the Nash wage elasticity is therefore a strong predictor of the effect of wage rigidity on the volatility of unemployment. The relationship between the Fundamental Surplus and the Nash wage elasticity is so close that it is very difficult for a model in the class we study to deliver a high volatility of unemployment unless it either has a low Nash wage elasticity or has a high value of $\frac{z}{w}$.

We assume, as in Section 3.4.3 that all variables follow approximate random walks and set, for each variable x, that $\hat{x}_t \simeq \hat{x}_{t-1} \simeq E_t[\hat{x}_{t+1}]$ – i.e. we consider the long run effects of an almost permanent shock. This is essentially the same as Ljundqvist and Sargent's approach of studying the comparative statics of model steady states with respect to parameters, and is a valid approximation insofar as shocks are highly persistent.

To derive the fundamental surplus formula, log-linearize the matching function, to

[19]This class was specified immediately above.

infer that (under the random walk assumption):

$$\hat{f}_t \simeq (1-\phi)(\hat{v}_t - \hat{u}_t) \tag{3.35}$$

where ϕ is the elasticity of the matching function with respect to unemployment, in the neighborhood of the steady state. Substitute (3.4) into (3.2) to eliminate $\mathscr{J}$ terms and log-linearize. Then, combine with (3.22) and (3.35) and use that $\hat{w}_t = \varepsilon_N \hat{w}_t^N$, and $\hat{z}_t = \varepsilon_z \hat{w}_t$, where ε^N is the Nash wage elasticity and ε_z is the procyclicality of z_t relative to wages. After rearrangement, we obtain:

$$-\frac{\hat{u}_t}{\hat{r}_t} \simeq \phi(1-u)\tilde{\Upsilon}\underbrace{\left(\frac{w+\delta_F h}{\alpha_0\delta_F h + (w-z)\tilde{\varepsilon}\left[1-\frac{\delta_F}{\delta}(1-\alpha_0)\right]}\right)}_{\text{Inverse Fundamental Surplus Ratio}}\hat{r}, \tag{3.36}$$

where

$$\begin{aligned}
\tilde{\varepsilon} &= \frac{\varepsilon_N}{1-\varepsilon_z\varepsilon_N\frac{z}{w}},\\
\delta &= 1-(1-f-s)(1-\rho),\\
\delta_F &= 1-(1-s)(1-\rho),\\
\alpha_0 &= 1-\frac{\kappa_1}{h}.
\end{aligned}$$

and

$$\tilde{\Upsilon}^{-1} = \left(\frac{\delta\tilde{\varepsilon}-\delta_F\tilde{\varepsilon}}{\left(\frac{1-\beta}{\beta}\right)\alpha_0\delta_F + \delta\tilde{\varepsilon} - (1-\alpha_0)\delta_F\tilde{\varepsilon}}\right)\phi + \left(\frac{\left(\frac{1-\beta}{\beta}\right)\alpha_0\delta_F + \alpha_0\delta\tilde{\varepsilon}}{\left(\frac{1-\beta}{\beta}\right)\alpha_0\delta_F + \delta\tilde{\varepsilon} - (1-\alpha_0)\delta_F\tilde{\varepsilon}}\right)(1-\phi).$$

The left-hand side of (3.36) is the size of response of unemployment to a shock to $\hat{r}_t$. Thus, squaring this equation gives the cyclical volatility of unemployment relative to $\hat{r}_t$. Ljungqvist and Sargent (2017) assume $\varepsilon_z = 0$, and consider cases where $\varepsilon_N = 1$ (Nash bargaining) and $\varepsilon_N = 0$ (the completely sticky wage of Hall (2005) discussed below). After some rearrangement, it can be shown that the values of $\tilde{\Upsilon}$ and of the Inverse Fundamental Surplus Ratio are exactly the same in these special cases as found by Ljundqvist and Sargent (using different notation). The 'Fundamental Surplus' refers to the reciprocal of the Inverse Fundamental Surplus Ratio.

It is immediate that wage behavior only enters the right-hand side of equation (3.36) via the Nash wage elasticity. As such, insofar as the random walk approximation is accurate, the Nash wage elasticity is an accurate summary statistic for the effect of wage rigidity on the cyclical volatility of unemployment.

We now discuss how this formula shows that the Fundamental Surplus term, and, in particular, the Nash wage elasticity and the ratio $\frac{z}{w}$ are the key drivers of the cyclical volatility of unemployment. Ljundqvist and Sargent show that $\tilde{\Upsilon}$ is bounded below by 1 and above by $\min[\phi; 1-\phi]^{-1}$, a result that can be straightforwardly seen to also hold in our setting by inspecting the expression for $\tilde{\Upsilon}^{-1}$. The standard view in the literature is that the data supports $\phi \simeq 0.5$ (Petrongolo and Pissarides, 2001) in which case $\tilde{\Upsilon} \in [1, 2]$. Then, the only way to get a high volatility of unemployment relative to $\hat{r}_t$ (which is the easiest way to get the model to produce large fluctuation in unemployment) is to make the Inverse Fundamental Surplus Ratio large and the Fundamental Surplus small. Given the very small size of the terms in δ and δ_F in the equation for the Fundamental Surplus, it is virtually impossible to make the Fundamental Surplus small unless the term $\frac{(w-z)\tilde{\varepsilon}}{w}$ is small – in other words, either the Nash wage elasticity (which is the main term in $\tilde{\varepsilon}$) is small, or z is close to w – workers are roughly indifferent between being unemployed and employed. This echoes the conclusion of Christiano et al. (2021) that wage rigidity is essential to allow SAM models without very high z to deliver large fluctuations in unemployment.[20]

3.5.2 Numerical Simulations

We build a very simple calibrated SAM model based on the model framework laid out in Section 3.5. Our simulations reveal that the NWE closely predicts the volatility of unemployment in the model, relative to the volatility of r_t, which is the driving shock.

We parametrize the model laid out in Section 3.5 in the simplest possible way while allowing the NWE to vary. In the next section we show results of simulations of this model that show that changes in the NWE predict the cyclical volatility of unemployment well, just as implied by the formal analysis of the Fundamental Surplus above.

[20]Ljundqvist and Sargent also argue that, for instance, large fixed hiring costs κ_0 will bring down the fundamental surplus. The formulation here, where κ_0 does not directly appear in the fundamental surplus formula, makes clear that large fixed hiring costs can significantly shrink the fundamental surplus (only) insofar as they increase the equilibrium steady state ratio $\frac{z}{w}$.

We assume that all matches are homogeneous and that the marginal revenue product of labor r_t follows an AR(1) process with quarterly autocorrelation equal to 0.97, roughly the autocorrelation of labor market tightness in our sample period. We are agnostic about whether changes in the marginal revenue product of labor are due to changes in markups (e.g. as a consequence of aggregate demand shocks with nominal rigidities in goods markets) or because of changes in productivity. The separation rate s_t is time invariant and set equal to the steady state separation rate in our empirical analysis.

Workers and firms match according to a Cobb Douglas matching function:

$$M_t = \overline{M} v_t^{1-\phi} u_t^{\phi},$$

where $\overline{M}$ is a constant and $\phi = 0.5$.

To examine the aggregate effects of our estimated level of wage rigidity, we assume that the wage satisfies:

$$\hat{w}_t = \gamma \hat{w}_t^N, \tag{3.37}$$

where γ is the Nash wage elasticity, and the Nash wage is determined by the log-linearized equation (3.22) that was derived in Section 3.2. We consider values of the Nash wage elasticity γ in the range from 0 to 1.

For simplicity, we assume linear utility, so that $\sigma = 0$. This entails that $\hat{c}_t$ drops out of the model equations, which allows us to avoid making assumptions about goods markets and the determination of aggregate consumption.

Finally, it is necessary to specify the cyclicality of the flow value of unemployment. For this we consider two cases, one case where z_t is acyclical, and one where it is proportional to w_t, in which case:

$$\hat{z}_t = \hat{w}_t \tag{3.38}$$

All other variables are calibrated in line with the steady state values we used in Section 3.3 above.

Table 3.5 shows the standard deviation of unemployment relative to the marginal revenue product of labor for different levels of the NWE, and for the case of acyclical and procyclical z. All moments are hp-filtered, in accordance with our empirical analysis. To adjudicate the accuracy of the fundamental surplus formula, (3.36), the volatility implied by that formula is also shown. We see that, except for very low

Table 3.5: NWE and Simulated Relative Unemployment Volatility

	Acyclical z		z Proportional to w	
NWE	Relative Volatility of Unemployment	Volatility Implied by FS formula	Relative Volatility of Unemployment	Volatility Implied by FS formula
0	21.1	46.1	21.1	46.1
0.05	8.6	13.5	8.4	13.2
0.1	5.4	7.9	5.2	7.6
0.2	3.1	4.3	2.8	3.9
0.3	2.2	3.0	1.9	2.6
0.4	1.7	2.3	1.4	1.9
0.5	1.3	1.8	1.0	1.4
0.6	1.1	1.5	0.8	1.1
0.7	1.0	1.3	0.7	0.9
0.8	0.9	1.2	0.5	0.7
0.9	0.8	1.0	0.4	0.6
1	0.7	0.9	0.4	0.5

values of the NWE, the fundamental surplus formula provides a relatively good guideline of the likely effect of wage rigidity on the volatility of unemployment.

The relative volatility of unemployment increases substantially as the Nash wage elasticity falls. With a Nash wage elasticity of 0.1, somewhat higher than most side of our estimates, the relative volatility of unemployment is more than 7 times as high as in the case of a Nash wage elasticity of one. Thus, our empirical findings suggest that wage rigidity may be increasing the cyclical volatility of unemployment more than sevenfold compared to what would be occurring under flexible wages.[21] In our data, the relative cyclical volatility of unemployment is roughly 11 times that of productivity. Thus, our estimates suggest that if r_t represented shocks to productivity alone (i.e. we ignored e.g. aggregate demand shocks) then the empirical level of wage rigidity can account for around half of the cyclical volatility of unemployment. Thus, wage rigidity goes a long way to explaining the 'Shimer puzzle' that unemployment is far more volatile relative to productivity than implied in a simple model with Nash bargaining: it can explain around half of the Shimer puzzle just with productivity shocks alone.

[21] A caveat with this analysis is that, in a richer model, changes in the flexibility of wages could have additional repercussions for labor demand and therefore the volatility of r_t. For instance, if a higher level of wage flexibility led to a less countercyclical capital-labor ratio, this might make r_t less volatile. Alternatively, if changes in wages affect the aggregate demand for goods, this could also affect r_t if goods markets feature nominal rigidities.

3.5.3 Implications of the NWE for Non-Nash Wage Models

We now study how far our estimates of the NWE are informative for various models of non-Nash bargaining. We investigate the implications of our NWE estimates for four non-Nash approaches to modeling wages from the recent literature, first models in which the labor market is constrained efficient, as in many models of directed search (Wright et al., 2021), then three models designed to generate rigid wages: the approaches of Hall (2005), Gertler and Trigari (2009) and Christiano et al. (2016). We show that the wage setting assumptions in these papers can, with small changes, be incorporated into the framework of Section 3.2, at least in the case of homogeneous firms and matches.

We show that the constrained efficient wage setting model delivers a wage that is weakly more procyclical than the Nash wage provided the matching function displays as much complementarity between unemployment and vacancies as the main matching functions considered in the literature. Therefore, if wages were set in a way consistent with constrained efficiency, we would expect to estimate a value of the NWE greater than 1. As such, our low estimates of the NWE indicate, first, that many directed search models are likely to struggle to explain the pattern of wages we see in the data and, second, that wages in the data appear to be more rigid than is consistent with constrained efficiency.

We show that each of the three approaches to rigid wages implies a wage setting equation where the aggregate wage is a function of the Nash wage, and (possibly) hiring costs and the flow value of unemployment.[22] We study the cyclical implications of these three approaches to rigid wages by incorporating the wage setting equation of each into the business cycle model studied in Section 3.5.2. Our simulation allows us to infer the values of the Nash wage elasticity implied by these rigid wage models, as well as the resulting cyclical volatility of unemployment.

The simplest approach to rigid wages of the three is the approach of Hall (2005). In this model, firms are homogeneous and each firm's wage is assumed to be fixed provided that the fixed wage is consistent with positive match surplus for both worker and firm. If the steady state wage rate is consistent with positive match surplus for both worker and firm, then this will continue to be true in the neighborhood of the steady state, and so the wage will remain fixed in the

[22] Consistently with the framework of Section 3.2 the wage setting equation implied by these three approaches to rigid wages does not depend on many other features of the economic environment such as frictions in goods markets.

neighborhood of the steady state. In that case, it follows that $\hat{w}_t = 0$, and so the Nash wage elasticity implied by the Hall (2005) model is exactly zero. This is not far from some of our estimates of the NWE in Section 3.4. Since this model is a special case of the model in Section 3.5.2, the row of Table 3.5 corresponding to an NWE of 0 shows the results implied by the wage setting assumption of Hall (2005).

We now study the cyclical properties of the other three non-Nash models mentioned: constrained efficient wages, staggered wage bargaining and alternating offer bargaining.

3.5.3.1 Constrained Efficient Wages and Directed Search

We suppose that the wage is set in such a way that, in the absence of goods market or financial market frictions, the equilibrium level of unemployment is constrained efficient. This allows us to infer the wage behavior implied by the many directed search models which entail constrained efficiency in the absence of frictions in other markets.[23]

We suppose that the matching function has an elasticity of substitution between unemployment and vacancies that is weakly less than 1, so that the elasticity of matches with respect to unemployment, ϕ_t, is weakly decreasing in the number of unemployed. This assumption nests the cases normally considered in the literature, including the common Cobb-Douglas matching function which has an elasticity of substitution of 1, as well as, for instance, urnball matching functions.

A constrained efficient allocation would set vacancies according to the following first order condition of a benevolent social planner:

$$
\begin{aligned}
p_t - z_t - \frac{\kappa_1}{1-\phi_t}\left(\frac{v_t}{u_{t-1}f_t}\right) + E_t\Bigg[\left(\frac{u'(c_{t+1})}{u'(c_t)}\right)\frac{\kappa_1}{1-\phi_t}\left(\frac{v_{t+1}}{u_t f_{t+1}}\right)(1-s_{t+1}-\phi_t f_{t+1}) \\
- \kappa_0 + \kappa_0\left(\frac{u'(c_{t+1})}{u'(c_t)}\right)(1-s_{t+1})\Bigg] = 0,
\end{aligned}
$$

where p_t denotes the marginal product of labor. The intuition for the first line of this first order condition is as follows. Suppose the planner creates enough extra vacancies at time t to create an extra position at t, and reduces vacancies at $t+1$ to leave employment at $t+1$ unchanged. The benefit of this at time t is that there is p_t extra output, but one fewer worker is unemployed so the flow value

[23] See Wright et al. (2021) for a discussion of the relationship between directed search and constrained efficiency.

of unemployment z_t is lost. Furthermore, the planner has to create $\frac{1}{1-\phi_t}\left(\frac{v_t}{u_{t-1}f_t}\right)$ vacancies at time t, because each vacancy has a filling rate of $\frac{v_t}{u_{t-1}f_t}$, and the elasticity of matches with respect to vacancies is $1-\phi_t$. On the other hand, fraction $1-s_{t+1}$ of the extra hired workers are still employed at $t+1$ so the planner can create correspondingly fewer vacancies then, but also has to create extra vacancies at $t+1$ in proportion to $\phi_t f_{t+1}$, since there will be one fewer unemployed at the start of $t+1$ per extra worker hired at t, and so, all else equal, this will lead to $\phi_t f_{t+1}$ fewer matches at $t+1$ because the elasticity of matches with respect to unemployment is ϕ_t. The intuition for the second line is simply that hiring an extra worker at t costs κ_0 but requires $1-s_{t+1}$ fewer hires at $t+1$ since $1-s_{t+1}$ of extra employees will still be employed then.

Now, suppose there are no goods or financial market frictions, so that the marginal revenue product of labor satisfies $r_t = p_t$. The constrained efficient wage is then one such that, given this wage, firms' optimal hiring decisions will achieve the same allocation as the planner. Substituting (3.4) and (3.2) into the planner's first order condition, to eliminate p_t, reveals that this implies that the constrained efficient wage, w_t^E should satisfy:

$$\begin{aligned}\frac{\phi_t}{1-\phi_t}\left(\kappa_0+\frac{\kappa_1 v_t}{u_{t-1}f_t}\right) =& w_t^E - z_t\\ &+E_t\left[(1-\rho)\frac{u'_{c_{t+1}}}{u'_{c_t}}(1-s_{t+1}-f_{t+1})\frac{\phi_t}{1-\phi_t}\left(\kappa_0+\frac{\kappa_1 v_{t+1}}{u_t f_{t+1}}\right)\right],\end{aligned}$$

Comparing this with equation (3.8) makes clear that the constrained efficient wage is the same as the Nash wage, except setting the worker bargaining power β equal to ϕ_t (i.e. the well known Hosios condition), and ignoring the fixed cost of hiring κ_0. Now, first consider the Cobb-Douglas matching function, which holds $\phi_t = \phi$ fixed. Then, since our results in Section 3.3 made clear that adding a fixed cost of hiring will make the Nash wage less cyclical (thereby raising estimates of the Nash wage elasticity) it follows that the constrained efficient wage will be more procyclical than the Nash wage if $\kappa_0 > 0$. Now, alternatively, with a matching function that has an elasticity of substitution strictly less than 1, ϕ_t will tend to be procyclical, since it is decreasing in the number of unemployed and increasing in vacancies. This adds an additional procyclical element to the efficient wage, which is increasing in ϕ_t. This further accentuates the tendency for the constrained efficient wage to be more procyclical than the Nash wage.

Thus, we may conclude that if wages were set in a way consistent with constrained efficiency, as in many directed search models, wages would be more procyclical than Nash wages and we would expect the Nash wage elasticity to be greater than 1. As a consequence of this, it follows that the low Nash wage elasticity we find in the data not only indicates that directed search models may have difficulty in matching the empirical behavior of wages over the business cycle, but also indicates that the movement of wages over the business cycle is likely to be more rigid than is consistent with constrained efficiency.

3.5.3.2 Staggered Wage Bargaining:

The first model of rigid wages we consider is a staggered wage bargaining model based on Gertler and Trigari (2009), henceforth GT. In this model, each firm pays all its workers the same wage. At the start of each period, each firm draws an idiosyncratic iid shock which determines whether it renegotiates its wages with its workers or not. Fraction λ of firms retain the same wage as they had in the previous period, while fraction $1 - \lambda$ of firms negotiate a new wage with all their workers according to Nash bargaining,

To isolate the effect of wage rigidity on unemployment fluctuations, we amend GT's staggered wage bargaining model so that it is consistent with the modeling framework outlined in Section 3.2, with as few additional assumptions as possible. This allows us to compare the wage implied by staggered wage bargaining with the Nash wage derived in Section 3.2.

To this end, we make one change to the wage bargaining framework in GT. In GT, the firm, when negotiating wages with its existing workers, takes into account that this wage will affect the wages of new workers it hires. The effect of this assumption in GT is to lead firms to bargain as if their discount rate is somewhat lower, and the effective firm discount rate is time varying and depends on the firm's expectations of its future hiring, and also of how its future hiring will be affected by the wage rate it negotiates. This adds considerable complexity to the bargaining problem, and also entails that the GT wage bargaining solution depends on convex costs of hiring, which are a feature of GT but are inconsistent with the framework of Section 3.2. To avoid this complexity and to maintain consistency with the framework of Section 3.2, we do not assume convex costs of hiring. Furthermore, we assume that when a firm renegotiates its wages with workers, the outcome of this negotiation depends only on the match surplus the firm earns from its current workers, and the match surplus of these workers, and does not depend on the effect of wages on future hiring.

As such, we assume that, when a firm renegotiates wages with workers, the renegotiated wage for each match k satisfies the Nash bargaining solution:

$$\mathscr{W}_t^k - \mathscr{U}_t = \beta[(\mathscr{W}_t^k - \mathscr{U}_t) + (\mathscr{J}_t^k - \mathscr{V}_t^i)] = \beta[(\mathscr{W}_t^k - \mathscr{U}_t) + \mathscr{J}_t^k],$$

where β is the worker bargaining share and the Bellman values $\mathscr{J}_t^k$, $\mathscr{V}_t^i$, $\mathscr{W}_t^k$ and $\mathscr{U}_t$ are as defined in Section 3.2 and so evolve according to the same Bellman equations as in Section 3.2.

Since all matches in the same firm are the same, we let $\mathscr{J}_t^i(w)$ denote the match surplus of the firm i if it pays the wage w. Likewise, $\mathscr{W}_t^i(w)$ is the match surplus of the worker in firm i if they are paid wage w.

Let $\overline{\mathscr{J}}_t^i$ be the expected value of a firm i at the start of the period t, before it discovers whether or not it will renegotiate its wages that period. That is:

$$\overline{\mathscr{J}}_t^i = \lambda\, \mathscr{J}_t^i(w_{t-1}^i) + (1-\lambda)\, \mathscr{J}_t^i(w_t^{*i})$$

where w_{t-1}^i is the wage paid by the firm in the previous period, and w_t^{*i} is the wage that would be negotiated if the firm renegotiates its wages.

Define $\overline{W}_t^i$ similarly.

The Bellman equations in Section 3.2 imply that:

$$\begin{aligned}
\mathscr{J}_t^i(w_t^*) &= \overline{\mathscr{J}}_t^i - (w_t^{*i} - \overline{w}_t^i) + \lambda(1-s_{t+1})E_t[m_{t+1}(\mathscr{J}_{t+1}^i(w_t^{*i}) - \overline{\mathscr{J}}_{t+1}^i)] \\
\mathscr{W}_t^i(w_t^*) &= \overline{\mathscr{W}}_t^i + (w_t^{*i} - \overline{w}_t^i) + \lambda(1-s_{t+1})E_t[m_{t+1}(\mathscr{W}_{t+1}^i(w_t^{*i}) - \overline{\mathscr{W}}_{t+1}^i)],
\end{aligned}$$

where the expected wage of at the start of the period (before it is known whether renegotiation will happen) is:

$$\overline{w}_t^i = \lambda w_t^{*i} + (1-\lambda) w_{t-1}^i.$$

When wages are renegotiated, the bargaining solution satisfies:

$$\mathscr{W}_t^i(w_t^{*i}) - \mathscr{U}_t = \beta[(\mathscr{W}_t^i(w_t^{*i}) - \mathscr{U}_t) + \mathscr{J}_t^i(w_t^{*i})],$$

Combining this with the previous two Bellman equations and rearranging, we obtain:

$$\left(\frac{\beta}{1-\beta}\right)\left(\overline{\mathscr{W}}_t^i\right) = \overline{\mathscr{J}}_t^i - \frac{w_t^{*i} - \overline{w}_t^i}{1-\beta} + \lambda(1-s_{t+1})E_t\left[m_{t+1}\left(\left(\frac{\beta}{1-\beta}\right)\left(\overline{\mathscr{W}}_{t+1}^i - \overline{\mathscr{J}}_{t+1}^i\right)\right)\right]$$

Averaging across all firms (and so dropping i superscripts), log-linearizing around the steady state and rearranging, we obtain:

$$-\frac{\hat{\beta}_t}{1-\beta} = \left(\frac{w}{\beta h}\right)\left(\frac{1}{1-\lambda(1-s)(1-\rho)} \cdot \frac{\lambda}{1-\lambda}\right)(\hat{w}_t - \hat{w}_{t-1})$$

Substituting in equations (3.21) and (3.16), we obtain:

$$\hat{w}_t - \hat{w}_{t-1} = (1-\delta)E_t[\hat{w}_{t+1} - \hat{w}_t] + \frac{\delta}{\psi}(\hat{w}_t^N - \hat{w}_t) \tag{3.39}$$

where

$$\begin{aligned} \delta &= 1-(1-f-s)(1-\rho) \\ \psi &= \frac{w-z}{\beta h}\left(\frac{1}{1-\lambda(1-s)(1-\rho)}\right)\left(\frac{\lambda}{1-\lambda}\right). \end{aligned}$$

Equation (3.39) is the wage setting equation for the staggered wage bargaining model. This resembles New Keynesian Phillips curve equations, in that the rate of growth of wages depends on the deviation of wages from the negotiated (Nash) level, and also depends on the expected rate of growth of wages next period.

We now study the business cycle properties of the staggered wage bargaining model. In particular, we keep the business cycle model assumptions unchanged from Section 3.5.2, except that we replace the wage equation (3.37) assumed there, and replace it with the wage equation (3.39). For simplicity, we limit attention to the $\hat{z}_t = 0$ case. In Table 3.6 below, we consider various values of λ. Since λ is the probability that a firm is unable to renegotiate its wages in a period, this parameter determines the level of wage rigidity in the staggered bargaining model.

For each value of λ, the second column of the Table 3.6 shows the estimated NWE obtained from simulating the model for 10,000 periods and performing an OLS regression of the wage on the Nash wage in the simulated data, in accordance with the first approach we used to estimate the Nash wage in Section 3.3. The remaining columns of Table 3.6 mirror those Table 3.5: they show the relative volatility of unemployment implied by the staggered bargaining model, and then they show the relative volatility predicted by the Fundamental Surplus formula, given the estimated NWE in the second column.

GT originally calibrated λ at 0.88. As Table 3.6 shows, this is consistent with an NWE of 0.03, which is rather lower than the majority of our empirical estimates. On the other hand, if the model is recalibrated with $\lambda = 0.66$, the implied NWE is slightly higher than the majority of our estimates. This suggests that the staggered bargaining model is consistent with the level of wage rigidity we estimate, provided that λ is calibrated at a somewhat lower level than assumed by GT.

The last two columns of Table 3.6 also reveal that the NWE implied by the staggered bargaining model provides a useful guideline to the cyclical volatility of unemployment in that model, using the Fundamental Surplus formula. Nevertheless, the staggered bargaining model delivers a rather lower level of unemployment volatility than one would expect given the Nash wage elasticity. This is because the staggered bargaining model delivers a low NWE only in the short run (recall that we estimated the NWE delivered by the model using hp-filtered simulated data). In the long run, the staggered bargaining model implies that the wage should be fully flexible. Consequently, firms may e.g. hire more in recessions than the low NWE of the staggered bargaining model would suggest, because they anticipate that while wages are sticky now, they will fall in future.

Table 3.6: NWE and Relative Unemployment Volatility under Staggered Bargaining

λ	Estimated NWE	Relative Volatility of Unemployment	Volatility Implied by FS formula
0.01	0.99	0.70	0.95
0.02	0.98	0.70	0.96
0.11	0.83	0.72	1.12
0.22	0.61	0.78	1.50
0.33	0.41	0.89	2.20
0.44	0.26	1.10	3.44
0.55	0.14	1.46	5.76
0.66	0.07	2.14	10.25
0.77	0.03	3.47	19.06
0.88	0.01	6.58	33.75
0.99	0.00	18.07	45.68

3.5.4 Alternating Offer Bargaining

Following Hall and Milgrom (2008) and Christiano et al. (2016) we consider a wage setting protocol in which wages are determined by an alternating offer bargaining game. The details of the bargaining game follow Christiano et al. (2016) (henceforth CET) exactly. We suppose that each period is divided into $M = 60$ sub-periods. At the start of the first sub-period, the firm makes an initial wage offer to the worker, which the worker can accept or reject. If the wage offer is rejected, play proceeds to the next sub-period. In odd sub-periods, if the firm and worker have not yet reached agreement, then the firm gets to make a wage offer to the worker, which the worker can accept or reject. Every offer the firm makes costs the firm γ in processing costs. In even sub-periods, if the firm and worker have not yet reached agreement, then the worker makes an offer to the firm, which the firm can accept or reject. If neither have reached agreement by the end of the last sub-period, the match terminates. Additionally, each time an offer is rejected, bargaining breaks down and the match is terminated with probability ς. In each sub-period in which the two sides have not reached agreement, the worker does not produce the flow value r_t^k and does not get paid, but receives the flow value of unemployment z_t.

CET show that the solution of the bargaining game is that the worker accepts the firm's offer in the first sub-period and the wage satisfies:

$$\mathscr{J}_t = \mu_1(\mathscr{W}_t - \mathscr{U}_t) - \mu_2\gamma_t + \mu_3(r_t - z_t)$$

where $\mu_i = \frac{\alpha_{i+1}}{\alpha_1}$ and

$$\alpha_1 = 1 - \varsigma + (1-\varsigma)^M$$

$$\alpha_2 = 1 - (1-\varsigma)^M$$

$$\alpha_3 = \alpha_2\left(\frac{1-\varsigma}{\varsigma}\right) - \alpha_1$$

$$\alpha_4 = \left(\frac{1-\varsigma}{2-\varsigma}\right)\frac{\alpha_2}{M} + 1 - \alpha_2$$

where ς is the probability that bargaining breaks down each day.[24].

Into this, we substitute the firm's Bellman equation to eliminate r_t, substitute that $\mathscr{W}_t - \mathscr{U}_t = \frac{\beta_t}{1-\beta_t} \cdot \mathscr{J}_t$ (where β_t is the worker's share of match surplus) to eliminate $\mathscr{W}_t$ and $\mathscr{U}_t$, and substitute the firm's optimal hiring decision $\mathscr{J}_t = h_t$ to eliminate

[24]Here, we have written CET's result in terms of our own notation

$\mathscr{J}_t$. We obtain the following form of the alternating offer bargaining solution:

$$h_t = \mu_1 h_t \cdot \frac{\beta_t}{1-\beta_t} - \mu_2 \gamma + \mu_3 (w_t + h_t - E_t[m_{t+1}(1-s_{t+1})h_{t+1}] - z_t)$$

Log-linearizing this around the steady state and using equations (3.16) and (3.21), we obtain:

$$\hat{w}_t = \left(\frac{\mu_1}{\mu_1+\mu_3}\right)(\hat{w}_t^N - \hat{w}_t^A) + \hat{w}_t^A - \left(\frac{\mu_3}{\mu_1+\mu_3}\right)(1-\delta)E_t\hat{w}_{t+1}^A + \left(\frac{\mu_3}{\mu_1+\mu_3}\right)(1-\delta)E_t\hat{w}_{t+1} \tag{3.40}$$

where, $\hat{w}_t^N$ is the Nash wage (deviation from the steady state) and $\hat{w}_t^A$ is the deviation of an alternative wage, given by:

$$\hat{w}_t^A = \left(\frac{h}{\mu_3 w}\right)\left[1-\mu_3-\frac{\beta\mu_1}{1-\beta}\right]\hat{h}_t + \frac{(1-\rho)(1-s)h}{w}E_t\left[(\sigma c_t - \sigma c_{t+1}) + \hat{h}_{t+1} - \frac{s\hat{s}_{t+1}}{1-s}\right] + \frac{z}{w}\hat{z}_t. \tag{3.41}$$

As with the staggered wage bargaining model above, we study the cyclical properties of the business cycle model in Section 3.5.2, replacing the wage equation there (equation (3.37)) with the wage setting equations (3.40) and (3.41). Again, we fix $\hat{z}_t = 0$. The key parameter that determines the rigidity of wages under alternating offer bargaining is the probability ς that bargaining breaks down. As we did for λ with the staggered bargaining model, we vary the level of this parameter, estimate the resulting NWE on model simulated data and compare the cyclical volatility of unemployment with what would be implied by the model's implied NWE and the Fundamental Surplus formula. The results are in Table 3.7 below. We see that, across parameter values, the alternating offer model delivers an NWE around 0.7, significantly higher than almost all our estimates. The volatility of unemployment predicted by the model is close to what one would expect from its NWE, based on the fundamental surplus formula.

Table 3.7: NWE and Relative Unemployment Volatility under Alternative Offer Bargaining

ς	Estimated NWE	Relative Volatility of Unemployment	Volatility Implied by FS formula
0.000	0.69	0.84	1.34
0.002	0.70	0.87	1.32
0.003	0.71	0.89	1.30
0.005	0.72	0.90	1.29
0.006	0.72	0.91	1.28
0.008	0.73	0.92	1.27
0.009	0.73	0.93	1.26
0.011	0.74	0.93	1.26
0.012	0.74	0.94	1.25
0.014	0.74	0.94	1.25
0.015	0.75	0.95	1.25
0.017	0.75	0.95	1.24

It is surprising that the alternating offer bargaining model generates an NWE so close to 1, given that a key purpose of the model was to generate wage rigidity. Inspection of equation (3.40) and (3.41) indicates that a major reason for the cyclicality of the wage under the alternating offer bargain is the cyclicality of hiring costs. CET likewise note that the alternating offer bargain fits their macroeconomic data far better with fixed rather than variable hiring costs. For this reason, we also study the alternating offer bargaining model with primarily fixed hiring costs. Specifically, we reduce κ_1 from 0.43 to 1, and recalibrate κ_0 to maintain the average total hiring cost. Results of the simulations of this model are presented in Table 3.8 below. In this case, the results are very sensitive to ς, but with values of ς close to zero, the model achieves an NWE close to 0.3 and correspondingly larger unemployment fluctuations. Thus, our results support CET's assertion that fixed hiring costs help the model fit the data. This is still higher than many of our estimates. As such, our estimates generally support a level of wage rigidity that is as great or greater than implied by the alternating offer bargaining model.

Table 3.8: NWE and Relative Unemployment Volatility under Alternative Offer Bargaining
with Mostly Fixed Hiring Costs

ς	Estimated NWE	Relative Volatility of Unemployment	Volatility Implied by FS formula
0.000	0.33	3.33	3.99
0.002	0.42	2.71	3.19
0.003	0.50	2.36	2.75
0.005	0.56	2.14	2.48
0.006	0.61	1.99	2.30
0.008	0.64	1.87	2.16
0.009	0.68	1.79	2.06
0.011	0.71	1.72	1.98
0.012	0.73	1.67	1.92
0.014	0.75	1.63	1.87
0.015	0.77	1.60	1.83
0.017	0.78	1.57	1.80

3.6 Conclusion

In this paper, we develop a new measure of aggregate real wage rigidity, the Nash wage elasticity. The NWE is simply the elasticity of the measured marginal cost of labor with respect to the Nash wage, where the bargaining share is set to equal the actual wage in a steady state. A completely rigid wage implies an NWE of 0, and if wages were set by Nash bargaining then the NWE should in theory equal 1.

We build a broad modeling framework that encompasses a wide variety of cases studied in the literature, and show that the framework delivers equations for the worker share of match surplus and the Nash wage that can be calculated from empirical data. When taking these equations to US data from 1979-2012, we find that the worker share of match surplus is strongly countercyclical and that the Nash wage is substantially more procyclical than the observed cost of labor. These findings hold for a range of different measures of the cost of labor that range from practically acyclical to strongly procyclical.

Across 180 regressions, which variously use different wage measures and use either simple OLS or different instruments for the Nash Wage, we obtain estimates of the Nash wage elasticity that are mainly between 0 and 0.1. We only obtain Nash wage elasticity estimates above 0.65 for the most procylical series of the cost of labor (the user cost from the NLSY) and, even with this series, only for specifications that assume a relatively high value of fixed hiring costs and/or the opportunity cost of employment.

We investigate the business cycle implications of the our small estimated values for the NWE. We find that a small NWE makes an enormous difference to fluctuations in unemployment. We show that there is a tight link between the NWE and the Fundamental Surplus of Ljungqvist and Sargent (2017), with smaller values of the NWE greatly shrinking the Fundamental Surplus and increasing the volatility of unemployment in SAM models with shocks to the marginal revenue product of labor. In a simple SAM model with such shocks, an NWE of 0.1 yields fluctuations in unemployment that are more than seven times as large as occur when the NWE is 1. In this sense, the vast majority of cyclical movements in unemployment can be attributed to the effects of wage rigidity.

Finally, we compare our estimated NWE with the implications of various non-Nash models of wage setting. This includes constrained efficient wage setting, as in many directed search models, and three models of sticky wages. We find that our estimated NWE implies much more rigid wages than is consistent with the constrained efficient wage model. Our NWE estimates do suggest wages may be less rigid than assumed by Hall (2005) and Gertler and Trigari (2009), but more rigid than assumed by Christiano et al. (2016).

Appendix C

Appendix to Chapter 3

C.1 Data Sources

Table C.1: Data Sources

Name	Description	Source	ID	Notes
u	Unemployment Rate	Fred	UNRATE	
u^l	Unemployment Level	Fred	UNEMPLOY	
u^s	Short tern unemployment	Fred	UEMPLT5	unemployed for less than 5 weeks
e^l	Employment Level	Fred	CE16OV	
c	Personal Consumption	Fred	A794RX0Q048SBEA	
	labor Productivity	BLS	NFBUS	
w	NLSY New Hire Wage	Basu and House (2016)		
w^{CES}	Average Hourly Wage	Fred	AHETPI	
w^{UL}	NLSY User Cost of labor	Kudlyak (2014)		reported in Basu and House (2016)
z	Elasticity	Chodorow-Reich and Karabarbounis (2016)		
f	Finding Rate	Calculated		Eq. (3.26)
s	Separation rate	Calculated		Eq. (3.27)
v	Vacancy rate	Petrosky-Nadeau and Wasmer (2013)		
Forecasting VAR				
1-Year T-Rate	Market Yield on U.S. Treasury Securities at 1-Year Constant Maturity	Fred	GS1	
Real GDP	Real Gross Domestic Product	Fred	GDPC1	
GDP Deflator	GDP Implicit Price Deflator	Fred	USAGDPDEFQISMEI	
Shock Measures				
MP	Romer and Romer Narrative Series	Romer and Romer (2004b)		updated in Wieland and Yang (2020)

Chapter Four

The Role of Investor Sentiment in Commodity Price Behaviour

Evidence via Time-Varying Causality Tests

Introduction

One of the key issues facing academic economists and investors is whether the behaviour in the last fifteen years or so of commodity prices offers evidence that the returns in this asset class may have been driven by speculation. Abnormal behaviour of commodity contract futures arose during and in the immediate aftermath of the Global Financial Crisis ('GFC') of 2007-08. The spectacular run-up in many commodity prices and their sudden collapse prompted a wave of regulatory actions and enhanced scrutiny towards '*speculators*', defined as those myopic agents motivated only by short term profits. These speculators are investment funds, swap dealers and pension funds seeking a long-only exposure to commodities. Throughout this paper I will refer to them as 'index investors' since they often seek to replicate a commodity index through 'exchange traded funds' (ETFs) or over-the-counter (OTC) derivatives either by purchasing long futures and rolling them over near maturity or entering into return swap contracts.[1]

A key distinguishing features of index investment is its long-only perspective and its approach to commodity contracts for diversification. As a phenomenon, passive index investment emerged before the GFC, when, observing that the price behaviour in commodities is different from the asset prices in their portfolios, investors realized they could treat futures contracts in commodities as assets, including them as part of an improved, diversified portfolio. Such behaviour pertained to the futures contract rather because the investors rarely, if ever, held the physical commodities: investors would see contracts rolled over rather than mature into an obligation to buy and store the commodity.

The volume of index investments in commodity future contracts increased dramatically in the early 2000s and reached its zenith with the Great Financial Recession. Whilst ETFs have been popular from the early '90s, commodity indexed funds originated later, beginning from bullion market during its boom in the early 2000s, and have later expanded to oil, energy markets and agricultural commodities.[2] Oil based index investment has been particularly successful since its inception, fuelled by oil movements on the back of the Second Gulf War. Along with that surge in passive investments, it developed also the knowledge that basket investment was the main driver of the concomitant observed price increase in crude oil products and energy, i.e. that the market of commodities was 'financializing'. In my study the historical aspect is important, as I introduce a time-varying test that needs interpretation with the help of historical evidence.

The rift between speculators and commercial investor is captured by regulatory efforts made to prevent and limit excessive speculation. Interestingly the CFTC clearly distinguishes between index investors and hedgers (e.g. in setting contract limits). Index investment is seen by policymakers as a way of investing in commodities speculative in nature, whereas hedging is carried out for necessity by commercial and noncommercial investors driven by fundamentals. Kang et al.

[1]The Commodity Futures Trading Commission ('CFTC') defines index investment as a *'long-side exposure to a broad index of commodity prices as an asset class. This trading is not based on a view about current or expected individual commodity prices, as would be the case for most speculative trading'* (Commodity Futures Trading Commission, 2006).

[2]UBS sponsored 'Street Track Gold Trust' (now SPDR Gold) is reported to be the first commodity-based ETF (FT Article (a)) and in 2005 Oil Securities launched the first crude-only ETF (FT Article (b))

(2020) notes how this segmentation appears artificial given that that commercial hedgers' position may still have a 'speculative' component reflecting a view over future prices, e.g. they may hold a long position in the expectation of a longer term appreciation of the underlying commodity.

While market participants are aware that institutional investors, including hedge funds and commodity index traders, play an important role in determining commodity futures prices, it is not at all clear whether such activities lack a fundamental basis in supply and demand, or whether they are simply part of a mechanism whereby information on market fundamentals becomes incorporated in competitively determined prices.

A debate has taken place along two axes: between practitioners and academics and within the two groups. From the former group, debate was brought into focus by Masters (2009, 2008) and Soros (2008) who argued that commodities were becoming viewed by some as a distinct and mainstream asset class, leading to market outcomes where speculative positions taken by an influx of new traders into oil and other commodity futures markets became essential in explaining the run-up in commodity future prices seen in the years prior to the GFC.

This process has been described as the '*financialization*' of commodities markets (Fattouh et al., 2013; Tang and Xiong, 2012; Irwin and Sanders, 2011). Financialization means that the increased traded volume of commodity contracts results in higher future prices. A fundamental contribution to the financialization literature is Singleton (2014), who argues that index investment flows predict oil future prices.

The notion of financialization lies at the centre of an academic debate. The controversy can be deconstructed in 3 arguments: 1) whether increased financialization has been the major driver of commodity futures prices over the period since the GFC; 2) whether it has created '*excessive speculation*' manifesting itself as bubbles; 3) and whether the large positions taken by such investors have led to self-fulfilling movements in prices that are tantamount to '*market manipulation*'.

That the entrance of index investor in the commodity market has changed the composition of such markets is undisputed, for example, Tang and Xiong (2012) quote a CFTC staff report that estimated the total value of various commodity index-related instruments purchased by institutional investors to have increased from around USD 15 billion in 2003 to over USD 200 billion in mid-2008 and to over USD 300 billion in 2011. The Bank of International Settlements ('BIS') estimates a volume of commodity forward and swaps derivatives outstanding in 1H-2008 at USD 8 trillion (excluding gold and precious metals), an all-time high and significant one-off deviation from the pre-2008 average of USD 1.3 trillion and post-2008 of USD 1.7 trillion.[3]

The sheer increase in index investment volumes does not provide prima facie evidence of a process of financialization and the high frequency nature of trading data makes it difficult to establish a clear causal link based on a chronological chain of events. The financialization hypothesis posits that index investors move the market as to cause a demand-driven increase in prices. Lacking convincing evidence, one could argue for an opposite causal relation: long-only investors may have been attracted to commodity futures because of exogenous shifts in fundamentals.

There is no academic consensus on whether financialization causes artificially high commodity prices, in terms of bubble dynamics or other significant departures from fundamentals, but empirical evidence is somewhat mixed. The most studied commodity market is the crude oil one. In my opinion, the most influential paper in this literature is Singleton (2014), who find evidence of one-directional causality from index investment to crude oil prices. Offering a credible evidence-base, Singleton's paper was highly influential in the financialization debate that has, to some extent, relied upon testimony and reports based only on descriptive statistics.

With this paper, I aim to revisit Singleton's question by looking at a number of commodities, not just crude oil. Another issue of concern comes from recent developments in the econometric literature, where proper Granger causality testing is seen to need to take account of the type of data that was observed during and in the aftermath of the GFC. Extant causality testing methods based on stationarity or unit-root non-stationarity do not account for the behaviour in commodities over the sample period of interest, especially the run-up and abrupt collapse seen in commodities around 2008.

Accordingly, I will use a new methodology in causality testing, performing the time-varying Granger causality test introduced by Shi et al. (2018, 2020). This test builds upon the recently proposed PSY test based on Phillips et al. (2015a) and uses a recursive-evolving algorithm. My starting point is a recent paper by Gilbert (2018) who used a battery of static Granger causality tests to assess unidirectional

[3] BIS OTC Derivatives Outstanding Table D5.2

causality from two index investment proxies to twelve future contracts in non-ferrous metals, energy and agriculture commodities. I start with a replication of Gilbert's results, broadly confirming them while, at the same time, correcting a small but fundamental calculation error in creating the proxies. Using the more powerful and robust PSY testing strategy, I then apply a similar testing methodology to a number of commodities.

The results herein improve upon results in the literature: firstly, the time-varying aspect of the Shi et al. (2018). test reduces the risk in the application of static Granger causality tests of the atypical data around the GFC contaminating the results based on the whole sample. Secondly, the robustness of this test allows the results to be taken at face value. Contrasting the results with the results of the static test provides practitioners with an indication of the (lack of) efficacy of the static test.

The paper offers three main contributions. Evidence is provided to show that the results of Singleton (2014) are essentially due to the atypicality of the data on crude oil around the GFC, not financialization. In short, the abrupt behaviour of the data around then overwhelms the static Granger causality test based on the whole sample. Secondly, I show that the behaviour of some commodity futures prices does seem to provide limited evidence for financialization. The results therefore corroborate the results found by Gilbert (2018) using a different model; while milder, they still offer some meaningful refinement to mainstream results in the financialization literature. Thirdly, my work illustrates why commodity traders and practitioners should now be using the new, robust Shi et al. (2018) test when conducting Granger causality analyses.

4.1 Does Index Investment Impact Future Prices?

In this Chapter I ask whether the volume of index investment bears an impact on commodity future prices. There is no consensus on this research question even though historically high commodity prices have been moving in tandem with index investment. In this section I explore the most relevant studies that have rigorously endeavoured in tackling the question at hand with causal inference.

A key reference in this literature is Singleton (2014), which was the first paper to provide a credible evidence-based approach to the financialization debate. Singleton (2014) stands in stark contrast with the early discussion around excessive speculation that have motivated financial legislation.

Causality from index investment to bubble dynamics is taken for granted and stated outright in CFTC briefings and documentations without being underpinned by formal statistical interpretation and reasoning. This rhetoric, according to which nefarious and predatory speculation artificially inflated commodity prices in a bubble-like fashion, stems from the observed disastrous consequences of the financial crisis for the global economy and became dominant in the policy-making discussion. The policy debate is an interpretation around the phases of expansion and collapse of commodity prices pre and post GFC.

Hedge-fund manager Micheal W. Masters was amongst the first practitioners to voice a strong view regarding the causal relation linking long-only investments to future prices, contributing to its crystallisation in the regulatory consensus. In his hearings before the CFTC and the Commission on Homeland Security and Governmental Affairs, (Masters, 2009, 2008) he condemned index investors as the culprits of 2008 boom-bust.[4] The core of Master's argument is that the sheer weight of index investors' on future contracts demand caused a displacement in the future market, with an increase in prices well above fundamentals and solely motivated on irrational expectations. But again, this argument is made informally in a series of testimony given to governmental bodies.

Whilst regulators and market practitioners accept as a datum that the increase in volume in index investment resulted in artificially higher prices, academics are divided on whether this is (and was) indeed the case. An alternative explanation of wild commodity price movements might pertain to shifts to market fundamentals, which might have underpinned speculators and hedgers trades alike. Hence, the current literature developed around Singleton (2014) and in relation to it. I note how financial practitioners, on the other hand, have responded more to the Masters' arguments.

Academics are perhaps reluctant in embracing the view that index investment a cause of future prices as it contradicts the Efficient Market Hypothesis (EMH) that asset prices incorporate already all information available, as famously elaborated in Fama (1970). The latter explanation is more benign than the former and does not question the allocative efficiency of markets. But bubbles themselves are evidence of departures from the EMH, and there is a thriving literature on explosive behaviour

[4] For a detailed overview of this debate see Irwin and Sanders (2011)

of commodity prices.

Fanelli (2015) reviews a range of cases where the Efficient Market Hypothesis does not hold with respect of commodity prices and longer-run cointegration relations appear in certain commodity contracts.

Gilbert (2010b) contributed to establish the Granger-causality testing methodology applied to autoregressive-lag models in the index investment literature. He finds evidence of speculative bubbles in non-ferrous metals and energy prices. Gilbert and Pfuderer (2014a), however, find negative results from index investment to food prices through a conventional Granger-test. Gilbert (2010a); Gilbert and Pfuderer (2014b) notes that even if the EMH holds, there could be a higher-frequency (almost contemporaneous) interaction between index-investment and commodity prices, which is not necessarily apparent at longer lags (e.g. weekly).

The bulk of academic consensus have rejected the financialization hypothesis but with an important caveat: the GFC represented an anomaly during which commodity future prices and index investment displayed a heightened degree of correlation and predictability. Even if they end up rejecting a whole-sample causality linking index investment to commodity prices, many of the papers reviewed in this section find some evidence of non-linear and time varying relation. In particular, they acknowledge that the GFC may have represented a one-off anomaly that temporarily broke down the proper market functioning. For instance, Irwin and Sanders (2011) comment on a Gilbert (2010b) paper stating that identified bubble days are all around the GFC.

An early proponent of financialization evidence in commodity market is Phillips and Yu (2011), who show a statistically significant bubble migration from house prices to commodity prices between the bust of the housing bubble in December 2007 and the onset of the GFC in September 2008. They do not find evidence of bubbles in agricultural commodities within the same time frame. Phillips and Yu (2011) take on bubble migration is similar to the point of view expressed by Soros (2008), in stating that further to the subprime crisis, speculative flows have flown from the housing market to commodities, where perceived risk was low and expected returns high. Phillips and Yu (2011) paper, however, does not offer a causal view, but is rather concerned with identifying behavioural linkages behind bubble origination and migration.

In a series of papers, Irwin and Sanders; Irwin tackle directly the Master's argument against index investment. Irwin and Sanders (2011) review much of the literature on index investment-commodity prices causality pointing out logical inconsistencies in the Master's hypothesis as well as statistical shortcomings with quantitative papers testing causality. In essence, they find the Master's hypothesis not convincing.

Similarly, Chari and Christiano (2019) reject that financialization has an impact on spot prices. At the same time they refute the CFTC consensus on speculators/hedgers, showing in a structural model that hedgers and other players trade to insure each other against demand shocks. Basak and Pavlova (2016) build a structural model for index investing and find that this type of investing is associated with higher future prices. Hence, the fragmentation of the literature in regards to financialization is even more apparent as there is little consensus on the stylised facts a structural model should be able to explain.

A fundamental paper in this literature is Singleton (2014), who finds a strong and significant relation between index investment and excess returns on the crude oil future. This finding is consistent with a theoretical model where heterogeneous agents disagree on economic fundamentals generating bubble inflation/bust.

Hamilton and Wu (2015) regression analysis using CFTC data does not find evidence of financialization outside the GFC, when there was a heightened correlation between index investment and future prices.

Fattouh et al. (2013) break down the Masters hypothesis arguing that the question at hand is much more suited for a structural model rather than a qualitative exercise as the one put forward by Masters. They also do not find Singleton (2014) results conclusive in any way, as the co-movement of index-investment and future prices does not imply causality even if index-investment moves ahead of prices.

Sanders and Irwin (2011) reject the null hypothesis of Granger-causality in U.S. grains futures (corn, soybeans, CBOT and KCBT wheat) using an autoregressive-distributed lag model ('ADL'). On the same set of commodities and with an instrumental ADL model, Gilbert and Pfuderer (2014a) detect Granger-causality in the same grains and oilseeds complex.

Another strand of literature is concerned with commodity risk premiums: Hamilton and Wu (2014) note how the oil futures risk premium has decreased with time, probably reflecting the rise in index investors' share.

Whilst Hamilton and Wu tackle financialization and risk premiums change as two separate research questions, others have studied the issue of financialization as a fact that affects risk premiums. Cortazar et al. (2021) take index-bundled commodities as a standalone asset class and state a factor model to compute their risk premium on an indexed basis.

Adams and Glück (2015) highlight a 'style effect' of portfolio management. The fact that during and in the immediate aftermath of the GFC commodity futures have established themselves as an asset class for long-only investors represents a structural break and a permanent departure from the earlier status quo. I note how the concept of 'risk spillovers' mentioned in Adams and Glück (2015) resembles the notion of 'bubble migration' of Phillips et al. (2015a), i.e. in 2008 commodities attracted investors fleeing from other asset classes and this inflow made commodity returns more volatile. This phenomenon extended past the crisis as commodities became more correlated to the stock market as institutional investors included them in their portfolio construction and rebalancing.

My research question relates to the work of Peter C. B. Phillips, as I use a very recent statistical methodology that he developed with other co-authors with the aim of uncovering causal relations and their changing points. In a similar research strand, Phillips et al. (2015a,b) devised another testing strategy for the real time detection of asset bubbles based on recursive computations of the Augmented Dickey-Fuller Test, developing on an earlier article (Phillips and Yu, 2011). In a host of papers, (Figuerola-Ferretti et al., 2015; Figuerola-Ferretti and McCrorie, 2016; Figuerola-Ferretti et al., 2020) apply this novel test to most traded commodities: oil, precious metals and non-ferrous metals. Although they are concerned with bubble behaviour, their work uncover some evidence circa investor exuberance on commodities as an individual asset class. They find that precious metals sporadically displayed the hallmarks of speculative bubbles, but when they did, it was around the launch of certain specific ETF for silver and palladium. For non-ferrous metals, the authors paint a more nuanced picture: they find mild explosivity in copper, nickel, lead, tin and zinc, all commodity contracts analysed but aluminium. The departure from fundamental value of non-ferrous metals may be explained by the availability of their supply, which is driven by mining and industrial processes and therefore cannot easily adjust to periods of increased demand. This demand-supply interaction can give the rise to bubble dynamics, as opposed to energy and other commodities.

4.2 Global Shocks and Commodity Price Cycles

The starting point for my analysis is naturally the Gilbert's dataset, which is comprehensive of the weekly returns of twelve commodity future contracts and two index investment proxies. Gilbert's sample goes from 11 April 2006 to 27 December 2016, so I retain this original sample for the sake of replicating Gilbert's analysis but my first extension of his work is along the time dimension, as I update the series to reach December 2020.

The first thing I notice in the data is the presence of several distinct commodity price crashes that highlight the presence of short-lived commodity cycles within the 10 years sample. There is a vast literature about commodity prices swings based on trend-cycle decomposition techniques. Extremely low-frequency movements in commodity prices are evidence of a 'super-cycle' that spans decades from peak to trough and is a fundamental driving factor for the individual commodities. Fernandez et al. (2020) estimate that the current super-cycle had peaked around the financial crisis and, similarly, Reinhart et al. (2016) concludes that the most recent non-oil super-cycle peaked in 2011 and had yet to bottom out when the paper was written. Hence my analysis will partly cover the current super-cycle, starting four years before its peak and covering the entirety of its decline. This includes the aftermath of the GCF and the beginning of the Covid-Crisis. Although daily data are available up to present day, I decide to cut my sample in December 2020, to exclude recent shocks in commodity prices that might be still ongoing.[5]

I retain a focus on a much shorter time period than several decades, therefore I am much more concerned with higher-frequency episodes in the commodity market, which directly affect my analysis.[6] The onset of the GFC in September 2008 represents the peak of both the index investment proxy and most of the commodity price in analysis and marks the start of a period of heightened correlation between index investment and commodity prices (see Fig. 4.3). This empirical fact is also the motivation for this paper.

Fernandez et al. (2020); Reinhart et al. (2016) are also investigating the linkages from a commodity super-cycle to the macroeconomy, with a particular concern towards emerging economies. Smaller markets tend to be more reliant on fewer export

[5]Notably the US presidential election in November 2020 and all the inflation expectations that were created around then, the unfolding of the Covid-Crisis and, more recently, the War on Ukraine.

[6]Sadly the starting point for my empirical analysis is dictated by the availability of index investment data, from 2006

products and therefore their economy rests more exposed to boom-bust dynamics that in turn has the potential of spilling over to exchange rate and currency Végh (2013).

For that reason, also macro- investors and multilateral development institutions closely follow commodity prices and have developed their own indexes. I report the major four indexes in Fig. 4.1. These are namely: the S&P GSCI Index, the Bloomberg Commodity Index, IMF All Commodities and World Bank Energy and Non-Energy.

S&P GSCI is the most popular commodity index and often taken as a benchmark by index investors, so its analysis is relevant to this paper as index investor take long position on synthetic products replicating the GSCI.[7]

The composition (and shape in the chart below) of S&P GSCI indicator reflects the higher weighting of crude oil products, whereas the Bloomberg Commodity index shows a more persistent decline after the GFC since it has a bigger share of grains and the distribution across energy commodities is flatter.

In general, commodity prices dropped in several other instances besides the GFC, albeit by never as much as during the financial crisis. If I pick the days in which the S&P GSCI Index dropped by more than 6%, I can distinguish few other commodity crashes that are significant for my analysis (see Fig. 4.1).

The bulk of days in which S&P lost more than 6% is clustered around the GFC: between October 2008 and February 2009. In May 2011 commodity index reached its highest level before dropping on the back of political unrest in Libya and concerns for a contagion effect to the Middle East. November 2014: China GDP Deceleration; The last big commodity prices drop in my sample coincides with the World Health Organisation declaring a Pandemic in March 2020.

[7] Another complication for index investment analysis stems from the way these baskets are rolled forward when the underlying commodity contracts are near expiration. Hamilton and Wu (2015) specifically test causality in the roll window.

Figure 4.1: Commodity Price Indexes.

Vertical red lines represent days in which the S&P GSCI Index lost more than 6%. All the indexes attribute a higher weighting to energy commodities.

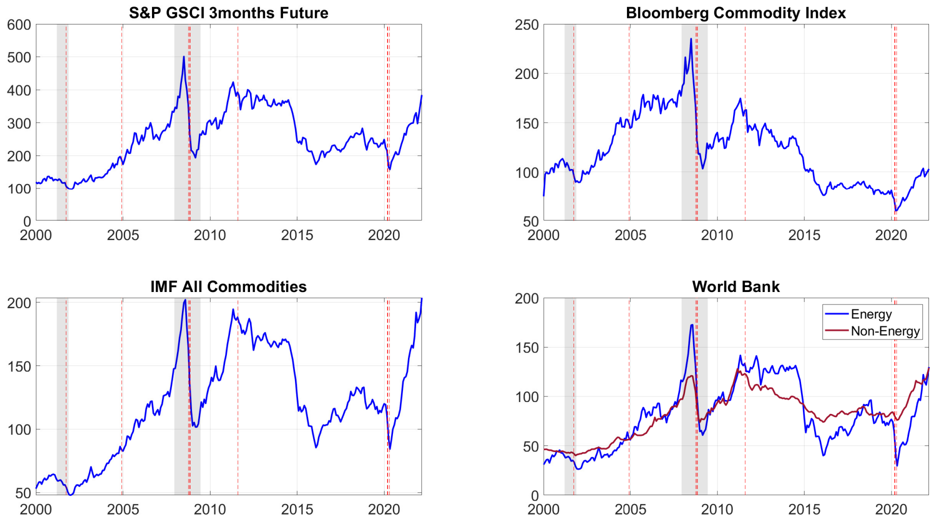

4.3 The Data

4.3.1 Commodity Contract Prices

The present work builds on Gilbert (2018) analysis of Granger causality from index investment proxies to prices of twelve future contracts. These commodities are: soft wheat, corn, soybeans and soybean oil, aluminium, copper, lead, nickel and zinc, ICE Brent and West Texas Intermediate (WTI) crude oil and natural gas.

The only difference between my analysis and that of Gilbert is that I source the data from Datastream, which reports CBOT prices at settlement time rather than at closing.[8] This decision helps me dodge issues related to rolling (which I touched on briefly in footnote 7). Datastream series are available in form of a continuous time series splicing together distinct future contracts. This necessity arises as future contracts have maturity dates and one has to switch from the expiring contract to the nearest.[9]

In my statistical analysis I use log-returns of weekly commodity prices to assess

[8]This makes a minimal difference for the statistical analysis.

[9]London Metal Exchange Contracts use a different averaging method, to which I am neutral as the continuous series is available in Datastream.

their percentage change. Although daily settlement data are available, the weekly frequency is required in order to have data at the same frequency as the index investment proxy.

I present summary descriptive statistics in Table 4.2. Returns have naturally 0 mean. Most volatile commodities are NYM WTI and Natural Gas. When the whole sample is considered, natural gas becomes the most volatile commodity. The series have all positive kurtosis, resulting in heavy tails and a degree of skew. In the full sample these results appear exacerbated: the empirical distributions have more skew and are more platykurtic. I tabulate histograms and (Fig. D.3) QQ plots for the log returns (Fig. D.2) finding strong visual indication of non-normality.

Hence, I carry out some additional tests to verify whether the commodity prices are normally distributed: the Jarque-Bera and Kolmogorov-Smirnov. To these two tests I append a Ljung-Box test for serial autocorrelation. In all instances I reject the null hypothesis of non-normality, furthermore, Soybeans Oil, Copper and ICE Brent and NYM WTI are autocorrelated. The autocorrelation result is stronger across the whole sample than between 2006 and 2016.

Table 4.1: Normality and Autocorrelation Hypothesis Testing

Test	Null Hypothesis	Alternative Hypothesis
Jarque-Bera	Data are Normally Distributed	Data are not Normally Distributed
KSS	Data are Normally Distributed	Data are not Normally Distributed
Ljung-Box	Data are Not Autocorrelated	Data are Autocorrelated

These features of financial return series are in accordance with the wider literature on market prices. The fact that returns are not normally distributed was reported in Mandelbrot (1963, 1972, 1967).[10]

[10]The cited studies are concerned with commodities: Mandelbrot (1963) uses the daily spot prices of cotton as primary example, citing how his results extend to wheat and other edible grain. Mandelbrot (1967) extends the analysis to what and railroad stocks.

Table 4.2: Descriptive Statistics and Normality Tests

	Aluminium	Copper	Lead	Nickel	Zinc	ICE Brent	NYM WTI	Natural Gas	Wheat	Corn	Soybeans	Soybeans Oil
Descriptive Statistics (March 2006- December 2016)												
Mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Standard Dev.	0.03	0.04	0.05	0.05	0.04	0.05	0.06	0.05	0.05	0.05	0.04	0.03
Skewness	0.02	-0.33	-0.09	-0.06	-0.39	0.06	-0.15	0.50	0.21	-0.13	-0.51	0.09
Kurtosis	4.14	5.04	5.00	4.59	4.71	5.09	4.94	4.51	3.91	5.02	4.22	3.92
Maximum	0.12	0.15	0.21	0.21	0.14	0.23	0.22	0.30	0.18	0.18	0.10	0.15
Minimum	-0.12	-0.17	-0.18	-0.19	-0.19	-0.17	-0.25	-0.18	-0.18	-0.21	-0.15	-0.12
Descriptive Statistics (March 2006 - December 2020)												
Mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Standard Dev.	0.03	0.03	0.05	0.05	0.04	0.05	0.06	0.07	0.05	0.04	0.03	0.03
Skewness	0.09	-0.38	-0.07	-0.06	-0.38	-0.22	0.55	0.42	0.33	-0.13	-0.45	0.09
Kurtosis	4.45	5.50	5.65	4.85	4.89	16.54	32.76	5.13	4.29	5.51	4.52	4.13
Maximum	0.12	0.15	0.21	0.21	0.14	0.41	0.69	0.33	0.20	0.18	0.11	0.15
Minimum	-0.12	-0.17	-0.18	-0.19	-0.19	-0.43	-0.55	-0.25	-0.18	-0.21	-0.15	-0.12
Normality Tests												
Normality Tests (March 2006 - December 2016)												
Jarque-Bera Test												
Statistic	30.490	107.375	94.564	59.107	82.503	102.226	90.113	76.813	23.626	96.503	59.434	20.331
P-value	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***
KSS Test												
Statistic	0.464	0.457	0.444	0.441	0.451	0.447	0.439	0.434	0.447	0.443	0.459	0.460
P-value	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***
Ljung-Box Test												
Statistic	28.716	29.764	12.622	14.840	17.298	28.144	31.662	15.177	17.431	20.237	20.451	33.258
P-value	0.093*	0.073*	0.893	0.785	0.634	0.106	0.047**	0.766	0.625	0.443	0.430	0.031**
Normality Tests (March 2006 - December 2020)												
Jarque-Bera Test												
Statistic	68.92	218.66	226.29	110.00	132.91	5889.25	28459.94	168.65	67.61	203.49	99.83	41.64
P-value	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***
KSS Test												
Statistic	0.463	0.459	0.449	0.443	0.453	0.446	0.440	0.431	0.449	0.447	0.461	0.463
P-value	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***	<0.001***
Ljung-Box Test												
Statistic	28.18	33.77	14.07	21.01	19.75	38.24	35.85	22.20	24.66	23.42	22.37	36.23
P-value	0.105	0.027**	0.827	0.397	0.474	0.008***	0.016**	0.330	0.215	0.269	0.321	0.014**

4.3.2 The Index Investment Proxy

Starting from January 2007, the CFTC started to un-bundle index investment positions from commercial and noncommercial ones in the framework of a wider review of its published reports (Commodity Futures Trading Commission, 2006). Information on index investment integrated the already existing weekly report called 'Commitment of Traders' (COT), [11], which states the aggregate trading positions of traders across selected commodity contracts. The noncommercial index traders are managed funds, pension funds and other institutional investors and commercial ones OTC hedgers (Commodity Futures Trading Commission, 2006).

The Supplemental Commitment of Traders ('SCOT') lists the open-interest positions of index traders in the twelve future contracts that the CFTC deems particularly impacted by index investments (i.e. that are frequently included in benchmark indexes).[12] SCOT headline series is available weekly, it's published on Friday with data updated as at the Tuesday prior and discloses open interest positions of commercial, noncommercial and index traders.[13] The drawback of the SCOT data is that they are only available for selected agricultural commodities.

A second data series on index investors is the Special Call, so called because it aggregates all index positions of investors who have received a 'special call' from the CFTC's Division of Market Oversight.[14]

The Special Call reports also the dollar value of open index investment interest positions besides the equivalent number of future contracts and is not only limited to agricultural commodities. Hence, the Special call is more detailed than the SCOT as the classification of index investment activity is made at the reporting entities level. The trade-off is that the Special Call was initially published quarterly (between September 2009 and June 2009) and then it became monthly. Since it was time-consuming for the CFTC to gather and publish these data, its publication was stopped in October 2015 also as a consequence of the reduced interest of market participants after the GFC.[15]

[11] The CFTC publishes COT reports since 1962

[12] Other than in Gilbert, CFTC index data have been used widely in the literature (inter alia by Sanders and Irwin (2011); Hamilton and Wu (2015); Kang et al. (2020)), surprisingly, Shahzad et al. (2021) the paper most similar to ours in the econometric analysis, does not use CFTC data to inform their time-varying Granger Test.

[13] Link: (Commitment of Traders)

[14] The CFTC issued 43 special calls to commodity index funds and swap dealers. (CFTC Explanatory Notes)

[15] CFTC Press Release 7282-15

In his econometric analysis, Gilbert (2018) uses the dis-aggregated data on index investment as proxies for overall index activity. These are proxies because they do not provide a complete figure of all index investment but rather the best effort of the Regulator to distinguish index investors from other players. The key limitation of the SCOT is that trading positions are classified in accordance to what is the predominant activity of the reporting entity. The Special Call reports are provided by the entities specifically prompted to provide these figures, and hence may exclude other important parties.

To uncover the portion of index investment carried out through OTC swap agreements, the CFTC issued a 'Special Call' to index funds and swap dealers known to engage in index investment. These investors are required to disclose the notional value of their index investment positions and equivalent number of future contracts.

Gilbert uses the SCOT data to derive his proxies for index investment in the following fashion: first, he retains four individual SCOT series for the respective agricultural contracts listed above, secondly he aggregates the SCOT across all commodities to have an aggregate measure of index investment. In both cases, prices are kept constant as at the first Special Call observation, making the investment proxies Laspeyres indexes. I follow exactly his aggregation method.

$$Index_t = \sum_{j=12}^{12} x_{j,t} p_j \tag{4.1}$$

Where the index investment proxy in time t is the sum across the twelve commodity contracts of the number of contracts x times their price p as at December 2007 relative to CBOT wheat. Commodity future prices and the SCOT proxy are plotted in Figure 4.3. I note that implementing this indexing methodology, I correct a small but fundamental error in the Special Call reported by Gilbert (2018). Whilst intending to do as stated above, he applied to commodity contracts the price-weights offset by one position, resulting in erroneous weights for all commodity and in particular 0 weights for WTI oil and natural gas, thus severely underestimating the impact of index investment in the energy sector.

Figure 4.2: Commodity Prices (blue line - LHS) plotted Against the Index Investment Proxy (red line - RHS)

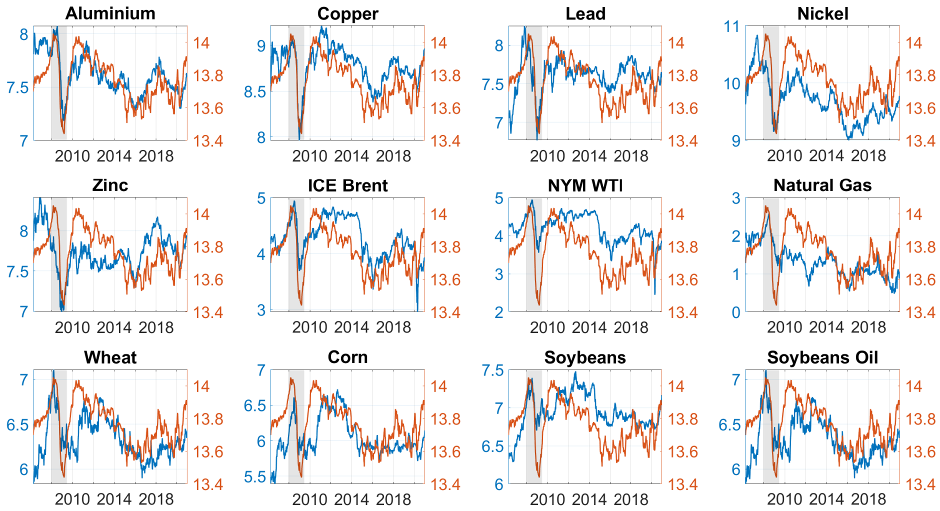

Figure 4.3: Commodity Prices (blue line - LHS) plotted Against the Special Call Investment Proxy (red line - RHS)

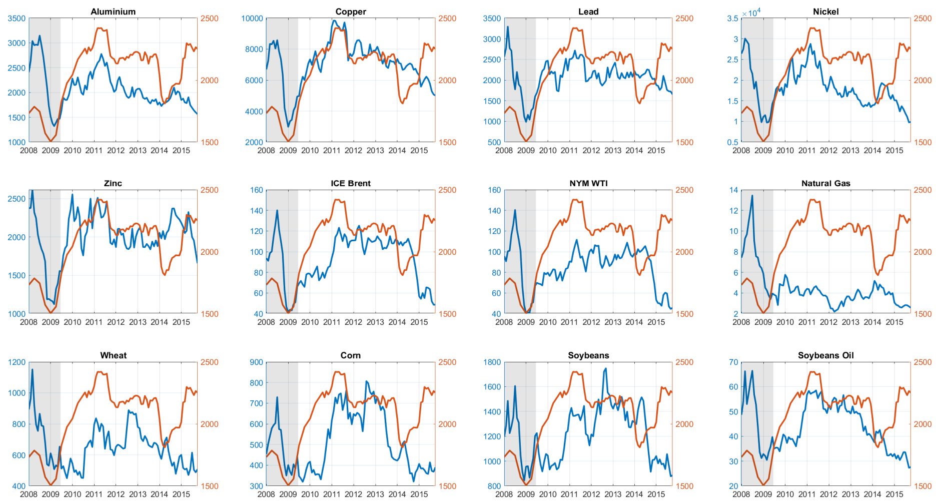

4.4 Replication of Gilbert (2018)

The first step that I take towards a more robust modelling of time-varying causality is the replication of Gilbert (2018). Re-appraising Gilbert's results offer a natural starting point for this Chapter as I can point out strengths and weaknesses of his approach and compare his findings from a standard Granger test against a PSY test.

As in Gilbert (2018), I perform a static conventional Granger causality based on the F-test using an Autoregressive Distributed Lag (ADL) Model for the twelve commodity contracts. This type of testing strategy is common in the literature. Gilbert's ADL takes the form of the following linear regression for the j-th commodity contract:

$$\Delta ln(Price_j) = \mu + \sum_{i=1}^{p} \alpha_{j,i} \Delta ln(Price_{j,t-1}) + \sum_{i=1}^{p} \beta_{j,i} \Delta ln(Index_{j,t-1}) + \varepsilon_t \tag{Unrestricted Model}$$

Where p is the maximum lag order. *Index* is the selected proxy for Index Investment. A contract specific proxy is only available for Agricultural commodities. In this context, the Granger causality test is an F-test that compares that model with its restricted version, where $\beta_j = 0$.

$$\Delta ln(Price_j) = \mu + \sum_{i=1}^{p} \alpha_{j,i} \Delta ln(Price_{j,t-1}) + \varepsilon_t \tag{Restricted Model}$$

And the F-Stat is:

$$\text{F-Stat} = \frac{(RSS_{restricted} - RSS_{unrestricted})/p}{RSS_{unrestricted}/(T - 2p - 1)} \tag{4.2}$$

Where RSS is the Residuals Sum of Squares. Whole sample Granger test results reported in Gilbert (2018) are below successfully replicated (See Table 4.3). I also fail to reject the null of no Granger causality for Aluminium, Copper, Nickel, Zinc, ICE Brent, Corn and Soybeans Oil.

Table 4.3: Replication of Gilbert Tables 1,2,3.

F-statistic based Granger Causality Test. Same ballpark results but slightly different since I use a different dataset of prices at settlement instead that at closing.

	Aluminium	Copper	Lead	Nickel	Zinc	ICE Brent	NYM WTI	Natural Gas	Wheat	Corn	Soybeans	Soybeans Oil
SCOT												
F-Stat	3.874	4.120	1.940	2.989	3.356	3.794	1.448	0.0458	0.365	3.015	1.449	7.739
P-Value	0.002***	0.006***	0.164	0.051*	0.067*	0.0519**	0.236	0.830	0.546	0.0295**	0.229	0.000***
DoF Nominator	5	3	1	2	1	1	2	1	1	3	1	2
DoF Denominator	549	553	557	555	557	557	555	557	557	553	557	555
Special Call												
Fstat	9.252	1.079	1.665	4.517	11.28	2.281	4.63	1.31	0.634	0.886	2.159	0.892
P-Value	0.003***	0.392	0.094*	0.036**	0.001***	0.019**	0.034**	0.251	0.533	0.349	0.099*	0.545
Dof Nominator	1	12	12	1	1	11	1	8	2	1	3	10
Dof Denominator	90	68	68	90	90	70	90	76	88	90	86	72

I arrive at the same results as Gilbert, finding Granger Causality from index-investment proxies to Aluminium, Copper, Lead, Nickel, Zinc and ICE Brent and NYM WTI for the period comprised between 2006 and 2016. There is Granger-causality from the SCOT proxy to Corn and Soybeans Oil, and weak causality from the Special Call to Soybeans. I note that the way non-ferrous metals respond to investment proxies is consistent with Figuerola-Ferretti and McCrorie (2016), i.e. hard constraint in metals mining and production result in more volatile prices when the demand picks up.

I find a number of issues with this approach i) Gilbert (2018) did not use any diagnostic for residuals autocorrelation, and I know from the previous section that log-returns in commodity prices display strong indications of non-normality. ii) I correct a typo in Gilbert's calculation of the Special Call, which resulted in a wrong index proxy severely underestimating the impact of energy prices on the overall series. This means that my Special Call results are novel in the literature and I draw comfort from the fact that they confirm and are consistent with the SCOT results. iii) I aggregate the Special call at his lower frequency (monthly) rather than interpolating it as in Gilbert (2018) and and (iv) It has to be recognized that the test statistic involved in applying a Granger causality test to the ADL model has different statistical properties to a test statistic based on an autoregressive model, owing to the presence of the lagged dependent variable. Nevertheless, as I now show, when I apply the same analysis to a VAR model, Gilbert's essential results carry over.

4.5 Time Varying Granger Causality Test

In this section I use similar data as Gilbert to gauge the efficacy of standard Granger causality testing when the data behave in ways such as seen around the GFC and using the recently proposed PSY causality test to examine the extent to which Gilbert's conclusions still hold. To do so, I have i) to recast the time-series model into a vector autoregression; ii) perform an array of evolving-recursive calculations to detect when the causal relation becomes significant.

I offer a novel contribution expanding on the Gilbert (2018) analysis on two dimensions: one is to support the argument that the results based on the time-varying test are the ones that should be taken at face-value. Based on this new econometric evidence, I reinterpret known views and perspectives in the financialization debate. The second contribution is the argument that the PSY test should be preferred to the standard Granger test for statistical reasons. This Chapter shows how the results from the two styles of testing are different and how the PSY test is more insightful for the practitioner.

I use the PSY test intended as the recursive evolving econometric procedure applied to Granger causality testing in Shi et al. (2018). The PSY test is a computation intensive strategy, according to which a standard test is performed within a double recursion in order to detect changes to the causal structure of the data and find in real time the changing points of the causal relation. This test is explicitly designed for data with discontinuities such as the ones seen around the GFC. For a standard Granger test, structural breaks result in a loss of statistical power, whereas in the PSY style of tests, the robustness is drawn from its recursive nature and data discontinuities help detecting the changing points in causality.

The choice to use a VAR model is because VAR model behaviour under non-stationarity and non-normality is well known in the context of causality testing and I am therefore able to mitigate data characteristics that will make the estimation and inference not robust. Specifically, in the case at hand data are stationary but non-normal and I can take advantage on a heteroscedastic consistent version of the Wald test.

4.5.1 The VAR Model

I extend the Gilbert framework by using the PSY algorithm for change detection in causality relations. The PSY algorithm uses a reduced form vector autoregressive model (VAR). In practice, I pair the SCOT index with each individual commodity indexed by i and I test causality from the SCOT and Special Call index investment proxies to the selected commodity log-return. Hence $y_t = [Proxy, Price_i]'$. The VAR model has the conventional textbook representation:[16]

$$y_{i,t} = c + \sum_{j=1}^{p} A_{i,p} y_{i,t-j} + u_{i,t} \tag{4.3}$$

The key difference from the ADL model is that now the index investment proxies are endogenous and respond to commodity prices. A first issue is to select the

[16] For an exhaustive reference: Lütkepohl (2005).

right lag order. Bayesian Information Criterion (BIC) is the most parsimonious yet consistent selection criterion and it would suggest adding just 1 lag across all the contracts. When minimizing the selection criterion, I follow Ng and Perron (2005) and I keep the sample size constant.

I noted in the previous section how the log-return data display strong evidence of non-normality and their empirical distribution resembles more a Student t. Heteroscedasticity is however a well-known feature of financial series and it does not pose a problem for the Granger test since Shi et al. (2018) have devised a heteroscedastic-consistent Granger test based on limit theory.[17]

4.5.2 Granger Test

In a twin set of papers Shi et al. (2018, 2020) lay the theory for using both a standard and a heteroscedastic-consistent Wald statistic based Granger-test in an evolving-recursive fashion to derive a time series of Supremum Norm Wald statistics. PSY heteroscedastic-consistent Granger test takes the following form:

$$W = T_w[\mathrm{R}\,vec(\hat{A})]'[\mathrm{R}((\mathrm{V}^{-1}\hat{\Omega}\mathrm{V}^{-1})\mathrm{R}')]^{-1}[\mathrm{R}\,vec(\hat{A})] \tag{4.4}$$

Where $\mathrm{V} = \hat{\mathrm{Q}} \otimes \mathrm{I}_n$ and $\hat{Q} = \frac{1}{T_w}\sum_{t=Tf_1}^{Tf_2} x_t x_t'$,and $\hat{\Omega} = \sum_{t=Tf_1}^{Tf_2} \hat{\xi}_t \hat{\xi}_t'$ with $\hat{\xi} = x_t \otimes \hat{u}_t$.

$\hat{A}$ represent the coefficient matrix; X is the matrix of lagged y and R is a $[n \times (k^2 p + k)]$ selection matrix that selects the zero constraints, where p are the lags, k the VAR dimension and n the number of restrictions to be tested.

I use 4.4 to test for Granger Causality across two sub-samples (2006-2016 and 2006-2020) and using the Special Call, the SCOT and the Commodity-Specific SCOT. The VAR based heteroscedastic Granger test yields similar results as the ADL model (see Table 4.3). With the SCOT, I find evidence of a Granger-causal relation again in Aluminum, Copper, Zinc and Soybeans Oil, so in fewer commodities as before (I cannot reject the null for ICE Brent and Corn). Using the Special Call I find stronger evidence of a causal relation for all non-ferrous metals and weak significance in NYM WTI. I find that I cannot reject the null for any of the agricultural commodities.

[17] As the alternative for a plain Wald test as in Lütkepohl (2005)

Table 4.4: Heteroscedastic VAR Granger Test

	Aluminium	Copper	Lead	Nickel	Zinc	ICE Brent	NYM WTI	Natural Gas	Wheat	Corn	Soybeans	Soybeans Oil
SCOT Index to Commodity Prices (2006 - 2016)												
χ^2	5.77	10.15	1.72	2.12	2.89	2.03	0.60	0.048	0.32	1.25	1.30	7.19
P-Value	0.016**	0.00***	0.19	0.14	0.09*	0.15	0.43	0.82	0.56	0.26	0.25	0.00***
SCOT Index to Commodity Prices (2006 - 2020)												
χ^2	6.98	12.86	1.75	1.49	3.23	3.97	2.07	0.015	0.19	0.52	0.61	3.54
P-Value	0.00***	0.00***	0.18	0.28	0.07*	0.04**	0.15	0.90	0.65	0.46	0.43	0.06*
Special Call Index Proxy to Commodity Prices												
χ^2	6.86	10.01	10.63	3.49	6.73	7.14	3.03	0.246	0.223	0.67	1.123	1.99
P-Value	0.00***	0.00***	0.06*	0.09*	0.01**	0.00***	0.08*	0.619	0.632	0.412	0.268	0.157
Commodity Specific Index Proxy to Commodity Prices (2006 - 2016)												
χ^2									2.143	2.308	0.106	2.839
P-Value									0.143	0.129	0.745	0.09*
Commodity Specific Index Proxy to Commodity Prices (2006 - 2020)												
χ^2									4.61	0.78	0.01	2.31
P-Value									0.03**	0.37	0.91	0.12
DoF	1	1	1	1	1	1	1	1	1	1	1	1

4.5.3 Time Varying Results

4.5.3.1 Time-Varying Granger Test: the PSY Algorithm

Given a fractional minimum regression sample of f_0 observations, the algorithms is initialised at time f_1, when a first regression is estimated on a sample of length f_0T that goes from 1 to f_2. After this first step I save and store the associated Suptemum norm Wald statistics. With the evolving step the starting point of the algorithm is shifted one period forward at $f_2 + 1$, then the recursion takes place as now two VAR regressions are estimated with two Wald Statistics of length f_0T and $f_0T + 1$. The higher Wald Statistic is then saved as the Supremum Norm. The evolution step is then repeated up to the last observation in the dataset and at that point sub-regression expanding backwards to the first observation.[18]

4.5.3.2 Bootstrap

The critical values are obtained by the bootstrap method laid out in Shi et al. (2018, 2020): i) I estimate a whole sample VAR and store its residuals; ii) fit a second VAR model to a subset of the data (bootstrap window) to which I add the randomly drawn residuals from the firs step; iii) compute the PSY Supremum norm Wald statistics on the bootstrap sub-sample. I repeat these steps 499 times and I use the 95% percentiles of the Supremum norm Wald statistic as my critical values. This testing method is geared to solve the multiplicity issue associated with running a battery of tests for each step forward in the PSY algorithm and ensure that the empirical sizes approximates the nominal size of the Granger test.

[18] I report an example of the algorithm in Matlab pseudo code in Figure D.5.

4.5.3.3 Optimal length of f_0

The testing strategy hinges on a key modelling decisions: the minimum regression window f_0. Such parameter also determines how many observations I have to discard at the beginning of the sample in order to derive a meaningful Supremum norm Wald statistic for the regression starting the algorithm. I would not want to throw away too many observations, as I am chiefly interested in the Great Financial Crisis, which happens relatively early within the sample. I face another trade-off, as longer regression windows may understate more transient episodes of Granger causality. Shi et al. (2018) define the length of the minimum regression window as a fraction of the whole sample f_0T, setting it at 0.24 for the sake of a simulation exercise for which the duration of the causal relation is known ex-ante. They however note how the optimal size of the minimum regression window is data dependent and should ideally be as long as the 'true' causal episode. I do not know ex-ante what is the length and strength of the underlying causal relation. A f_0 which is longer than the duration of the actual causal relation will not result in more testing power. I settle on a minimum window of 104 weeks, conditioning on my expectations of a causal relation of approximately 2 years, consistent with the brewing and the eruption of the GFC in the biennium 2007-09. Hence my f_0 is approximately 18% of the Gilbert sample and 15% of the 2006-2020 sample. I note how the approach to the PSY set up depends on the data characteristics. In Shi et al. (2018), f_0 is 20% of the whole period, controlled over a 3 years bootstrap window. In Shi et al. (2020), the minimum regression window is set at 72 monthly observations, i.e. 10.8% of the whole sample. Shi et al. (2020) control over a period of 1-year. A shorter window will capture more short-lived episodes as the collapse in commodity prices in 2014/15, but downplays the GFC as the W-stat is very high for very short. I investigate alternative minimum regression sizes in the section below on Robustness Checks.

4.5.3.4 Gilbert's Sample

Gilbert reported a statistically significant static causal relation in the non-ferrous metals and oilseeds markets. I confirm this finding as in my time-varying model aluminium, copper and soybean oil are Granger-caused by the index investment proxy for almost the entirety of the 2006-2016 sample (Fig. 4.4).

There is also one-way causality from the SCOT proxy to Zinc and Nickel in a more intermittent way and concentrated at the beginning of the time period. The shape of the time-varying Wald statistic is similar for the two crude products, as they experience an increase in magnitude of the Granger causality relationship coincidentally with the GFC, which quickly tapers and stays below the critical value for the remainder of the sample period.

These findings represent a first sense check for my results. Gilbert's findings carry over to a time varying setting despite of the data properties that make the conventional Granger test less robust than the PSY alternative.

Figure 4.4: Time Varying Wald Statistic

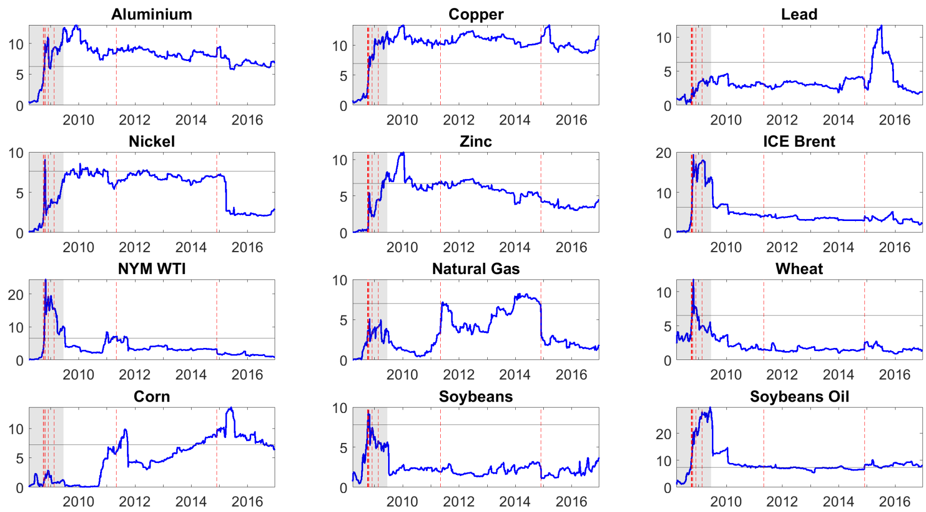

The picture is more mixed for the other two commodity classes: energy and agriculture. Causality in ICE Brent and NYM spikes with the GFC to be re-absorbed by 2009. Evidence of occasional causality appears again in 2011 for NYM, on the back of geopolitical unrest in Libya and a fear of contagion to the Middle East. On the other hand, causality in the Natural Gas is much stronger in 2011 and 2015 than it is during the GFC. This change points in the causal relation can be ascribed to narrative evidence: in 2011, and in 2015 gas prices fell because of a boom in US shale production and a warmer winter than expected.

Next, I substitute the SCOT aggregate investment proxy with the commodity specific proxy. I arrive to similar results, fleshing out more the output of the previous exercises.

Interestingly, as I regress agricultural commodities log-returns and their specific

SCOT index investment proxy on their lagged values, I end up with slightly differently shaped Supremum norm Wald statistic (Fig. 4.5). Wheat has a central peak in 2011 as corn, whereas soybeans and soybeans oil display a different causality profile, with the first displaying a single peak in 2013 and the second having a strong Granger-causality relation at the beginning of the sample, clustered between 2008 and 2010. For Soybeans oil, the strength of the causal relation is significant only during the GFC and wanes in the rest of the sample.

The 2011 spike in corn appears magnified when I consider causality coming from the corn specific index. That event out-weights the 2015 uptick, which in the aggregate case shows with a larger magnitude.

Figure 4.5: Time Varying Wald Statistic

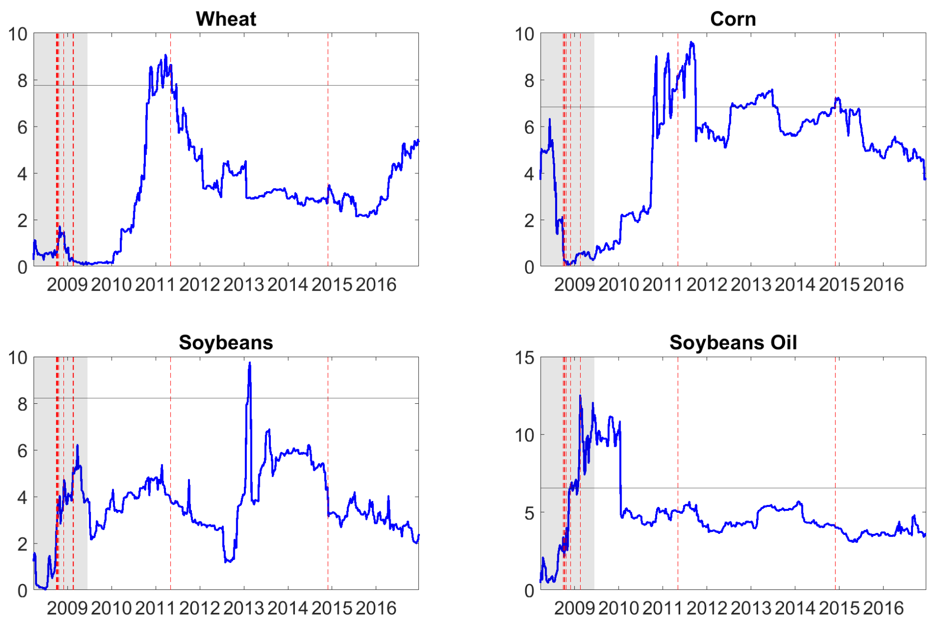

4.5.3.5 Full Sample

The first takeaway of extending Gilbert's sample forward to end of 2020 is that there is no significant change in the causal relation profile between index investment and commodity prices. There are no new episodes in the causal relation linking index returns and future prices in the period 2016-2020. The Covid pandemic has little impact on the strength of the Granger-causal relation as the Supremum norm W statistic does not exceed the critical value for any of the twelve future contracts in 2020 (see Fig. 4.6). This finding qualifies the GFC as a recession different in its own rights, as it stands out more prominently in a 14 years sample. This result is consistent with the bubble migration hypothesis put forward in Phillips et al. (2015a), which I corroborate with clearer evidence of financialization in selected commodities.

Once again this extended sample show a much more significant response of non-ferrous metals, with the Granger-causality of SCOT to Aluminium confirmed throughout the whole sample together with Copper.

Figure 4.6: Full Sample Time Varying Wald Statistic

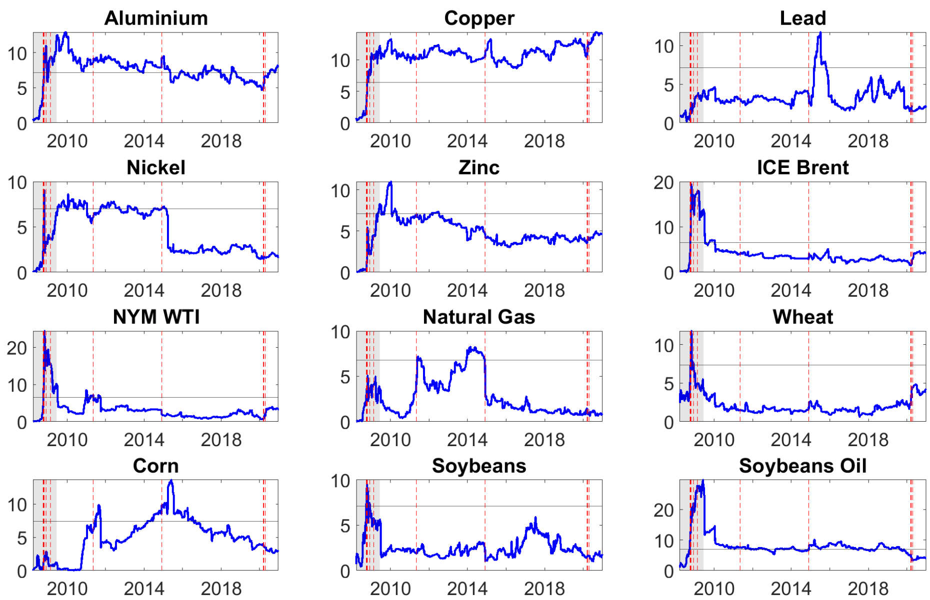

Again, updating the sample period does not yield significantly different causality results also when considering the commodity specific SCOT. The outset of Covid does not trigger a change in causality per se, but movements in commodity could be attributed to idiosyncratic and commodity-specific factors. As opposed to the restricted period, in the full sample I show an additional episode of heightened causality in the wheat market from May 2019, this coincides with an all-time high of the wheat future prices.

I desire to draw the Reader attention to the fact that the time-varying Granger-test results the GFC marks the first point of causal change for many of the commodity considered. Exceptions are lead, natural gas, and corn.

Figure 4.7: Time Varying Wald Statistic

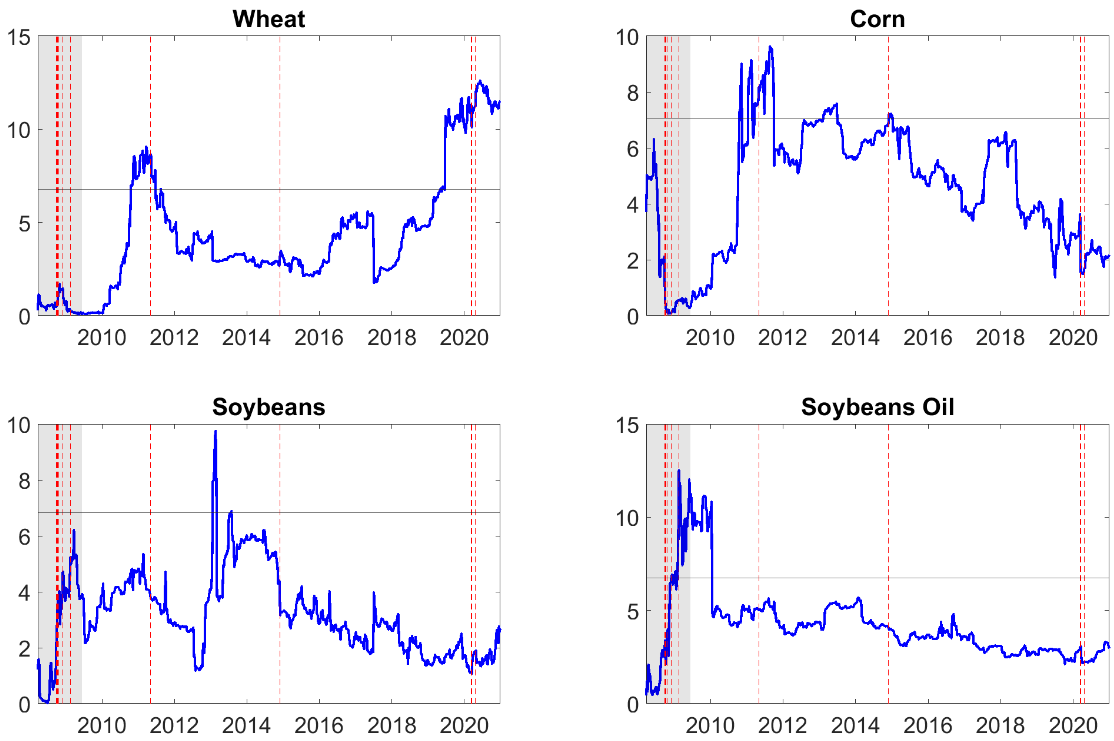

4.5.3.6 Robustness

In addition to the baseline specification, I check for robustness to shorter regression windows, trying two additional specifications, setting f_0T to one year and six months. This specification remains robust with a minimum regression sample of 52 weeks (with equal bootstrap sample). In this latter case however, causality is general more persistent than in the baseline specification. Stronger and almost continuous causality is detected across the whole non-ferrous metals compartment. The causal profile of oil is fairly similar to the baseline specification. I fail to detect causality in the wheat market but I find overwhelming evidence of a strong causal relationship in Corn, Soybeans and Soybeans oil. Interestingly, using a regression window of 1 year instead than 2, shows a cluster of statistically significant Supremum Norm Wald stats starting in February 2015 and ending in June 2015, suggesting that a shorter window may capture occasional episodes of causality that are averaged-out in the baseline specification. However, also this robustness check shows heightened causality during the GFC.

Performing the time-varying Granger causality test using an even shorter window

($f_0 T = 30$ observations), yields a more volatile profile for the Supremum norm statistic. Results are overall in lie with what I found above, but more occasional spikes become visible between 2010 and 2015, years which continue standing apart as the strongest causal episodes. On the other hand, the increased volatility of the Supremum norm may mean that the length of actual causal episodes is indeed longer than six months.

I run the analysis of Granger causality from SCOT to spot prices finding approximately the same time-varying pattern as in Fig. 4.6 but in a weaker form.[19]

4.6 Discussion: Prices and Change Detection

I argue that the heightened correlation between index investment and commodity returns at specific points in time is driving the static results in Table 4.4. Taking an historical perspective and with the help of the time-varying Supremum norm indicator, I uncover some key differences between commodities. This helps me qualify better the Singleton and Gilbert results since elements that are structural to individual commodities are affecting them in different ways. The causality profile and its strength is different across commodity classes.

Secondly I confirm that the GFC is a structural break and was the catalyst for the increase in Granger causality at the beginning of the sample. This stylised fact is accepted by the wide literature I cited in Section 4.1. The explanation for this finding may be the bubble migration which led to the emergence of index investment as a standalone asset class Phillips et al. (2015b); Adams and Glück (2015). Regrettably my index investment proxy starts in 2006 and I cannot verify other episodes of Granger causality connected to the increase in index investment volume in the early 2000.

Taken together these results give us evidence to support a mild re-appraisal of Singleton (2014) result of financialization, extending his findings outside the oil markets.

For oil financialization is localised at a specific point in time, which Phillips et al. (2011) have connected to bubble migration dynamics around the collapse of housing prices. But apart from that oil does not show any other significant events of causality. Non-ferrous metals and Soybeans Oil are seemingly more 'financialized' than the

[19]Results available upon request.

others, for them Granger causality was detected at more frequent intervals and was overall more persistent. Figuerola-Ferretti and McCrorie (2016) note how the rigidity in non-ferrous metals supply might have been a driver for their prices to rise between 2003-07 at times in an explosive-fashion. Given the hard constraints on metals mining and productions, it makes sense that their future prices are more responsive to index money flows.

For other commodities, the changing points in Granger Causality appear as more episodic and, if taken at face value, they show evidence of financialization that extends to agricultural markets.

Figure 4.8: Commodity Prices (blue line - LHS) and a Granger Causality Dummy (1 if Wald stat > Critical Value - RHS)

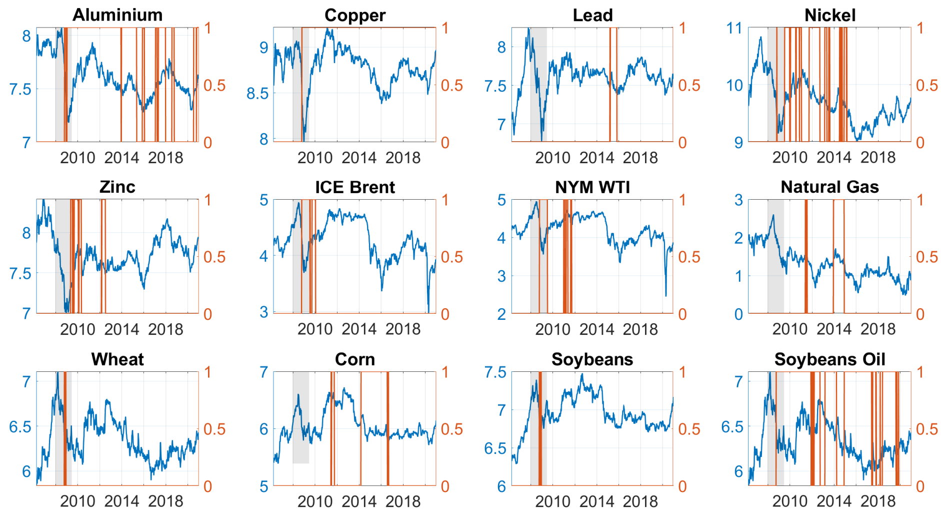

Figure 4.9: Commodity Prices (blue line - LHS) and a Granger Causality Dummy (1 if Wald stat > Critical Value - RHS)

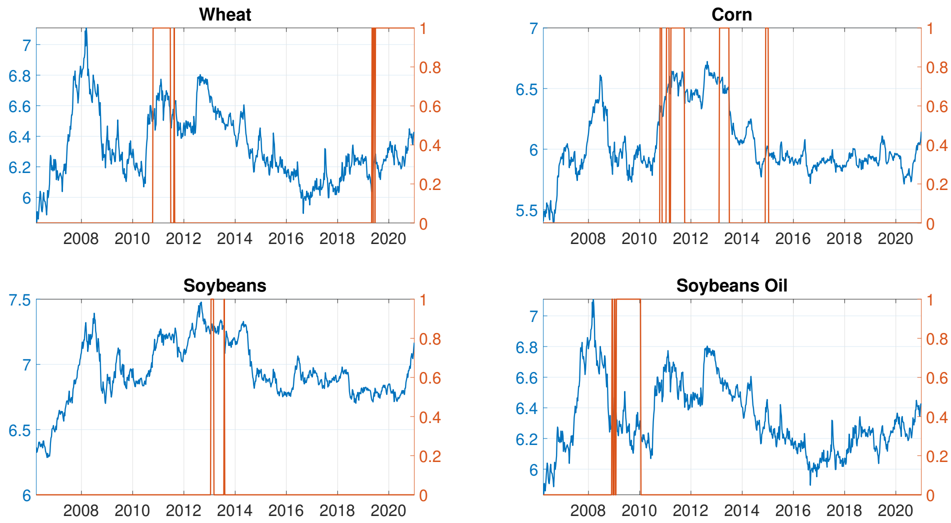

My results help to better nuance Gilbert's results regarding causality in grains and oilseed markets Gilbert and Pfuderer (2014b). They found a causal relation when accounting for contemporaneous causality vis-à-vis a much larger literature that rules out such a link. There is time-varying causality from the aggregate SCOT proxy to the soybeans contract around the GFC, whereas the soybeans oil is intermittently Granger-caused by the same proxies along the entirety of the sample.

4.7 Conclusion

In this paper I analyse the strength and changes in significance in the causal relation linking index investment to commodity futures log-returns. My work is aimed to inform the literature about the evolution of causal relation in such markets and to do so I use a novel and more powerful time-varying Wald test as opposed to the standard static Granger test generally applied in the literature. I do so because whether the commodity market is efficient or not is an open question. Regulators and practitioners have accepted the hypothesis of financialization, whereas the academic literature has failed to reach a consensus.

The literature has settled on using a standard Granger causality test, which can overstate the results in presence of structural breaks. I note how the extreme movements in commodity prices may incorrectly reflect in a static Granger test and I use a different testing strategy to gain a more comprehensive insight on the relevance of index investment flows in the commodity future price formation.

This Chapter replicated Gilbert (2018), correcting a small but important error in calculating the Special Call index proxy. I then applied a standard Granger test based on a Wald statistics in a VAR, providing limited but nonetheless discernible evidence of financialization in non-ferrous metals and some agricultural commodities. This result ties to the important result offered by Singleton (2014), extending it to contracts other than crude oil. My final contribution is applying a more robust testing strategy based on the Shi et al. (2018) (PSY) test to the same twelve commodities, producing a richer set of results that allows us to see (or measure) how the Granger causal relationships have evolved over time.

My results are novel in the literature and they are useful in supporting established facts. Aluminium and Copper display an essentially constant Granger causality, which is intermittent for Nickel, Lead and Zinc.

For crude oil, the heightened causality appears only during the GFC, whilst Natural Gas follows its own cycle. In the case of Natural Gas and agricultural commodities, the causal relation may be reconciled with narrative identifiable evidence: the unidirectional relation stemming from changes in index investment materialises in relation with peak and through of a high frequency commodity cycle.

APPENDIX D

APPENDIX TO CHAPTER 4

D.1 Data Sources

Table D.1: Data Sources

Prices	Exchange	Datastream Ticker
Aluminium	London Metal Exchange	LAH3MTH(P)
Copper	London Metal Exchange	LCP3MTH(P)
Lead	London Metal Exchange	LED3MTH(P)
Nickel	London Metal Exchange	LNI3MTH(P)
Zinc	London Metal Exchange	LZZ3MTH(P)
NYM Oil	New York Mercantile Exchange	NCLCS00(PS)
NYM Natural Gas	New York Mercantile Exchange	NNGCS00(PS)
ICE Brent	Intercontinental Exchange	LLCCS00(PS)
IPE Brent	Intercontinental Exchange	LCRCS01(PS)
Soyabeans	eCBOT	CSNCS00
Corn	eCBOT	CCFCS00
Wheat	eCBOT	CWFCS00
Soyabeans	eCBOT	CZSCS00
Soyabeans	eCBOT	CS.CS00

- Supplemental Commitment of Traders Data (SCOT): `https://www.cftc.gov/MarketReports/CommitmentsofTraders/HistoricalViewable/index.htm`

- Special Call Data: https://www.cftc.gov/MarketReports/IndexInvestmentData/index.htm

D.2 Descriptive Statistics

Figure D.1: Commodity Prices (blue line - LHS) plotted Against the Commodity Specific Index Investment Proxy (red line - RHS)

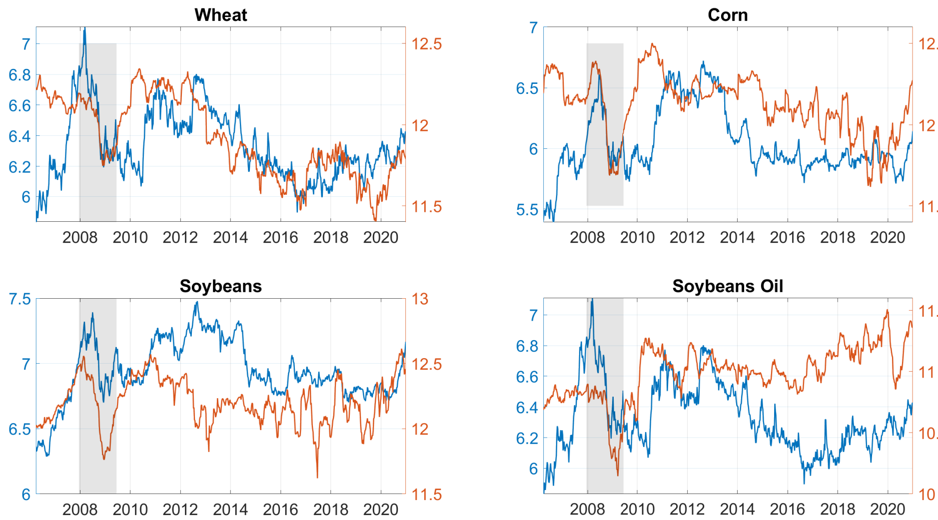

Figure D.2: Q-Q Plots of Empirical Log Returns against t-Location Scale Distribution Quantiles

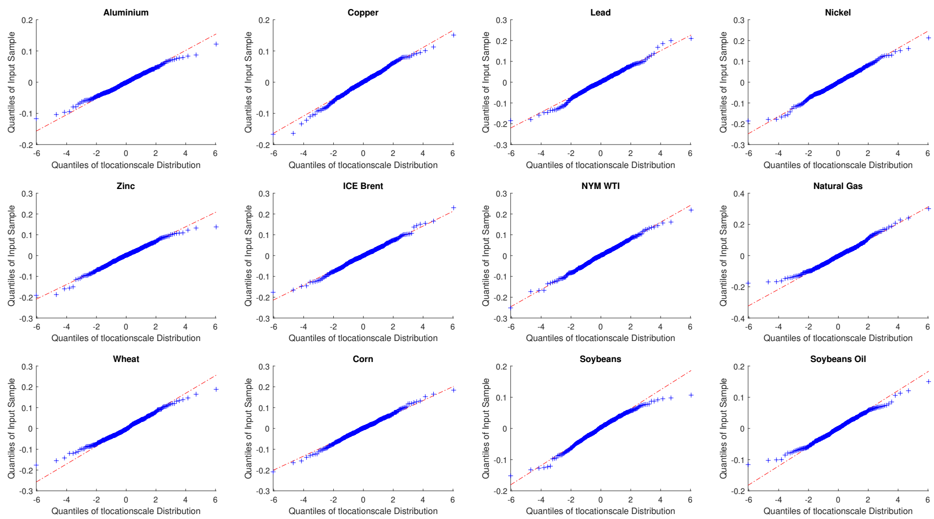

Figure D.3: Histogram of Log Returns against t-Location Scale Distribution

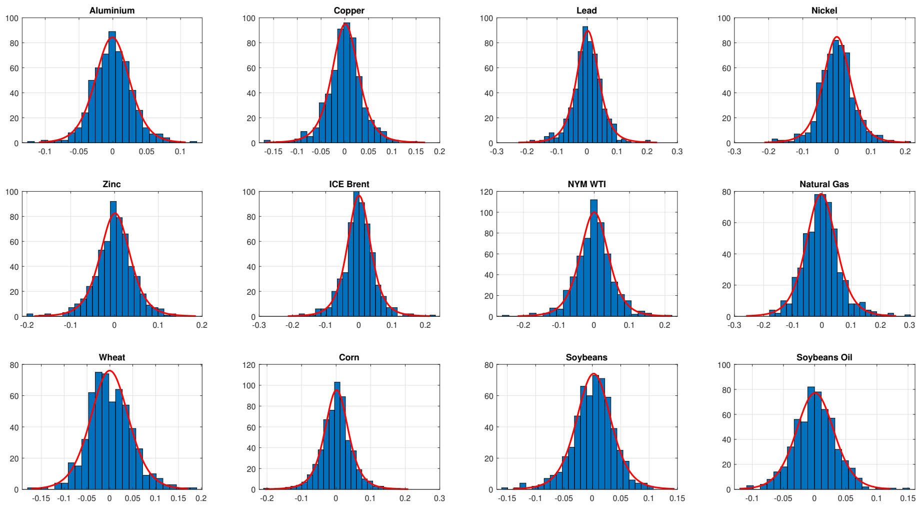

D.3 Recursive Evolving Algorithm

Figure D.4: Evolving Recursive Algorithm Code Example

```
% Set Parameters
Y = [Price, Index];
T = length(Y);
window = minimum regression sample;

%% Start R-E Algorithm
for r2=window:1:T % Ending point of the regression moves forward
     dimW=r2-window+1; % Cut-off point for backward window,
         ensures that the minimum regression sample is respected
     backwardFstat=zeros(dimW,1);
 % For each movement forward of r2, the test is performed
     recursively on a backward window.
      for r1=1:1:dimW
          backwardFstat(r1,:) = Perform Granger Test on sub-sample
              Y(r1:r2)
      end
      if r2==window
          Wstat(r2-window+1,:)=backwardWstat; % W-Stat for first
              period is the F-Stat of regression Y(1:window)
      else
          [Wstat(r2-window+1,:),ind]=max(backwardWstat);% W-Stats
              for subsequent periods are the sup W-Stat of
              regressions Y(r1:r2)
      end
end

% Result
W-stat is the vector of sup W-Stat
```

D.4 Robustness

Figure D.5: Full Sample Granger Test on 52 weeks window

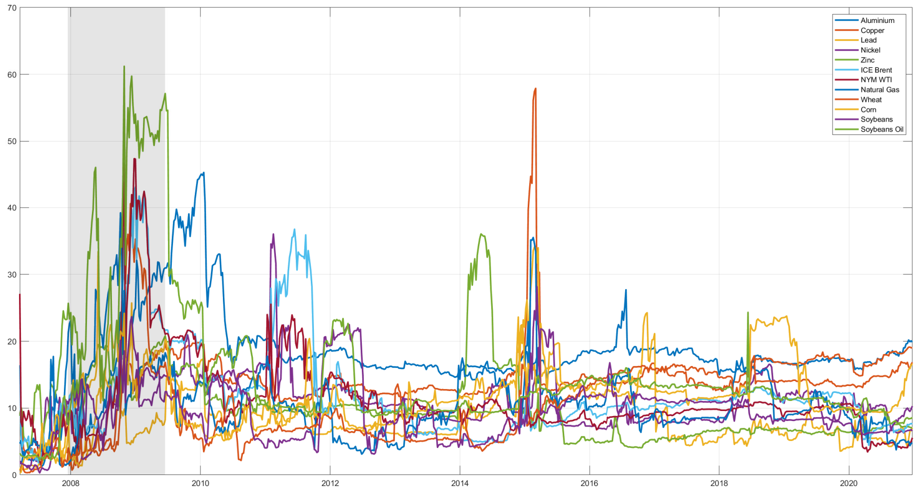

CONCLUDING REMARKS

In this thesis I have presented four Chapters on macroeconomics focussing on instances in which markets behaviour deviated from what we would normally expect under textbook assumptions. My research contributes to the existing literature in three areas: asset price targeting, labour market frictions and commodity prices behaviour.

In Chapter 1 and 2 I analyse deleverage and defaults dynamics following monetary policy shocks. I show empirically how defaults increase rapidly on impulse of an exogenous surge in the interest rate while household debt and house prices follow a more sluggish adjustment path. This suggests that part of the household deleverage is not benign and happens because households are not able to pay back their debt, rather than by a structured deleverage through pre-payments or refinancing.

The fact that debt moves at a lower frequency and insolvencies have a higher frequency component has policy implications, as some papers in the literature have put forward the case for an activist monetary policy stance geared to curb credit cycles (dubbed 'leaning against the wind' or 'LATW').

In Chapter 2 I construct a model with mortgagors and entrepreneurs with a banking intermediary extending loans to the former while taking deposits from the latter. I then use this representative agent model to evaluate several types of monetary policy rules. I do not find compelling evidence for house price targeting and I find that targeting house prices and inflation together may result in model indeterminacy, especially when debt is not indexed and payable in nominal terms.

With this modelling exercise, I aim to contribute to the academic literature on credit frictions with a New-Keynesian model that explicitly factors in the supply and demand of credit and a net-worth channel for monetary policy. In terms of transmission mechanisms, I show that asset price targeting is similar to output targeting, given the fact that house prices are central for the aggregate demand.

My results could understate the effects of LATW policy as mortgage loans are modelled as a 1-period adjustable rate bonds. This is obviously different from longer term mortgages, which in the real world could go up to 30 years. So similar models could try to assess the transmission mechanism in presence of longer term debt contracts. Longer maturity mortgages may result in a levered IS schedule with higher sensitivity to interest rates.

In the thesis I highlight how defaulting of contractual obligations is a business cycle stylised fact that is difficult to capture in a micro-founded model. Nevertheless, I use a costly-state verification friction to derive a default rate, an additional control variable that responds to other macro-variables. The response of the default threshold are qualitatively correct, but the model does not capture the volatility of the default rate observed in the business cycle. Again, the modification mentioned above could translate in a stickier debt stock that is difficult for the mortgagors to re-absorb upon a negative shock.

The central theme of the first two Chapters of this thesis is the conduct of monetary policy. I do various policy experiments using the model I developed as a sandbox to observe the performance of various monetary policy rules. Among other findings, I find that interest rate smoothing shifts the efficient frontier of the output-inflation variance trade-off.

Interestingly, there is no consensus on the optimal degree of smoothing in the United Kingdom or if the Central Bank smooths the interest rate at all. DSGE estimates of the Taylor rule often find a high degree of smoothing, but micro-econometric estimates show that highly significant interest rate smoothing might be spurious and due to the endogenous regressor problem. In the first Chapter I deal with endogeneity presenting a VAR identification that hinges on the retrieval of exogenous shock. In Chapter 2 I remain agnostic on the optimal degree of smoothing, commenting just on how that changes my results. An interesting extension of Chapter 2 could be around Taylor rules, as those enter into DSGE models in a formulaic way. Factoring in different laws for expectation formation may offer a different case for LATW or other monetary policy rules targeting financial stability. Having different laws for expectation formation may provide a more natural framework to ask a follow up question such as can monetary policy stabilise the housing market by playing a role in anchoring expectations. A model in which the Central Bank has to learn the workings of the economy could also provide a case for smoothing, which is already quantitatively important for a model with rational expectations.

In Chapter 3 we show how the US wage is very sticky. To formalise this notion of stickiness, we compare different measures for the user cost of labour to a model implied Nash wage. The actual wage is much stickier relative to the Nash wage. We call this empirical measure of elasticity 'Nash Wage Elasticity' (NWE) and we suggest that NWE can be calibrated in DSGE models to match the volatility of unemployment. Overall, we show that wage rigidity amplifies the business cycle.

This finding can be used in future DSGE model efforts as it helps matching the volatility of unemployment in the business cycle. Our NWE measure can provide a basis for comparison across other models with rigid wages and help to inform macro-modellers about the actual degree of wage rigidity in the data.

In Chapter 4 I test the hypothesis of financialization in the commodity market contributing to a wide and thriving literature concerned with causal linkages from the increase in traded volume of commodity indexes to future prices. Such one-directional causality is often referred to as 'financialization'. I apply a novel Granger-testing methodology suitable to detect changing points in the causal relation and its strength from index investment to a panel of commodities. Through this method I can confirm that the Great Financial Crisis (GFC) represented a key changing point in causal relation, but I also uncover differences among commodities that are due to idiosyncratic factors.

The empirical results provide mild evidence of financialization in oil and agricultural markets. However, my estimates do support the hypothesis of financialization, although heightened causality is clustered at specific points in time. Our inference is more robust than alternative testing methodology and thus can be taken at face value despite the outliers due to having the GFC early in the sample.

This chapter helps to shed some light on financial markets behaviour, providing some compelling evidence supporting occasional departures from the efficient market hypothesis. The fact that data exhibit a heightened correlation at certain points in time and thus a time-varying causal profile in the Granger sense, means that there is evidence that investor sentiment is a major driver for certain observed financial market patterns.

I provide a re-interpretation of Singleton (2014) on financialization of the crude oil market and appraisal of the work of Christopher Gilbert and co-authors. I can do so using a testing strategy that is robust to the behaviour of data around the Great Financial Crisis, when the structural break associated with wild price movement can bias standard test results. I argue for using a time-varying test as a robust alternative for the static Granger causality

References

Jonathan Acosta-Smith, Marzio Bassanin, and Ivy Sabuga. Dynamics of the Output Floor: A Model-Based Assessment. 2021.

Klaus Adam, Pei Kuang, and Albert Marcet. House price booms and the current account. *NBER Macroeconomics Annual*, 26(1), 2011. ISSN 08893365. doi: 10.1086/663990.

Zeno Adams and Thorsten Glück. Financialization in commodity markets: A passing trend or the new normal? *Journal of Banking and Finance*, 60, 2015. ISSN 03784266. doi: 10.1016/j.jbankfin.2015.07.008.

Malin Adolfson, Stefan Laséen, Jesper Lindé, and Lars E.O. Svensson. Optimal monetary policy in an operational medium-sized DSGE model. *Journal of Money, Credit and Banking*, 43(7), 2011. ISSN 00222879. doi: 10.1111/j.1538-4616.2011.00426.x.

David Aikman, Andrew G. Haldane, and Benjamin D. Nelson. Curbing the credit cycle. *Economic Journal*, 125(585), 2015. ISSN 14680297. doi: 10.1111/ecoj.12113.

Jonas E. Arias, Juan F. Rubio-Ramírez, and Daniel F. Waggoner. Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications. 2014. URL www.federalreserve.gov/pubs/ifdp/.www.ssrn.com.

Jonas E. Arias, Dario Caldara, and Juan F. Rubio-Ramírez. The systematic component of monetary policy in SVARs: An agnostic identification procedure. *Journal of Monetary Economics*, 2019. ISSN 03043932. doi: 10.1016/j.jmoneco.2018.07.011.

Guido Ascari and Argia M. Sbordone. The macroeconomics of trend inflation, 2014. ISSN 00220515.

Adrien Auclert, Bence Bardóczy, and Matthew Rognlie. Mpcs, mpes, and multipliers: A trilemma for new keynesian models. *The Review of Economics and Statistics*, pages 1–41, 2021.

Regis Barnichon, Davide Debortoli, and Christian Matthes. Understanding the size of the government spending multiplier: It's in the sign. *The Review of Economic Studies*, 2021.

Suleyman Basak and Anna Pavlova. A Model of Financialization of Commodities. *Journal of Finance*, 71(4), 2016. ISSN 15406261. doi: 10.1111/jofi.12408.

Susanto Basu and Christopher L House. Allocative and remitted wages: New facts and challenges for keynesian models. In *Handbook of macroeconomics*, volume 2, pages 297–354. Elsevier, 2016.

Andriana Bellou and Bariş Kaymak. The cyclical behavior of job quality and real wage growth. *American Economic Review: Insights*, 3(1):83–96, 2021.

Jonathan Benchimol and André Fourçans. Central bank losses and monetary policy rules: A DSGE investigation. *International Review of Economics and Finance*, 61, 2019. ISSN 10590560. doi: 10.1016/j.iref.2019.01.010.

Jaromir Benes and Michael Kumhof. Risky bank lending and countercyclical capital buffers. *Journal of Economic Dynamics and Control*, 58, 2015. ISSN 01651889. doi: 10.1016/j.jedc.2015.06.005.

Ben Bernanke and Mark Gertler. Agency costs, net worth, and business fluctuations. *American Economic Review*, 79(1), 1989. ISSN 00028282. doi: 10.2307/1804770.

Ben S Bernanke and Mark Gertler. Inside the Black Box: The Credit Channel of Monetary Policy Transmission. *Journal of Economic Perspectives*, 9(4):27–48, 1995. ISSN 0895-3309. doi: 10.1257/jep.9.4.27.

Ben S. Bernanke, Mark Gertler, and Simon Gilchrist. The financial accelerator in a quantitative business cycle framework. In *Handbook of Macroeconomics*, volume 1, pages 1341–1393. 1999. doi: 10.1016/S1574-0048(99)10034-X.

Mark Bils, Peter J Klenow, and Benjamin A Malin. Resurrecting the role of the product market wedge in recessions. *American Economic Review*, 108(4-5):1118–46, 2018.

Mark J Bils. Real wages over the business cycle: evidence from panel data. *Journal of Political economy*, 93(4):666–689, 1985.

Olivier Blanchard and Jordi Galí. Real wage rigidities and the new keynesian model. *Journal of money, credit and banking*, 39:35–65, 2007.

Alan S. Blinder and Ricardo Reis. Understanding the Greenspan standard. In *Proceedings - Economic Policy Symposium - Jackson Hole, Federal Reserve Bank of Kansas City*, pages 11–96, 2005.

Kristina Bluwstein, Michal Brzoza-Brzezina, Paolo Gelain, and Marcin Kolasa. Multi-Period Loans, Occasionally Binding Constraints and Monetary Policy: A Quantitative Evaluation. *SSRN Electronic Journal*, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3230122.

Claudio Borio. The financial cycle and macroeconomics: What have we learnt? *Journal of Banking and Finance*, 2014. ISSN 03784266. doi: 10.1016/j.jbankfin.2013.07.031.

Claudio E. V. Borio and Philip William Lowe. Asset Prices, Financial and Monetary Stability: Exploring the Nexus. *SSRN Electronic Journal*, 2011. doi: 10.2139/ssrn.846305.

Ben Broadbent. Debt dynamics Speech given by. *London Business School - Speech*, (January):1–20, 2019.

Tobias Broer, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg. The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies*, 87(1):77–101, 2020.

MichaŁ Brzoza-Brzezina, Marcin Kolasa, and Krzysztof Makarski. The anatomy of standard DSGE models with financial frictions. *Journal of Economic Dynamics and Control*, 37(1), 2013. ISSN 01651889. doi: 10.1016/j.jedc.2012.06.008.

Giacomo Candian and Mikhail Dmitriev. Default recovery rates and aggregate fluctuations. *Journal of Economic Dynamics and Control*, 121, 2020. ISSN 01651889. doi: 10.1016/j.jedc.2020.104011.

Cristiano Cantore, Filippo Ferroni, and Miguel León-ledesma. The missing link : monetary policy and the labor share. *Bank Of England - Staff Working Paper*, (857), 2020.

Charles T. Carlstrom and Timothy S. Fuerst. Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis. *American Economic Review*, 87(5), 1997. ISSN 00028282. doi: 10.2307/2951331.

Anabela Carneiro, Paulo Guimarães, and Pedro Portugal. Real wages and the business cycle: Accounting for worker, firm, and job title heterogeneity. *American Economic Journal: Macroeconomics*, 4(2):133–52, 2012.

Mark Carney. the Grand-Unifying Theory and Practice of Macroprudential Policy (Speech). 2020. URL https://www.bankofengland.co.uk/speech/2020/mark-carney-speech-at-university-college-london.

Ambrogio Cesa-Bianchi, Gregory Thwaites, and Alejandro Vicondoa. Monetary Policy Transmission in the United Kingdom: A High Frequency Identification Approach. *European Economic Review*, page 103375, 2020. ISSN 00142921. doi: 10.1016/j.euroecorev.2020.103375.

Varadarajan V Chari, Patrick J Kehoe, and Ellen R McGrattan. Business cycle accounting. *Econometrica*, 75(3):781–836, 2007.

Varadarajan Venkata Chari and Lawrence J. Christiano. Financialization in Commodity Market. 2019.

Xu Cheng, Xu Han, and Atsushi Inoue. Instrumental Variable Estimation of Structural VAR Models Robust to Possible Non-Stationarity. *SSRN Electronic Journal*, 2019. doi: 10.2139/ssrn.3376494.

Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1):451–488, 2009.

Gabriel Chodorow-Reich and Loukas Karabarbounis. The cyclicality of the opportunity cost of employment. *Journal of Political Economy*, 124(6):1563–1618, 2016.

Lawrence J. Christiano, Roberto Motto, and Massimo Rostagno. Risk shocks. *American Economic Review*, 2014. ISSN 00028282. doi: 10.1257/aer.104.1.27.

Lawrence J Christiano, Martin S Eichenbaum, and Mathias Trabandt. Understanding the great recession. *American Economic Journal: Macroeconomics*, 7(1):110–67, 2015.

Lawrence J Christiano, Martin S Eichenbaum, and Mathias Trabandt. Unemployment and business cycles. *Econometrica*, 84(4):1523–1569, 2016.

Lawrence J Christiano, Martin S Eichenbaum, and Mathias Trabandt. Why is unemployment so countercyclical? *Review of Economic Dynamics*, 2021.

Kai Christoffel and Tobias Linzert. The role of real wage rigidity and labor market frictions for inflation persistence. *Journal of Money, Credit and Banking*, 42(7): 1435–1446, 2010.

Richard Clarida, Jordi Galí, and Mark Gertler. Monetary policy rules in practice some international evidence. *European Economic Review*, 42(6), 1998. ISSN 00142921. doi: 10.1016/S0014-2921(98)00016-6.

Laurent Clerc, Alexis Derviz, Caterina Mendicino, Stephane Moyen, Kalin Nikolov, Livio Stracca, Javier Suarez, and Alexandros P. Vardoulakis. Capital regulation in a macroeconomic model with three layers of default. *International Journal of Central Banking*, 11(3), 2015. ISSN 18157556. doi: 10.2139/ssrn.2629093.

James Cloyne and Patrick Hürtgen. The macroeconomic effects of monetary policy: A new measure for the united kingdom. *American Economic Journal: Macroeconomics*, 8(4):75–102, 2016. ISSN 19457715. doi: 10.1257/mac.20150093.

David Cobham and Yue Kang. Time horizons and smoothing in the bank of England's reaction function: The contrast between the standard GMM and Ex ante forecast approaches. *Oxford Bulletin of Economics and Statistics*, 75(5), 2013. ISSN 03059049. doi: 10.1111/j.1468-0084.2012.00709.x.

Commodity Futures Trading Commission. Commission Actions in Response to the "Comprehensive Review of the Commitments of Traders Reporting" (June 21, 2006) Program. Technical report, 2006. URL https://www.cftc.gov/sites/default/files/idc/groups/public/@commitmentsoftraders/documents/file/noticeonsupplementalcotrept.pdf.

Gonzalo Cortazar, Hector Ortega, Maximiliano Rojas, and Eduardo S. Schwartz. Commodity index risk premium. *Journal of Commodity Markets*, 22, 2021. ISSN 24058513. doi: 10.1016/j.jcomm.2020.100156.

Vasco Curdia, Andrea Ferrero, Ging Cee Ng, and Andrea Tambalotti. Evaluating interest rate rules in an estimated DSGE Model. 2011.

Spencer Dale and James Talbot. Forward Guidance in the UK, 2013. URL https://voxeu.org/article/forward-guidance-uk.

Markus Demary. The interplay between output, inflation, interest rates and house prices: International evidence. *Journal of Property Research*, 27(1), 2010. ISSN 09599916. doi: 10.1080/09599916.2010.499015.

John T Dunlop. The movement of real and money wage rates. *The Economic Journal*, 48(191):413–434, 1938.

Stéphane Dupraz, Emi Nakamura, and Jón Steinsson. A plucking model of business cycles. Technical report, National Bureau of Economic Research, 2019.

Colin Ellis, Haroon Mumtaz, and Pawel Zabczyk. What lies beneath? A time-varying FAVAR model for the UK transmission mechanism. *Economic Journal*, 124 (576):668–699, 2014. ISSN 14680297. doi: 10.1111/ecoj.12147.

Eugene F. Fama. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2), 1970. ISSN 00221082. doi: 10.2307/2325486.

Viviana Fanelli. Commodity-Linked Arbitrage Strategies and Portfolio Management. In *Handbook of Multi-Commodity Markets and Products*. 2015. doi: 10.1002/9781119011590.ch20.

Roger E. A. Farmer. The Indeterminacy School in Macroeconomics. In *Oxford Research Encyclopedia of Economics and Finance*. 2020. doi: 10.1093/acrefore/9780190625979.013.511.

Bassam Fattouh, Lutz Kilian, and Lavan Mahadeva. The role of speculation in oil markets: What have we learned so far?, 2013. ISSN 01956574.

Andres Fernandez, Stephanie Schmitt-Grohé, and Martín Uribe. Does the Commodity Super Cycle Matter? *Nber Working Paper Series*, 53(9), 2020. ISSN 1098-6596.

Isabel Figuerola-Ferretti and J. Roderick McCrorie. The shine of precious metals around the global financial crisis. *Journal of Empirical Finance*, 38, 2016. ISSN 09275398. doi: 10.1016/j.jempfin.2016.02.013.

Isabel Figuerola-Ferretti, Christopher L. Gilbert, and J. Roderick Mccrorie. Testing for Mild Explosivity and Bubbles in LME Non-Ferrous Metals Prices. *Journal of Time Series Analysis*, 36(5), 2015. ISSN 14679892. doi: 10.1111/jtsa.12121.

Isabel Figuerola-Ferretti, J. Roderick McCrorie, and Ioannis Paraskevopoulos. Mild explosivity in recent crude oil prices. *Energy Economics*, 87, 2020. ISSN 01409883. doi: 10.1016/j.eneco.2019.05.002.

Daria Finocchiaro and Virginia Queijo Von Heideken. Do Central Banks React to House Prices? *Journal of Money, Credit and Banking*, 45(8), 2013. ISSN 00222879. doi: 10.1111/jmcb.12065.

Chiara Forlati and Luisa Lambertini. Risky mortgages in a DSGE model. *International Journal of Central Banking*, 7(1):285–335, 2011. ISSN 18154654. doi: 10.2139/ssrn.1714303.

Milton Friedman and Anna Jacobson Schwartz. *A Monetary History of the United States 1867 - 1960*. Princeton, 1963.

Jordi Galí. *Monetary policy, inflation, and the business cycle: An introduction to the new Keynesian framework*. 2009. ISBN 9780691133164. doi: 10.1111/j.1475-4932.2009.00606.x.

Leonardo Gambacorta and Federico M. Signoretti. Should monetary policy lean against the wind?. An analysis based on a DSGE model with banking. *Journal of Economic Dynamics and Control*, 2014. ISSN 01651889. doi: 10.1016/j.jedc.2014.01.016.

Paolo Gelain, Kevin J. Lansing, and Caterina Mendicino. House prices, credit growth, and excess volatility: Implications for monetary and macroprudential policy. *International Journal of Central Banking*, 9(2), 2013. ISSN 18154654. doi: 10.24148/wp2012-11.

Paolo Gelain, Kevin J. Lansing, and Gisle James Natvik. Leaning against the credit cycle. *Journal of the European Economic Association*, 16(5), 2018. ISSN 15424774. doi: 10.1093/jeea/jvx043.

Andrea Gerali, Stefano Neri, Luca Sessa, and Federico M. Signoretti. Credit and banking in a DSGE model of the Euro area. *Journal of Money, Credit and Banking*, 2010. ISSN 00222879. doi: 10.1111/j.1538-4616.2010.00331.x.

Elena Gerko and Hélène Rey. Monetary policy in the capitals of capital. *Journal of the European Economic Association*, 15(4):721–745, 2017. ISSN 15424774. doi: 10.1093/jeea/jvx022.

Mark Gertler and Peter Karadi. Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 2015. ISSN 19457715. doi: 10.1257/mac.20130329.

Mark Gertler and Antonella Trigari. Unemployment fluctuations with staggered nash wage bargaining. *Journal of political Economy*, 117(1):38–86, 2009.

Mark Gertler, Christopher Huckfeldt, and Antonella Trigari. Unemployment fluctuations, match quality, and the wage cyclicality of new hires. *The Review of Economic Studies*, 87(4):1876–1914, 2020.

Christopher L. Gilbert. Commodity Speculation and Commodity Investment. *Journal of Commodity Markets and Risk Management*, (20), 2010a.

Christopher L. Gilbert. Speculative Influences on Commodity Futures Prices 2006-2008. Technical Report March, UNCTAD, 2010b.

Christopher L. Gilbert. Investor sentiment and market fundamentals: the impact of index investment on energy and metals markets. *Mineral Economics*, 31(1-2), 2018. ISSN 21912211. doi: 10.1007/s13563-017-0135-6.

Christopher L. Gilbert and Simone Pfuderer. The financialization of food commodity markets. In *Handbook on Food: Demand, Supply, Sustainability and Security*. 2014a. doi: 10.4337/9781781004296.00012.

Christopher L. Gilbert and Simone Pfuderer. The Role of Index Trading in Price Formation in the Grains and Oilseeds Markets. *Journal of Agricultural Economics*, 65(2), 2014b. ISSN 14779552. doi: 10.1111/1477-9552.12068.

Sílvia Gonçalves and Lutz Kilian. Bootstrapping autoregressions with conditional heteroskedasticity of unknown form. *Journal of Econometrics*, 123(1), 2004. ISSN 03044076. doi: 10.1016/j.jeconom.2003.10.030.

Marvin Goodfriend and Robert G. King. The New Neoclassical Synthesis and the Role of Monetary Policy. *NBER Macroeconomics Annual*, 12, 1997. ISSN 08893365. doi: 10.1086/654336.

Charles A. E. Goodhart and Dimitrios P. Tsomocos. The Mayekawa Lecture : The Role of Default in Macroeconomics. *Monetary and Economic Studies*, 29(November 2011):49–72, 2011. URL http://www.imes.boj.or.jp/research/abstracts/english/me29-4.html.

Charles A.E. Goodhart. The Monetary Policy Committee's Reaction Function: An Exercise in Estimation. *Topics in Macroeconomics*, 5(1), 2005. doi: 10.2202/1534-5998.1240.

Eleonora Granziera and Sharon Kozicki. House price dynamics: Fundamentals and expectations. *Journal of Economic Dynamics and Control*, 60, 2015. ISSN 01651889. doi: 10.1016/j.jedc.2015.09.003.

John Grigsby, Erik Hurst, and Ahu Yildirmaz. Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review*, 111 (2):428–71, 2021.

Luca Guerrieri and Matteo Iacoviello. Collateral constraints and macroeconomic asymmetries. *Journal of Monetary Economics*, 90, 2017. ISSN 03043932. doi: 10.1016/j.jmoneco.2017.06.004.

Christian Haefke, Marcus Sonntag, and Thijs Van Rens. Wage rigidity and job creation. *Journal of monetary economics*, 60(8):887–899, 2013.

Marcus Hagedorn and Iourii Manovskii. The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review*, 98(4):1692–1706, 2008.

Marcus Hagedorn and Iourii Manovskii. Job selection and wages over the business cycle. *American Economic Review*, 103(2):771–803, 2013.

Robert E Hall. Employment fluctuations with equilibrium wage stickiness. *American economic review*, 95(1):50–65, 2005.

Robert E Hall and Paul R Milgrom. The limited influence of unemployment on the wage bargain. *American economic review*, 98(4):1653–74, 2008.

James D. Hamilton and Jing Cynthia Wu. Risk premia in crude oil futures prices. *Journal of International Money and Finance*, 42, 2014. ISSN 02615606. doi: 10.1016/j.jimonfin.2013.08.003.

James D. Hamilton and Jing Cynthia Wu. Effects of index-fund investing on commodity futures prices. *International Economic Review*, 56(1), 2015. ISSN 14682354. doi: 10.1111/iere.12099.

Jonathon Hazell and Bledi Taska. Downward rigidity in the wage for new hires. *Available at SSRN 3728939*, 2020.

Boris Hofmann and Gert Peersman. Is there a debt service channel of monetary transmission? *BIS Quarterly Review*, (December):Bank for International Settlements, December, 2017. URL https://ideas.repec.org/a/bis/bisqtr/1712e.html.

Matteo Iacoviello. House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review*, 95(3):739–764, 2005. ISSN 00028282. doi: 10.1257/0002828054201477.

Matteo Iacoviello and Stefano Neri. Housing market spillovers: Evidence from an estimated DSGE model. *American Economic Journal: Macroeconomics*, 2(2), 2010. ISSN 19457715. doi: 10.1257/mac.2.2.125.

Scott H. Irwin. Commodity index investment and food prices: Does the "masters hypothesis" explain recent price spikes? *Agricultural Economics (United Kingdom)*, 44(SUPPL1), 2013. ISSN 01695150. doi: 10.1111/agec.12048.

Scott H. Irwin and Dwight R. Sanders. Index funds, financialization, and commodity futures markets. *Applied Economic Perspectives and Policy*, 33(1), 2011. ISSN 20405804. doi: 10.1093/aepp/ppq032.

Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises-The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2), 2020. ISSN 19457715. doi: 10.1257/mac.20180090.

Wenjin Kang, K. Geert Rouwenhorst, and Ke Tang. A Tale of Two Premiums: The Role of Hedgers and Speculators in Commodity Futures Markets. *Journal of Finance*, 75(1), 2020. ISSN 15406261. doi: 10.1111/jofi.12845.

George Kapetanios, Riccardo M. Masolo, Katerina Petrova, and Matthew Waldron. A time-varying parameter structural model of the UK economy. *Journal of Economic Dynamics and Control*, 106, 2019. ISSN 01651889. doi: 10.1016/j.jedc.2019.05.012.

Greg Kaplan, Giovanni L. Violante, and Justin Weidner. The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, Spring 201, 2014. ISSN 15334465. doi: 10.1353/eca.2014.0002.

Mervyn King. Changes in UK monetary policy: Rules and discretion in practice. *Journal of Monetary Economics*, 39(1):81–97, 1997. ISSN 03043932. doi: 10.1016/S0304-3932(97)00009-3.

Nobuhiro Kiyotaki and John Moore. Credit cycles. *Journal of Political Economy*, 1997. ISSN 00223808. doi: 10.1086/262072.

Felix Koenig, Alan Manning, and Barbara Petrongolo. Reservation wages and the wage flexibility puzzle. 2021.

Michael U Krause and Thomas A Lubik. The (ir) relevance of real wage rigidity in the new keynesian model with search frictions. *Journal of Monetary Economics*, 54 (3):706–727, 2007.

Marianna Kudlyak. The cyclicality of the user cost of labor. *Journal of Monetary Economics*, 68:53–67, 2014.

Luisa Lambertini, Caterina Mendicino, and Maria Teresa Punzi. Expectations-driven cycles in the housing market. *Economic Modelling*, 60, 2017a. ISSN 02649993. doi: 10.1016/j.econmod.2016.10.004.

Luisa Lambertini, Victoria Nuguer, and Pinar Uysal. Mortgage default in an estimated model of the U.S. housing market. *Journal of Economic Dynamics and Control*, 76, 2017b. ISSN 01651889. doi: 10.1016/j.jedc.2017.01.007.

Stefan Laseen and Ingvar Strid. Debt Dynamics and Monetary Policy: A Note. *SSRN Electronic Journal*, 2018. doi: 10.2139/ssrn.2456215.

Lien Laureys, Roland Meeks, and Boromeus Wanengkirtyo. Optimal simple objectives for monetary policy when banks matter. *European Economic Review*, 135, 2021. ISSN 00142921. doi: 10.1016/j.euroecorev.2021.103719.

Zheng Liu, Pengfei Wang, and Tao Zha. Land-Price Dynamics and Macroeconomic Fluctuations. *Econometrica*, 81(3), 2013. ISSN 0012-9682. doi: 10.3982/ecta8994.

Lars Ljungqvist and Thomas J Sargent. The fundamental surplus. *American Economic Review*, 107(9):2630–65, 2017.

Helmut Lütkepohl. *New introduction to multiple time series analysis*. 2005. ISBN 3540401725. doi: 10.1007/978-3-540-27752-1.

James M Malcomson and Sophocles Mavroeidis. Bargaining and wage rigidity in a matching model for the us. *Oxford Bulletin of Economics and Statistics*, 79(6): 997–1017, 2017.

Benoit Mandelbrot. The Variation of Certain Speculative Prices. *The Journal of Business*, 36(4), 1963. ISSN 0021-9398. doi: 10.1086/294632.

Benoit Mandelbrot. The Variation of Some Other Speculative Prices. *The Journal of Business*, 40(4), 1967. ISSN 0021-9398. doi: 10.1086/295006.

Benoit Mandelbrot. Correction of an Error in "The Variation of Certain Speculative Prices" (1963). *The Journal of Business*, 45(4), 1972. ISSN 0021-9398. doi: 10.1086/295487.

Pedro S Martins, Gary Solon, and Jonathan P Thomas. Measuring what employers do about entry wages over the business cycle: a new approach. *American Economic Journal: Macroeconomics*, 4(4):36–55, 2012.

Micheal W. Masters. Testimony before the committee on homeland security and governmental affairs. Technical report, 2008.

Micheal W. Masters. Testimony before the commodity futures trading commission. Technical report, 2009.

Christian Merkl and Heiko Stüber. Wage cyclicalities and labor market dynamics at the establishment level: Theory and evidence. Technical report, Institute of Labor Economics (IZA), 2017.

Karel Mertens and Morten O. Ravn. The dynamic effects of personal and corporate income tax changes in the United States. *American Economic Review*, 103(4):1212–47, 2013. ISSN 00028282. doi: 10.1257/aer.103.4.1212.

Pascal Michaillat and Emmanuel Saez. Beveridgean unemployment gap. *Journal of Public Economics Plus*, 2:100009, 2021.

Silvia Miranda-Agrippino and Giovanni Ricco. Bayesian Vector Autoregressions. *Bank Of England - Staff Working Paper*, Staff Work(756), 2018. URL www.bankofengland.co.uk/working-paper/staff-working-papers.

Tommaso Monacelli. New Keynesian models, durable goods, and collateral constraints. *Journal of Monetary Economics*, 56(2), 2009. ISSN 03043932. doi: 10.1016/j.jmoneco.2008.09.013.

Andrew Mountford. Leaning into the wind: A structural VAR investigation of UK monetary policy. *Oxford Bulletin of Economics and Statistics*, 67(5):597–621, 2005. ISSN 03059049. doi: 10.1111/j.1468-0084.2005.00133.x.

Toshihiko Mukoyama. The cyclicality of job-to-job transitions and its implications for aggregate productivity. *Journal of Economic Dynamics and Control*, 39:1–17, 2014.

Jamie Murray. A Small Model of the UK economy. 2012.

Marco Del Negro, Giorgio E. Primiceri, Michele Lenza, and Andrea Tambalotti. What's up with the Phillips curve? *Brookings Papers on Economic Activity*, 2020 (Spring), 2020. ISSN 15334465. doi: 10.1353/ECA.2020.0004.

Christopher J. Nekarda and Valerie A. Ramey. The cyclical behavior of the price-cost markup. *Journal of Money, Credit, and Banking: 50th Anniversary Issue*, 2021.

Serena Ng and Pierre Perron. A note on the selection of time series models, 2005. ISSN 03059049.

Barbara Petrongolo and Christopher A Pissarides. Looking into the black box: A survey of the matching function. *Journal of Economic literature*, 39(2):390–431, 2001.

Nicolas Petrosky-Nadeau and Etienne Wasmer. The cyclical volatility of labor markets under frictional financial markets. *American Economic Journal: Macroeconomics*, 5(1):193–221, 2013.

Nicolas Petrosky-Nadeau and Lu Zhang. Unemployment crises. *Journal of Monetary Economics*, 2020.

Peter C. B. Phillips and Jun Yu. Dating the timeline of financial bubbles during the subprime crisis. *Quantitative Economics*, 2(3), 2011. ISSN 17597323. doi: 10.3982/qe82.

Peter C.B. Phillips, Yangru Wu, and Jun Yu. Explosive Behavior In The 1990S Nasdaq: When Did Exuberance Escalate Asset Values? *International Economic Review*, 52(1), 2011. ISSN 00206598. doi: 10.1111/j.1468-2354.2010.00625.x.

Peter C.B. Phillips, Shuping Shi, and Jun Yu. Testing for multiple bubbles: Historical episodes of exuberance and collapse in the SP 500. *International Economic Review*, 56(4), 2015a. ISSN 14682354. doi: 10.1111/iere.12132.

Peter C.B. Phillips, Shuping Shi, and Jun Yu. Testing for multiple bubbles: Limit theory of real-time detectors. *International Economic Review*, 56(4), 2015b. ISSN 14682354. doi: 10.1111/iere.12131.

Michele Piffer. Monetary policy and defaults in the united states. *International Journal of Central Banking*, 14(4):327–358, 2018. ISSN 18157556.

Christopher A Pissarides. The unemployment volatility puzzle: Is wage stickiness the answer? *Econometrica*, 77(5):1339–1369, 2009.

Sarah Zubairy Ramey, Valerie A. Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, 126(2), 2018.

VA Ramey. Macroeconomic shocks and their propagation. *Handbook of Macroeconomics*, page 71, 2016a.

Valerie A. Ramey. Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, pages 71 – 162. 2016b.

Federico Ravenna and Carl E Walsh. Vacancies, unemployment, and the phillips curve. *European Economic Review*, 52(8):1494–1521, 2008.

Carmen M. Reinhart, Vincent Reinhart, and Christoph Trebesch. Global cycles: Capital flows, commodities, and sovereign defaults, 1815-2015. In *American Economic Review*, volume 106, 2016. doi: 10.1257/aer.p20161014.

Rafael Repullo and Javier Suarez. Entrepreneurial moral hazard and bank monitoring: A model of the credit channel. *European Economic Review*, 44(10):1931–1950, 2000. ISSN 00142921. doi: 10.1016/S0014-2921(99)00069-0.

Riksbank. The effects of monetary policy on household debt. In *Monetary Policy Report*, number February 2014, pages 45–49. 2014.

Ørjan Robstad. House prices, credit and the effect of monetary policy in Norway: evidence from structural VAR models. *Empirical Economics*, 54(2):461–483, 2018. ISSN 03777332. doi: 10.1007/s00181-016-1222-1.

Christina D. Romer and David H. Romer. Choosing the federal reserve chair: Lessons from history, 2004a. ISSN 08953309.

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

Leena Rudanko. Labor market dynamics under long-term wage contracting. *Journal of monetary Economics*, 56(2):170–183, 2009.

Marek Rusnak, Tomas Havranek, and Roman Horvath. How to Solve the Price Puzzle? A Meta-Analysis. *Journal of Money, Credit and Banking*, 2013. ISSN 00222879. doi: 10.1111/j.1538-4616.2012.00561.x.

Dwight R. Sanders and Scott H. Irwin. New Evidence on the Impact of Index Funds in U.S. Grain Futures Markets. *Canadian Journal of Agricultural Economics*, 59(4), 2011. ISSN 00083976. doi: 10.1111/j.1744-7976.2011.01226.x.

Thomas j. Sargent. *Macroeconomic Theory*. Academic Press, New York, 1979.

Daniel Schaefer and Carl Singleton. The extent of downward nominal wage rigidity: New evidence from payroll data. Technical report, Department of Economics, University of Reading, 2021.

Benjamin Schoefer. The financial channel of wage rigidity. Technical report, National Bureau of Economic Research, 2021.

Farrukh Shahzad, Elie Bouri, Khaled Mokni, and Ahdi Noomen Ajmi. Energy, agriculture, and precious metals: Evidence from time-varying Granger causal relationships for both return and volatility. *Resources Policy*, 74, 2021. ISSN 03014207. doi: 10.1016/j.resourpol.2021.102298.

Shuping Shi, Peter C.B. Phillips, and Stan Hurn. Change Detection and the Causal Impact of the Yield Curve. *Journal of Time Series Analysis*, 39(6):966–987, 2018. ISSN 14679892. doi: 10.1111/jtsa.12427.

Shuping Shi, Stan Hurn, and Peter C B Phillips. Causal Change Detection in Possibly Integrated Systems: Revisiting the Money–Income Relationship*. *Journal of Financial Econometrics*, 18(1):158–180, 2020. ISSN 1479-8409. doi: 10.1093/jjfinec/nbz004.

Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American economic review*, 95(1):25–49, 2005.

Christopher A Sims. The Econometric Society. *The Economic Journal*, 42(166):331, 1980. ISSN 00130133. doi: 10.2307/2223855.

Kenneth J. Singleton. Investor flows and the 2008 boom/bust in oil prices. *Management Science*, 60(2), 2014. ISSN 00251909. doi: 10.1287/mnsc.2013.1756.

George Soros. Testimony before the U.S. Senate Commerce Committee Oversight Hearing on FTC Advanced Rulemaking on Oil Market Manipulation, 2008.

James H. Stock and Mark W. Watson. Vector Autoregressions. *Journal ofEconomic Perspectives*, 15(4):101–115, 2001. doi: 10.1002/9780470996249.ch33.

James H. Stock and Motohiro Yogo. Testing for weak instruments in Linear Iv regression. In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*. 2005. ISBN 9780511614491. doi: 10.1017/CBO9780511614491.006.

Lars E.O. Svensson. Inflation targeting and "leaning against the wind". In *International Journal of Central Banking*, 2014.

Lars E.O. Svensson. Against the wind: Costs and benefits, effects on debt, leaning in DSGE models, and a framework for comparison of results. *International Journal of Central Banking*, 13(3):385–408, 2017. ISSN 18157556.

Ke Tang and Wei Xiong. Index investment and the financialization of commodities. *Financial Analysts Journal*, 68(6), 2012. ISSN 0015198X. doi: 10.2469/faj.v68.n6.5.

Lorie Tarshis. Changes in real and money wages. *The Economic Journal*, 49(193): 150–154, 1939.

Robert M. Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, 21(2), 1979. ISSN 10957235. doi: 10.1016/0022-0531(79)90031-0.

Harald Uhlig. A Toolkit for Analysing Nonlinear Dynamic Stochastic Models Easily. In *Computational Methods for the Study of Dynamic Economies*. 2003. doi: 10.1093/0199248273.003.0003.

Harald Uhlig. What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics*, 2005. ISSN 03043932. doi: 10.1016/j.jmoneco.2004.05.007.

Harald Uhlig. Shocks, Sign Restrictions, and Identification. In Bo Honore, Ariel Pakes, Monika Piazzesi, and Larry Samuelson, editors, *Advances in Economics and Econometrics*, pages 95–127. Cambridge University Press, Cambridge, 2017. doi: 10.1017/9781108227223.004. URL https://www.cambridge.org/core/product/identifier/CBO9781108227223A013/type/book_part.

Carlos A. Végh. *Open Economy Macroeconomics in Developing Countries*. The MIT Press, Cambridge, Massachussets, 2013. ISBN ISBN: 9780262018906.

Carl E. Walsh. *Monetary Theory and Policy, Third Edition*. 2010.

Johannes Wieland and Mu-Jeung Yang. Financial dampening. *Journal of Money Credit and Banking*, pages 79 – 113, 2020.

Michael Woodford. Fiscal Requirements for Price Stability. *Journal of Money, Credit and Banking*, 33(3), 2001. ISSN 00222879. doi: 10.2307/2673890.

Michael Woodford. Optimal interest-rate smoothing. *Review of Economic Studies*, 70 (4), 2003a. ISSN 00346527. doi: 10.1111/1467-937X.00270.

Michael Woodford. *Interest and prices: Foundations of a theory of monetary policy*. Princeton University Press, Princeton and Oxford, 2003b. ISBN 0691010498. doi: 10.1111/j.0013-0427.2005.430_2.x.

Randall Wright, Philipp Kircher, Benoît Julien, and Veronica Guerrieri. Directed search and competitive search equilibrium: A guided tour. *Journal of Economic Literature*, 59(1):90–148, 2021.