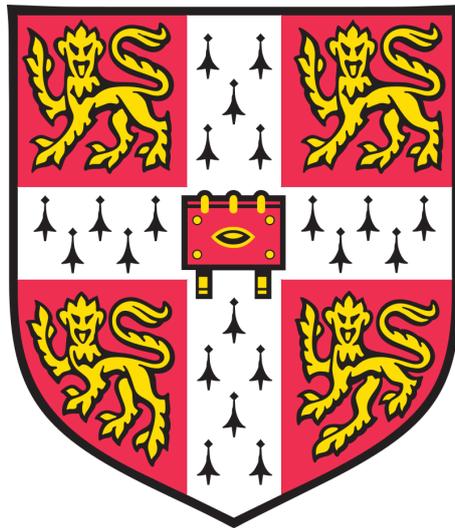


Behavioural Macro-Financial Cycles

Alfred Fitzgerald Lake

Faculty of Economics and Girton College
University of Cambridge



This thesis is submitted for the degree of Doctor of Philosophy

January 2020

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. I certify that Chapter 2 was co-authored with Dr Laurent Maurin and that I contributed over 50% of the work.

This thesis is not substantially the same as any that I have submitted, or is being concurrently submitted, for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

It does not exceed the prescribed word limit of 60,000 words from the Economics Degree Committee.

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made.

Abstract of Behavioural Macro-Financial Cycles

Alfred Fitzgerald Lake

In the first chapter I study the optimal inflation expectations that it is possible for agents to estimate and the difference between these and rational expectations. It is typically not feasible for an agent with a macroeconomic sample of realistic length to estimate a conditionally unbiased predictor of future variables, such as rational expectations. It is also often not optimal, in terms of forecast error, to minimise the conditional biases imposed, as using statistically simple expectations will often reduce forecast variance sufficiently to outweigh forecast bias. I therefore introduce optimal feasible expectations, the expectations that are predicted to minimise the relevant measure of forecast error out of the set of expectations that agents can estimate, as a realistic alternative to infeasible rational expectations. I then empirically estimate the optimal conditional biases when forecasting US inflation using a factor weighted ridge approach. I find it is optimal to impose large conditional biases: one should essentially only use information on past changes in price indices, despite several other variables and factors having economically and statistically significant associations with future inflation. I then compare these to the conditional biases in US household forecast surveys. I find that many of the conditional biases are similar, although households also make errors that reduce their forecast performance compared to feasible empirical alternatives. Therefore a combination of optimal feasible expectations and behavioural errors appear to explain US household inflation forecasts.

In the second chapter I study whether increases in asset price convergence and the quantity of cross-border asset holdings, common measures of financial integration, imply high quality changes in financial integration, i.e. changes that are likely to produce the largest net economic benefits. I use a new methodology based on a Bayesian FAVAR to overcome the econometrically challenging setting and test three aspects of the quality of changes in financial integration measures. I apply this methodology to the EU in the 21st century and find that there is a common factor that drives a wide range of price and quantity integration measures. However the changes in financial integration are primarily cyclical, as long-term cyclical strongly outweighs deterministic and stochastic trends, and dependent on macroeconomic conditions, as virtually all sign identified economic shocks cause large corresponding effects on financial integration. This suggests that increases in financial integration have not been high quality: they actually appear most closely related to cyclical changes in the underlying risks of European assets and aversion to these risks.

In the third chapter I introduce a new test of whether house prices are always equal to their fundamental value, adjusted to account for contractual rigidities and search frictions, based on the speed of their reaction to monetary shocks. I justify this test with two conceptual frameworks and references to existing empirical work on the transmission mechanisms of monetary policy. I then apply this test to house prices in the US using narrative monetary shocks in a local projections approach. I find that real house prices do not react to monetary shocks when contractual rigidities stop binding, however they have economically and statistically significant reactions at horizons over a year. This result is inconsistent with house prices always being equal to their fundamental value, but is consistent with agents either not fully observing monetary shocks or not incorporating these shocks into their expectations rationally. I also use a sign decomposition based on the conceptual frameworks to identify the relative importance of proximal drivers of house price cycles: I find that consumption demand is the most important driver but asset demand is also relatively important. Therefore housing cycles are likely to arise from the partially behavioural reactions to changes in housing demand.

Acknowledgements

Firstly, I want to thank Sean Holly for supervising me. I am grateful to him for his advice and for giving me the intellectual freedom to conduct applied economic research in my own way, despite being in a department that mainly focuses on abstract theory. I also want to thank my co-author Laurent Maurin for his advice and work, which was central to the paper which became the second chapter of this PhD, and Coen Teulings, my advisor and MPhil supervisor, for all his early advice. There are many others, both in Cambridge and outside, to whom I am also grateful to for useful economic discussions. This includes members of the Cambridge faculty, my fellow Cambridge PhD students, my colleagues in Luxembourg and those who taught me economics before my PhD.

I must also express my gratitude to the Cambridge Trust and Girton College, for fully funding me during my time in Cambridge, and to the European Investment Bank for offering me a paid research placement with them in Luxembourg.

Producing this PhD has been challenging at times and the people around me have been hugely supportive. I am lucky to have had friends in the department and the university squash team in Cambridge, as well as in my department in Luxembourg, that made my days more enjoyable. I am also grateful for the great friends I made in Wolfson Court and for my oldest friends back in North London: the three years would not have been the same without them. I want to thank El so much for her love, support and all night trips to Luxembourg: they mean so much to me. Finally I would like to thank my family, in particular my parents. They have been an invaluable source of support and advice all my life and I would not have got to this point without them.

Contents

Thesis introduction	1
1 Optimal feasible expectations in our uncertain economy	13
1.1 Introduction	13
1.2 Conceptual and econometric setup	21
1.3 Macroeconomic data and factors	30
1.4 Conditional biases in optimal feasible inflation expectations	37
1.5 Conditional biases in household inflation forecasts	44
1.6 Conclusion	47
2 Asset price convergence, international asset holdings and the quality of financial integration	51
2.1 Introduction	51
2.2 Econometric methodology	57
2.3 Dataset and estimated financial integration factor	62
2.4 The quality of EU financial integration	66
2.5 Conclusion	72
3 Behavioural finance at home: house price cycles in the USA	74
3.1 Introduction	74
3.2 Testing whether house prices are consistent with fundamentals using monetary shocks	80
3.3 Conducting the test with narrative monetary shock data	87
3.4 New stylised facts on the drivers of housing cycles	93
3.5 Conclusion	102

A Supplement: Behavioural financial cycles as the cause of perpetual business cycles	104
B Appendix of ‘Optimal feasible expectations in our uncertain economy’	114
B.1 Variables used to construct factors	114
B.2 Training set size robustness	118
B.3 Forecast performance measure robustness	122
C Appendix of ‘Asset price convergence, international asset holdings and the quality of financial integration’	124
C.1 Gibbs sampler	124
C.2 Signed impulse response functions	126
C.3 Convergence robustness	130
C.4 Prior robustness	132
C.5 Extended dataset	136
D Appendix of ‘Behavioural finance at home: house price cycles in the USA’	138
D.1 Conceptual frameworks	138
D.2 Cycle periodicity robustness	147
D.3 Housing data source robustness	153
Bibliography	156

Thesis introduction

This thesis studies behavioural macro-financial cycles. Financial cycles and their links to the macroeconomy are an important topic, as information on financial cycles can help to detect financial crises in real time (Borio, 2014) and the downturns of financial cycles are associated with serious recessions (Claessens et al., 2012). A behavioural approach to this topic is needed, as it is not appropriate to believe that all agents make financial decisions by maximising utility functions with rational expectations when most individuals cannot correctly answer basic questions on economic and financial literacy (Lusardi and Mitchell, 2014). This thesis is primarily made up of three self-contained chapters, each of which analyses topics which are related to macro-financial cycles and either explicitly analyses behavioural insights or suggests them as a potential underlying cause of the results. The three main chapters are followed by four appendices and a bibliography.

The first chapter is titled ‘Optimal feasible expectations in our uncertain economy’. In this chapter I analyse the optimal macroeconomic expectations that agents could feasibly estimate and any differences between these optimal feasible expectations and their actual expectations. It particularly focuses on inflation expectations due to their importance in macroeconomics.

There is a large macroeconomic learning literature, surveyed in Evans and Honkapohja (2012), that models how agents estimate inflation expectations. Papers in this survey often implicitly assume that agents have access to extremely large quantities of relevant data. In this case using rational expectations given the data available to agents is feasible and, since rational expectations are the true conditional expectation (Sheffrin, 1996), optimal in terms of minimising measures of forecast error. Some papers in this literature acknowledge that in reality structural change over

time means that limited samples are available (Orphanides and Williams, 2007), however they usually use models with so few variables that it is still possible to estimate conditionally unbiased estimators. The first contribution of this chapter is to discuss clearly why estimating conditionally unbiased expectations, such as rational expectations, will usually be impossible in reality. It is usually statistically impossible as there are a far greater number of potentially relevant macroeconomic variables than there are relevant time series observations of each of these variables. This seems very unlikely to change in the foreseeable future, as the rise of big data makes even more variables available and the economy undergoes structural change as a result of the response to Covid-19. It also clearly discusses why it may often not be optimal to try and limit the degree of conditional biases imposed. This is because it may often be worth shrinking expectations towards those that are statistically simpler to reduce the conditional variance of forecasts even if it imposes greater conditional biases in forecasts. This insight is at the heart of many modern machine learning approaches to forecasting, such as those in Medeiros et al. (2020). It therefore introduces optimal feasible expectations, namely the expectations that are predicted to minimise the relevant measure of forecast error out of the set of expectations that agents can estimate.

The econometric optimal inflation forecasting literature finds that very simple inflation forecasts perform best in forecasting horse races (Faust and Wright, 2013). Therefore the optimal forecasts suggested by this literature do not incorporate any associations between the vast majority of macroeconomic variables and future inflation. However this literature does not typically try to estimate the true associations between macroeconomic predictors and future inflation and so does not examine the conditional biases in optimal econometric inflation forecasts. The second contribution of this chapter is to estimate both the true associations between a set of macroeconomic variables and future inflation and the associations that it is optimal to use in forecasting. The difference between the two is the conditional bias that is applied to the macroeconomic variables by using optimal forecasts.

There is, however, a literature that estimates the conditional biases in surveys of agents forecasts (Coibion et al., 2018). The third contribution of this chapter

is to estimate the conditional biases in surveyed US household expectations but to also assess the extent to which these conditional biases are similar to the estimated optimal conditional biases. This offers insights into how similar household forecasts are to optimal feasible expectations. I also assess whether household forecasts match the forecasting performance of feasible empirical alternatives to complement this analysis.

The baseline specification for inflation forecasting that I start with is an extremely simple direct model which only includes lagged inflation. I then consider adding different macroeconomic variables to this specification in turn. The macroeconomic variables that I use include monthly measures of inflation cycles, business cycles and financial cycles; all of which are factors estimated with principal components from many underlying indicators. They also include information on exchange rates, wages and monetary shocks. They therefore include equivalents of all of the variables traditionally used for forecasting discussed by Stock and Watson (2008) and cover many of the main series typically used in macroeconomic models.

I estimate the true association between each variable and future inflation by adding each variable to the baseline case in turn and estimating the model with ordinary least squares (OLS). I estimate the association between each variable and future inflation that is optimal to use in forecasting as follows. I split the sample into many overlapping training sets and in each of them consider the model with each variable included in turn. However in each case I estimate the model many times using weighted ridge and applying different levels of shrinkage. The most shrunken case corresponds to estimating the baseline case with OLS and the least shrunken case corresponds to estimating the model with a specific variable included with OLS. I then take the optimal association between each variable and future inflation to use in forecasting as the one given by the level of shrinkage that minimises measures of forecast error in the test sets. The difference between the estimated optimal association for forecasting and the estimated true association for each macroeconomic variable is the estimate of the conditional bias applied to that variable by using optimal forecasts.

The results show it is optimal in terms of forecast error to apply very high levels of

shrinkage to many macroeconomic variables. Indeed, it is essentially optimal to only use information on past price changes when producing inflation forecasts, although in a non-naive manner that incorporates partial shrinkage. The results also show that several of the macroeconomic variables analysed have statistically and economically significant associations with future inflation. Together these results suggest that optimal inflation forecasts incorporate large conditional biases with respect to the information in common macroeconomic variables.

I then estimate the conditional biases in surveyed household inflation expectations. I do this by comparing the estimated associations between each macroeconomic variable and future inflation produced using the method described above and the equivalent association between each macroeconomic variable and household forecasts of future inflation. This involves adding each variable to the baseline case in turn and estimating the model with OLS but comparing the results produced when using future inflation or household forecasts of future inflation as the dependent variable. I find that most conditional biases in household forecasts are very similar to the equivalent conditional biases in optimal expectations. However some are clearly mistaken: for instance households forecasts of inflation significantly rise with a fall in the financial cycle indicator, when they should fall. Household forecasts are also clearly beaten in terms of pseudo out of sample forecasting performance by the baseline empirical model, suggesting that they are not optimal feasible expectations.

Therefore these results suggest that the optimal feasible inflation expectations, which are statistically simple, contain large conditional biases with respect to a set of major macroeconomic variables. They therefore appear to be very different to rational expectations given available macroeconomic data. Household inflation forecasts appear to be well explained by a combination of optimal feasible expectations and behavioural errors that decrease forecast performance.

These empirical results and the underlying discussions suggest that important conditional biases are likely to be present in the optimal feasible expectations of many macroeconomic variables. This implies that agents learning optimally from macroeconomic data are generally likely to estimate expectations that are very different from rational expectations. This is a very important result, as it undermines

one of the major arguments used to try and justify the rational expectations revolution. This implies that we should generally conceive of macroeconomic expectations as optimal feasible inflation expectations with the addition of behavioural errors in settings in which agents don't act optimally. This suggests that the many current models based on rational expectations or very minor deviations from them are likely to be seriously mis-specified. It also suggests promising areas for future research, such as examining whether optimal feasible expectations vary between individuals and applying this new type of expectations to different areas of macroeconomics. For instance, the supplement to this thesis contained in the first appendix suggests how optimal feasible expectations could contribute to cycles in asset prices.

The second chapter is titled 'Asset price convergence, international asset holdings and the quality of financial integration'. In it I assess the quality of recent changes in financial integration between EU countries.

This chapter builds on the existing empirical financial integration literature. Many papers in this literature produce measures of financial integration in Europe based on the unconditional convergence of asset prices and/or the proportion of international assets held in portfolios: see Hoffmann et al. (2019) for an overview. However increases in these measures do not just capture increases in the policy definition of financial integration, i.e. reductions in international financial frictions that cause differences in access to or investment of capital on the basis of location (Coeure, 2013), they also capture changes in the different underlying risks factors of assets and aversion to them. Increases in policy financial integration are usually suggested to have economic benefits by improving the efficiency of capital allocation and risk sharing (Baele et al., 2004). However, since empirical measures of financial integration do not just capture changes in policy financial integration, increases in them they may not indicate capital allocation and risk sharing benefits. Additionally, increases in measured financial integration, whether driven by policy financial integration or underlying risk components, might impose economic costs by increasing the risks of financial contagion and instability (Stiglitz, 2010).

As a result most policy papers that update financial integration measures have also begun to discuss the quality of changes in financial integration measures, i.e.

the extent to which financial integration is likely to have net economic benefits (European Central Bank, 2016; European Investment Bank, 2017). The chapter contributes to the existing literature by providing a methodology for statistically analysing three aspects of the quality of changes in financial integration measures that are discussed in the existing literature. These are: whether there are jointly driven changes in price and quantity measures of integration (Coeure, 2013), whether the changes in financial integration are permanent (European Investment Bank, 2019) and whether the changes in financial integration are robust to shocks that affect macroeconomic conditions (European Central Bank, 2018). It also contributes by applying this methodology to analyse the three aspects of the quality of financial integration changes in the EU since 2000 and so provides indirect evidence of the extent to which this integration is likely to have produced economic benefits.

To analyse whether price and quantity integration measures have a joint driver I need to construct a financial integration factor from both price and quantity indicators of integration. This is econometrically challenging, as I also need to allow for permanent changes in the factor that may cause non-stationarity and insert the factor into a system with macroeconomic and financial variables to analyse its response to shocks to macroeconomic conditions. Therefore I cannot use typical principal components methods, which require stationary variables. Maximum likelihood methods, such as the Kalman filter, could overcome this issue but are computationally challenging given the number of parameters used, so I use a Bayesian approach. Specifically I jointly estimate a Bayesian financial integration factor from many price and quantity integration indicators and the dynamics of a vector autoregressive system that includes macroeconomic and financial variables as well as the factor. This estimation is achieved using Markov chain Monte Carlo methods: specifically a Gibbs sampler with a Carter-Kohn step to generate the factor. Applying this method to data for the EU yields a measure of financial integration that has large increases in the early 2000s but then plateaus and has primarily decreased since 2008, although has recovered a little in the last few years.

I then test the sign of the factor loadings on price and quantity integration indicators to assess whether there is a joint driver, test the size of the deterministic and

stochastic trend changes to assess the permanence of any increases in the indicator and use sign restrictions to test the extent to which shocks that affect macroeconomic conditions also affect integration. I find that the financial integration factor is a strong joint driver of a wide range of price and quantity measures of integration. The factor loads particularly positively on virtually all integration measures in the bank lending, corporate bond and government bond markets, but loads far less strongly on indicators in equity markets. However the changes in the financial integration factor are virtually all cyclical, as long-term cyclicity strongly outweighs deterministic and stochastic trends, and vulnerable to shocks to macroeconomic conditions, as virtually all sign identified economic shocks cause large corresponding effects on financial integration.

Therefore it appears that changes in financial integration are primarily cyclical and vulnerable to macroeconomic shocks. As a result, most changes in financial integration in the EU since 2000 appear unlikely to have provided large net economic benefits: they actually appear most closely related to cyclical changes in the underlying risks of European assets and aversion to these risks. These results only covers the relatively short period of time since the millennium, so it would be interesting to update the results in the future when the effects of major recent events such as Brexit and Covid-19 can be studied.

The third chapter is titled ‘Behavioural finance at home: house price cycles in the USA’. In it I primarily assess whether US aggregate house prices are equal to their fundamental value, allowing for the effect of search frictions and contractual rigidities in housing markets. I also analyse the proximal drivers of housing market cycles.

This chapter partly builds on the existing literature testing whether house prices are equal to their fundamental value i.e. the rationally expected discounted sum of their rents given the state of the macroeconomy (Glaeser and Nathanson, 2015), and hence whether they are efficient in incorporating information. The patterns of strong correlations and persistence in house price changes and excess housing returns, dating to at least Case and Shiller (1989) with updates surveyed in Ghysels et al. (2013), suggests that housing markets are not efficient, even given time-varying

risk aversion. Therefore house prices do not appear to always be given by their fundamental value. The literature also suggests that housing market frictions may struggle to explain the size and cyclicity of these correlations and behavioural explanations may be needed (Glaeser and Nathanson, 2015), although this is hard to prove absolutely.

This chapter contributes to this literature by introducing a test for whether house prices are consistent with always being equal to the fundamental value of housing, adjusted to account for contractual rigidities and search frictions in housing markets, based on the speed of the reaction of house prices to monetary shocks. Monetary shocks should have a clearly signed impact on the fundamental value of housing, even allowing for rigidities and frictions, as soon as contractual rigidities no longer bind. Survey data from Ellie May indicates that this should be at horizons of one to two months. I support this concept with two conceptual frameworks and references to existing empirical work on the relevant transmission channels of monetary shocks and the frictions in housing markets. Therefore reacting significantly to monetary shocks within two months is a necessary, although not sufficient, condition for changes in house prices to be entirely caused by changes in the fundamental value of housing. Whereas if house prices deviate from their fundamental values, for instance because agents either do not observe monetary shocks or do not incorporate information on monetary shocks into their expectations rationally, then house prices may react to monetary shocks far more slowly. This is because they would only react once easily observable, noticeable and understandable information on monetary shocks becomes available, possibly as the macroeconomic effects of the shock are felt. This chapter also has a secondary contribution, as it also uses a sign decomposition based on the conceptual frameworks to identify the relative importance of proximal drivers of housing market cycles.

The approach used in this test is closely linked to the empirical monetary policy literature that attempts to produce impulse response functions of variables, such as house prices, to monetary shocks. This chapter uses methods similar to those in Coibion et al. (2017) to implement the test but additionally contributes to this literature by focusing on the particular horizon of interest and by using controls,

both in the generation of the shocks and in the production of the impulse response functions, that are specific to housing markets.

Specifically I implement the test using narrative shocks in the style of Romer and Romer (2004). However I include financial controls in the generation of the shocks to avoid any remaining endogeneity from central bank reactions to the strength of policy transmission as a result of financial conditions, which might be particularly important for an asset price like real house prices. I then use these shocks in a local projections approach that also includes housing specific controls, to obtain more accurate estimates. The local projection approach also means that the timing of responses can be directly estimated and so more accurately assessed than in indirect auto-regressive models.

The results show that there is no statistically or economically significant reactions of aggregate US real house prices to monetary shocks at either one or two month horizons. However impulse response functions show that there are significant reactions with plausible signs at horizons greater than a year. In particular a one percentage point reduction in the base rate is estimated to cause an increase in real house prices of approximately three percent after two years, but virtually no effect at long horizons. Therefore these results show that aggregate real house prices in the US fail the test of consistency.

I also implement a sign decomposition based on the conceptual frameworks to identify the relative importance of proximal drivers of the cyclical component of house prices. Empirically I use band-pass filters to identify the cyclical components of house price variables and then calculate the linear and rank correlations between them. There are strong positive correlations between the cyclical components of real house prices and housing starts, but only limited associations between either of these variables and the cyclical component of real rents. Therefore the frameworks suggest that changes in consumption demand are the most important proximal driver of house price changes, changes in asset demand are also relatively important and changes in housing supply are the least important.

The slowness of the response of real house prices to monetary shocks suggests that house prices are not consistent with always being equal to the fundamental

value of housing, even after this is adjusted for search frictions and contractual rigidities. Therefore this suggests that either agents do not observe information on shocks well or that they do not use this information rationally. Both are plausible, as even experts struggle to measure macroeconomic shocks and understand their effects (Ramey, 2016). On the basis of these results investors and policymakers should conceive of housing cycles as being the partly behavioural response of housing markets to shifts in housing demand. This suggests that current macroeconomic models of housing markets based on the actions of agents using rational expectations with full information on macroeconomic shocks may be seriously mis-specified. Future work should instead focus on incorporating the behavioural nature of housing markets into macroeconomics.

The first appendix is titled ‘Behavioural financial cycles as the cause of perpetual business cycles’ and is a supplement to the three main chapters. In it I draw on the ideas from the chapters and the wider behavioural finance literature to produce a conceptual framework showing how a realistic financial cycle that is hard to predict can cause a perpetual business cycle. This behavioural financial cycle is caused by agents who use optimal feasible expectations, as introduced in Chapter 1, and have time-dependent fear-based preferences over risk, which Guiso et al. (2018) provides evidence for. The behavioural financial cycle permanently fluctuates in a stochastic way and, through its effects on aggregate demand, causes a perpetual business cycle. The other three appendices contain additional information and robustness checks for each of the three main chapters.

The three chapters are linked by a subject area and a methodological approach. The methodological approach I use in this thesis is based on the fundamental scientific and statistical principles outlined in Box (1976). Firstly, scientific models typically have to be simplified versions of reality and so should seek to be at the efficient frontier of the trade-off between simplicity and realism. Secondly, the more approximations the model makes, the more approximate its conclusions must be: a highly realistic model can offer precise conclusions, whereas a stylised simplified model should only offer approximate ones.

In this thesis I use simple conceptual frameworks that are very general, so are

realistic for their level of simplicity. However, since they are relatively simple, they can only offer approximate conclusions which, in turn, can provide a framework for empirical work. This empirical work is based on a range of methods including natural experiments, surveys, sign restrictions, forecasting analysis and informative associations. This work is conducted with the aim of providing more precise quantitative answers to the questions addressed and therefore makes up the majority of the thesis.

This approach contrasts dramatically with the dynamic stochastic general equilibrium approach that dominates macroeconomic research in most academic institutions (Stiglitz, 2018). The dynamic stochastic general equilibrium approach uses far more complex and opaque models than my conceptual frameworks: typically taking many pages of calculus to derive and requiring software to solve. However these models are still very unrealistic, as they are based on very specific behaviour that contravenes important facts: a phenomenon which leads Romer (2016) to describe them as ‘post-real’ models. The clearest, but by no means only, example of this is that these models are built on the assumption that agents perfectly maximise utility functions using rational expectations, or minor deviations from such behaviour. This is despite the fact that the majority of people cannot correctly answer three extremely basic questions on economic and financial literacy: for instance in the United States only 34.3% of surveyed individuals were able to do so (Lusardi and Mitchell, 2014). Therefore, despite their complexity, these models are sufficiently unrealistic in important ways that one cannot be confident in either their precise or their approximate conclusions.

The three chapters also all analyse topics related to behavioural macro-financial cycles. More precisely, they each examine an aspect of the cyclical behaviour in an entire financial market, or group of markets, and its interaction with macroeconomic variables. These interactions can either concern the effects of macroeconomic variables on financial cycles or the effect of financial market cyclicity on the macroeconomy. In either case there is also typically a focus on analysing the role of behavioural phenomena.

Therefore there are several common elements between the three chapters. Firstly,

they all involve the measurement of financial cycles. In Chapter 1 I combine many indicators to develop financial cycle and business cycle indices using principal components; in Chapter 2 I combine many indicators to create an EU financial integration index and analyse its cyclical component in a Bayesian setting and in Chapter 3 I use band-pass filters to obtain the cyclical components of US housing market variables. Secondly, they all concern either the effect of financial cycles on macroeconomic variables or the effect of macroeconomic variables on financial cycles. In Chapter 1 I analyse the effects of financial cycles, amongst other variables, on inflation expectations; in Chapter 2 I estimate the effect of shocks that affect macroeconomic conditions on cyclical financial integration and in Chapter 3 I estimate the effects of monetary shocks on real house prices and estimate the relative importance of the proximal economic drivers of housing cycles. Thirdly, they all focus, at least partially, on the role of behavioural phenomena, in particular expectations and risk aversion. In Chapter 1 I show that optimal feasible inflation expectations in the real world are substantially different from rational expectations and examine the rationality of household inflation forecasts. Chapter 2 does not directly analyse behavioural phenomena, but cyclical behavioural risk aversion is suggested as one of the drivers of cyclical changes in European financial integration. In Chapter 3 I show that the effects of monetary shocks on real house prices are not consistent with house prices being equal to their fundamental value, even allowing for housing market frictions, in a way that suggests that agents do not have full information rational house price expectations.

Chapter 1

Optimal feasible expectations in our uncertain economy

1.1 Introduction

Understanding inflation expectations is central to macroeconomics. Inflation expectations drive inflation itself through wage bargaining and price setting, so affect nominal rigidities and the response of the real economy to aggregate demand shocks. Therefore the responsiveness of inflation expectations to available information on macroeconomic shocks affects the answer to crucial questions such as the reaction of unemployment to financial crises or the ability of government spending to boost output.

Rational expectations have been the most common approach to modelling how inflation forecasts are formed in academic economics in recent years, although empirical and theoretical work has suggested behavioural alternatives (Coibion et al., 2018). They are defined as agents' expectations being the conditional expectation of future variables, conditioning on available information¹. Sheffrin (1996) provides a full mathematical definition of this while the original definition is available in Muth (1961). They imply that agents' expectations should react to publicly available

¹Full information rational expectations, as described by Coibion et al. (2018) and commonly used in academic macroeconomics, also require that complete knowledge of the economy is available.

information in the same manner as future realised inflation reacts (Lovell, 1986).

Rational expectations given available information will not always be the optimal feasible expectations for agents when forecasting a variable. If rational expectations are feasible then they will be the optimal feasible expectations for an agent, using mean square forecast error to define optimality (Diebold, 2017). However they will only be feasible if the agent can either deduce or perfectly estimate the relevant parameters of the conditional distribution of the variable being forecasted. Given the complexity of modern economies it is simply not possible to deduce rational expectations without estimation from data in the vast majority of circumstances. Therefore estimation from data must be used to form expectations. This has driven a large learning literature studying whether expectations based on learning from data converge to rational expectations, which is surveyed by Evans and Honkapohja (2012). Since papers in this literature are primarily interested in convergence, they often assume agents have access to infinite observations of data. With infinite data an agent can use a conditionally unbiased and consistent estimate as the conditional expectation, such as that formed by approaches similar to regressing realised values of the variable being forecast on all available past information. Such an approach would converge to the conditional expectation, i.e. rational expectations, so rational expectations are feasible with infinite data.

However with a finite series of data rational expectations are not generally feasible, as conditionally unbiased estimators, such as those produced by regressing the variable being forecast on past information, will vary slightly around the conditional expectation as a result of estimation error. Any conditionally biased estimator will also contain clear deviations from rational expectations. Therefore rational expectations will not be the optimal feasible expectations in most settings. It also may not be optimal to try to limit the conditional biases that one imposes in expectations to make estimating them feasible. This is because it may be worth shrinking estimated expectations towards statistically simple expectations to reduce forecast variance, even though this introduces greater conditional biases. This insight is at the heart of modern machine learning (Ahmed et al., 2010) and Bayesian approaches to forecasting (De Mol et al., 2008) and is also present in frequentist forecasting approaches

(Bai and Ng, 2008).

I begin this chapter by discussing why some shrinkage is likely to be needed when forming expectations for the vast majority of macroeconomic variables, as a result of the large number of potentially relevant data series available relative to the number of observations of each series². I also discuss why it is very often likely that additional shrinkage towards statistically simpler specifications will improve the bias-variance trade-off of forecasts as a result of reducing estimation error and so reduce measures of forecast error. However the precise level of shrinkage in optimal feasible expectations in a particular setting is ultimately an empirical issue, so I then analyse the empirical importance of this shrinkage when forecasting US inflation.

Specifically I consider adding a number of different potential predictors of inflation to a baseline auto-regressive forecast of inflation. I estimate these forecasts in a number of training sets using weighted ridge specifications with different levels of statistical shrinkage applied to each additional predictor and then take the optimal degree of shrinkage as the one which minimises measures of forecast errors in test sets. The results suggest that a large degree of shrinkage should be applied to most variables³; indeed the optimal forecast virtually only uses information on past inflation and components of inflation. These results are closely linked to those from the empirical inflation forecasting literature, which show that univariate inflation forecasts are hard to beat in forecasting horse-races (Stock and Watson, 2008). However, unlike this literature, I also estimate equivalent specifications without shrinkage and find that inflation does have economically and statistically significant associations with some of the predictors. This implies that the high levels of optimal shrinkage do not just come from inflation being uncorrelated with past information, they also come from it being worth conditionally biasing inflation forecasts towards statistically simpler forecasts to reduce conditional forecast error. Therefore the optimal feasible inflation expectations are very different from rational expectations.

Finally I analyse the conditional biases in surveys of actual US household in-

²A phenomenon known as fat data (Koop, 2017)

³The results only consider the linear effects of the variables, but given the number of potential non-linear effects far more shrinkage would be needed to even make estimation with a wide range of non-linear effects feasible.

flation forecasts using an approach similar to that in the existing literature. I find that there are significant biases, particularly in the response to changes in past broad inflation and to financial cycle indicators. Many of these conditional biases appear to arise from using the same conditional biases as estimated optimal feasible expectations, such as the limited response to broad changes in past inflation. However household forecasts are shown not to be the optimal feasible expectations, as they perform worse in pseudo out of sample forecast comparisons than feasible empirical alternatives⁴. Therefore both optimal feasible expectations and behavioural mistakes are likely to have a role in explaining US household inflation forecasts.

I also suggest optimal feasible expectations as a new general class of expectations, formally defined as the expectations that are predicted to minimise the relevant measure of forecast error out of the set of expectations that it is feasible for agents in the real world to estimate. Optimal feasible expectations are likely to differ materially from rational expectations in most circumstances as they are likely to incorporate conditional biases associated with being statistically simple. Indeed, given the importance of parsimony in forecasting many variables (Kim and Swanson, 2018), despite their no doubt numerous true links to one another, optimal shrinkage is likely to cause optimal feasible expectations to be dramatically different to rational expectations in a large number of macroeconomic settings. I suggest that we should generally conceive of macroeconomic expectations as optimal feasible inflation expectations with the addition of behavioural errors in settings in which agents do not act optimally.

Work on how inflation expectations are formed has a long and important history in the macroeconomic literature that includes the discussions of money illusion in Keynes (1936), the adaptive inflation expectations in Friedman (1977), the model-specific rational price expectations in Lucas (1996) and the behavioural pricing in Akerlof (2002). However the work in this chapter is most closely related to, and contributes to, three relatively distinct branches of the existing literature.

Firstly, this chapter relates to the literature studying learning and expectations

⁴This is unsurprising given the clear evidence that many people have a poor understanding of inflation (Del Giovane et al., 2008)

in macroeconomics, as surveyed in Evans and Honkapohja (2012). In this literature work tends to investigate the implications of agents learning expectations from data in theoretical macroeconomic models. As the majority of this literature tends to focus on whether such learning behaviour leads to models converging to rational expectations equilibria, it is common practice to assume that agents have access to an infinite series of relevant data (Evans and Honkapohja, 2012). However, as described above, in this case it is feasible and optimal to use a conditionally unbiased and consistent estimate of rational expectations, such as that given by approaches based on least squares, which then simply implies that agents use rational expectations⁵. When agents have finite data it is not feasible to use rational expectations, as consistent estimators will not converge to the true conditional expectations. However it may be possible to use a conditionally unbiased estimator, such as approaches based on least squares similar to that in Orphanides and Williams (2007), which implies that agents use rational expectations plus noise. However the papers that are most closely related to this chapter are those in which agents with finite data use methods that give conditionally biased expectations. For instance Hommes et al. (2019) assume agents use least squares but only applied to an auto-regressive rule while Chung and Xiao (2013) assume that agents use a vector auto-regression with a subset of relevant variables.

The justification for these learning methods is that the authors are looking for a method that balances tractability in the theoretical model considered with being a good approximation for what some forecasters do in practice. I contribute to this literature by studying the optimal expectations that are feasible for an agent to use, rather than the feasible expectations that some agents may use in practice. My empirical approach frees me to do this and has the advantage of allowing me to study shrinkage in the real world. I demonstrate that in the case of US inflation forecasting the optimal feasible expectations contain large conditional biases, conditioning on important macroeconomic series and series that are often used as predictors of

⁵Note I am discussing whether an approach implies that agents use rational expectations from an infinite sample of relevant data, not whether an approach leads to a specific model converging to a rational expectations equilibria of that model.

inflation. This is because shrinking expectations towards simpler forecasts reduces conditional forecast variance sufficiently to more than offset the conditional forecast bias imparted. I also suggest why similar shrinkage is also likely to be used in the optimal feasible expectations in the vast majority of macroeconomic settings. This is hugely important as it implies that agents learning optimally will use expectations that are often very different from rational expectations, despite this being a fundamental justification of the rational expectations revolution (Coibion et al., 2018). I therefore suggest optimal feasible expectations as a new general class of expectations that are defined as the expectations that have the lowest predictable forecast error out of the set of expectations that agents in the real world could actually estimate. These are likely to be conditionally biased towards statistically simple specifications, so will usually be much statistically simpler than rational expectations.

Secondly, this chapter relates to the econometric literature on forecasting inflation. There are a very large number of papers that analyse different approaches for forecasting inflation, in terms of method and/or predictive variables, and compare pseudo out of sample forecast error measures. Reviews of this literature are provided for traditional econometric methods in Stock and Watson (2008) and Faust and Wright (2013), while Medeiros et al. (2020) extend this analysis to machine learning methods. A key message that emerges from these reviews is the importance of parsimony. Simple auto-regressive benchmarks forecast extremely well: they are hard to consistently out-perform and effectively impossible to consistently out-perform by a large margin at horizons less than two years. Those methods that do appear to out-perform them minimise and constrain additional estimation. These include factor models with a very limited number of macroeconomic factors (Stock and Watson, 2002), extensions to benchmark models that still only use price data but allow different components of inflation to have different effects (Stock and Watson, 2016) and very heavily pruned random forests that allow some heavily constrained effects of employment variables (Medeiros et al., 2020). Theoretical restrictions derived from DSGE models are not useful for improving forecasting performance (Giacomini, 2015), however central banks targets, or proxies for them, become the optimal forecasts at horizons much beyond two years (Faust and Wright, 2013). This ap-

appears sensible, as central banks aim to target inflation in the medium term, however at horizons of two years or less lags in the effects of monetary policy (Havranek and Rusnak, 2013) and central banks' preferences for gradual adjustment of interest rates (Coibion and Gorodnichenko, 2012b) suggest that inflation deviations from targets are forecastable.

This literature currently does not address precisely why it is not optimal to add information on particular variables to auto-regressive benchmarks and I contribute to this literature by studying why this is the case. I initially assess how much shrinkage is optimal to apply to a series of macroeconomic variables, that include the main variables often used in macroeconomic models and variables commonly used in inflation forecasting. In line with the existing literature I find that the majority of variables should have total shrinkage applied to them, implying that one should virtually only use information on price series to form inflation forecasts. This information should not be used naively though, as different types of inflation should be allowed to have effects that differ but are constrained to limit estimation error. However I go on to provide the first comparisons of the shrunken estimates of the association between each variable and future inflation that is optimal for forecasting and consistent OLS estimates of the equivalent actual association. This allows me to analyse whether the high optimal degree of shrinkage comes from the variables simply not having much of an association with future inflation or from the benefits of reducing the variance of the forecast despite this imposing conditional biases because the variables having strong associations with future inflation. The results suggest that for many variables, such as broad inflation and measures of business and financial cycles it is the former, although for variables like wages it is the latter. This is important as it suggests that it is primarily the high degree of uncertainty over the associations between some variables and future inflation that prevents them from being useful in forecasting inflation, rather than the variables simply not having much association with future inflation.

Thirdly, this chapter relates to the literature which tests for conditional biases, and hence deviations from rational expectations, in surveys of agents inflation expectations. The main method of testing this in the literature, and the approach

used in this chapter, is to test whether inflation forecasts and future realised inflation react differently to information that was publicly available at the time of the forecast. Coibion et al. (2018) surveys papers that take this approach. Variables that have been suggested to cause a different response in forecast and realised inflation include lagged forecast errors (Coibion and Gorodnichenko, 2015a), lagged changes in exchange rates (Pesaran and Weale, 2006), narrative shocks (Coibion and Gorodnichenko, 2012a) and lagged energy components of inflation (Coibion and Gorodnichenko, 2015b). Understanding which variables there is a conditionally biased reaction to is important, as this determines which nominal rigidities occur and so helps us to understand how large the nominal rigidities are for the transmission mechanisms of different macroeconomic shocks.

Suggested explanations usually focus on non-optimal behaviour⁶, often resulting from some combination of rational inattention or imperfect understandings of the economy (Coibion et al., 2018). This must be at least partly true, as Berge (2018) shows that agent's inflation forecasts can be beaten in pseudo out of sample forecasting by simple auto-regressive moving average models that would have been feasible for agents to use. However it is very important to understand whether some of the specific conditional biases actually arise from optimal feasible behaviour, and so could not be corrected, or if they all arise from potentially correctable behavioural errors. I contribute to this literature by providing what, to my knowledge, is the first evidence on this issue. Using methods similar to the existing literature I estimate the conditional biases in surveys of US household inflation forecasts with respect to a set of macroeconomic variables and show that household forecasts are not optimal feasible expectations as they can be beaten by simple auto-regressive benchmarks. However unlike the existing literature, I then go on to compare the conditional biases in household forecasts to the conditional biases in estimated optimal feasible expectations. I find that the conditional biases in the reaction to many variables, such as broad and narrow inflation, business cycles and exchange rates,

⁶Explanations for some variables, such as aggregate forecast revisions, also include that information on them might not be available in real time, but this is not an issue in this chapter as we only consider variables that are publicly available.

are consistent with suggested optimal feasible behaviour. However the reaction to financial cycle information and the amount of noise in household inflation forecasts do not appear to be consistent with optimal feasible behaviour and instead suggest behavioural mistakes. Therefore optimal feasible expectations and behavioural mistakes are each likely to explain part of US households' inflation forecasts.

The rest of the chapter is organised as follows. Section 1.2 lays out my conceptual and econometric framework, Section 1.3 describes the macroeconomic information used and how I combine some of this information into factors, Section 1.4 presents the estimates of the conditional biases in optimal feasible inflation expectations, Section 1.5 estimates the conditional biases in surveys of household inflation expectations and compares these to the estimated optimal conditional biases and Section 1.6 offers some concluding remarks.

1.2 Conceptual and econometric setup

To clarify the definitions that follow I begin by decomposing future inflation into a component based on public information that is currently available and a component that is unrelated to this information. I then also express forecast inflation in terms of public information that is currently available, as follows:

$$\pi_{t+h}^r = x_t \beta^r + \epsilon_{t+h} \quad (1.1)$$

$$\pi_{t+h}^f = x_t \beta^f \quad (1.2)$$

where π_{t+h}^r is inflation at time $t+h$, π_{t+h}^f is an agent's forecast at time t of inflation at time $t+h$, x_t is a vector of information that is publicly available at time t , β^r is a vector of true coefficients, β^f is a vector of coefficients that agents use in their forecasts and ϵ_{t+h} is the component of inflation at time $t+h$ that is unpredictable at a time t with public information.

This expression is very general, as x_t could include lagged information or information which is non-linear in underlying indicators. It could also include information that is unrelated to future inflation, so that some of the values in β^r could be zero.

The definitions of the terms used are then as follows. I define the set of feasible expectations as expectations based on choices of β^f that agents can actually use in realistic settings. For instance it would be feasible to use OLS to estimate the values based on past observations. It would also be feasible to choose to set the value on lagged inflation to one and all other values to zero. I define optimal feasible expectations as the specific expectations in the set of feasible expectations that ex ante can be predicted to minimise the relevant measure of out of sample forecast error. Rational expectations are defined following Sheffrin (1996), and originally Muth (1961), as expectations that are equal to the true conditional expectation of future variables, conditioning on available information. Applying this definition in this settings yields that rational expectations are the expectations given when $\beta^f = \beta^r$.

I now consider whether rational expectations will be the optimal feasible expectations in realistic settings. First it is important to note that if rational expectations are feasible then they will be optimal, as defined by the mean square forecast error⁷, since the conditional expectation statistically minimises mean squared forecast error (Granger and Newbold, 1986). If an agent had infinite relevant data to learn from then they could use any consistent estimator of β^r to obtain an estimate essentially equal to β^r that could then be used to construct rational expectations⁸. For instance one could use past observations to estimate Equation 1.1 using OLS with all potential predictors of inflation in x_t to obtain an estimate of β^r that is statistically perfect. The agent could then use this perfect estimate of β^r as β^f , so rational expectations are feasible in this scenario and hence they are also the optimal feasible expectation.

However in reality, agents clearly only have a finite sample of data available to them. Forecasts often need to be constructed at horizons of at least a year, however samples of relevant data are usually short relative to these horizons and will not

⁷A single point forecast can only generally minimise a single forecast accuracy measure and the mean square forecast error is one of the most common measures Diebold (2017).

⁸Technically this applies to stationary variables. One would need to difference non-stationary variables until stationarity was achieved before applying this process. Then the results of this process and the current values of the variables could then be used to construct rational expectations.

necessarily increase over time, as economies experience huge structural changes that decrease the relevance of older data. For instance, formal tests (Stock and Watson, 1996) and institutional change suggest that the economic dynamics of countries now are very different from the dynamics in the period before the 1980s, when most policymakers were fully Keynesian and the internet had not yet been invented. They are likely to be even more different to the dynamics from earlier periods when many of these countries engaged in active global wars with one another. Therefore data from previous structural eras is unlikely to be of significant quantitative relevance for an agent seriously engaged in inflation forecasting (Stock and Watson, 2008). There are also strong reasons to believe that this phenomenon will continue in the future. For instance it seems extremely likely that there will be significant structural economic change as a result of the rise of artificial intelligence, new shocks such as Covid-19 and the increased economic importance of countries like China.

In reality there are huge number of potential predictors that are likely to have some effects on inflation relative to samples of data of these lengths⁹, as any variable that affects how firms set prices will have some effect on future inflation at shorter horizons. Combining similar variables may reduce the number of series that could be used but lags and non-linear transformations will increase this number and it will remain very large in practice. For instance Refinitiv Datastream and similar services provides millions of macroeconomic data series yet even samples dating to World War 2 only contains hundreds of months of observations. Therefore using conditionally unbiased approaches is simply not feasible. For instance, OLS estimates of Equation 1.1 cannot be estimated while including many of the macroeconomic series that are available. Therefore agents will generally need to use an estimation approach that shrinks forecasts, partially or even absolutely, towards statistically simpler specifications for estimation to be feasible. This implies that in practice all feasible expectations are likely to contain conditional biases, so rational expectations will not be feasible.

Even if one had incorporated enough shrinkage to make estimation feasible it

⁹A phenomenon that has been more broadly described as big data in macroeconomics being 'fat' data, with many series but relatively few observations of each series (Koop, 2017)

may well be optimal to include more shrinkage. The optimal feasible approach needs to optimally balance conditional forecast bias against conditional forecast variance, conditioning on the information available. This can be seen most clearly when using the mean squared forecast error as the measure of forecast performance. Consider the following decomposition of the mean squared forecast error, where all expectations are conditional on the information in x_t and the decomposition uses Equation 1.1, into the components that contribute to it:

$$\begin{aligned}
MSFE &= E(\pi_{t+h}^f - \pi_{t+h}^r)^2 \\
&= E(\epsilon_{t+h})^2 + E(x_t\beta^f - x_t\beta^r)^2 - 2E(\epsilon_{t+h}(x_t\beta^f - x_t\beta^r)) \\
&= E(\epsilon_{t+h})^2 + E(x_t\beta^f - E(x_t\beta^f) + E(x_t\beta^f) - x_t\beta^r)^2 \\
&= E(\epsilon_{t+h})^2 + E(x_t\beta^f - E(x_t\beta^f))^2 + (E(x_t\beta^f) - x_t\beta^r)^2 \\
&= \text{unpredictable component} + \text{forecast's variance} + (\text{forecast's bias})^2
\end{aligned} \tag{1.3}$$

The choice of the parameters, β^f cannot change the unpredictable component but they will affect the conditional variance and the conditional bias. An approach that is just feasible, such as using OLS estimates of Equation 1.1 with as many series in x_t as observations may minimise conditional biases, but is also very likely to impart a large amount of estimation error that contributes to conditional variance. Whereas using an approach that did not involve estimation, such as assuming a random walk, would minimise conditional variance but is very likely to impart conditional bias. Therefore there is typically a bias-variance trade-off to consider in the choice of how much shrinkage an agent should use when choosing β^f .

The statistically simple specifications that it is optimal to shrink forecasts towards will not usually be given by theoretical macroeconomic models. On a purely empirical level this is currently true, as the literature survey in Giacomini (2015) shows that the full results of quantitative macroeconomic models are not useful for improving the forecasts of typical macroeconomic variables given by purely statistical approaches. Giacomini (2015) suggests that the limited results to the contrary

are a product of the data mining that is fundamental in creating a theoretical model of an economy based on recent experience and then testing its ability to forecast in a sample that includes the periods on which recent experience is based. On a more fundamental level it is likely to continue to be true as theoretical macroeconomic models usually only offer predictions conditional on structural shocks and state variables that are not well observed in practice (Chung and Xiao, 2013), so proxies for them may not have the predicted effects.

There are a limited number of cases where useful guesses of coefficients in β^f can be deduced without data¹⁰, some of which are discussed in Giacomini (2015). In a very limited number of cases these may even allow expectations that are close to rational to be used: for instance heavily shrinking long-term inflation forecasts in some countries towards the countries inflation target. However in the vast majority of cases where there is no such information available the natural choice to shrink coefficients in β^f towards is zero. The optimal degree of shrinkage can then be based on a combination of how relevant an agent thinks a variable is likely to be, for instance ruling out variables that are likely to have little association with the macroeconomy in question so are unlikely to have large effects, and empirical methods, such as pseudo out of sample tests or Bayesian model averaging.

I therefore suggest a new class of expectations: optimal feasible expectations. These are formally defined as the point expectations that are ex ante predicted to minimise the relevant measure of forecast error out of the set of expectations that it is feasible for agents to use in practice¹¹. Based on the above discussion I suggest that in the vast majority of macroeconomic settings optimal feasible expectations are likely to be statistically simpler than rational expectations, as many variables effects will be shrunk significantly towards zero, so they will generally incorporate

¹⁰However shrinking coefficient towards these values may actually increase the degree of shrinkage in optimal feasible expectations relative to shrinking them towards zero, as the same reduction in conditional forecast variance from absolute shrinkage could then be achieved with less conditional bias.

¹¹Optimal feasible expectations could therefore vary for different agents if they aim to minimise sufficiently different measures of forecast error in the same setting. However this is partly a product of analysing point forecasts and is not the focus of this paper, so is not explored here.

conditional biases. However the exact size of the conditional biases in optimal feasible expectations, and hence their differences with rational expectations, is an empirical question. It depends on the degree of shrinkage that is optimal to apply to variables that have large associations with future inflation. I therefore now turn to examining the degree of optimal shrinkage to apply to variables in US inflation forecasting. The specific variables I choose are ones that are thought to transmit shocks to inflation in many macroeconomic models and so have long been suggested in the literature as potentially having associations with future inflation.

It is worth noting however that even if empirically observed shrinkage is relatively small this could imply large deviations from rational expectations equilibria, as it may represent the endpoint of a feedback loop. For instance, consider agents applying shrinkage with regards to information on a macroeconomic shock, so that their expectations responded less than rational expectations would to information on the shock. This could in turn reduce the response of realised inflation itself to the shock, relative to rational expectations equilibria, which could lead to even greater differences between expectations and those in rational expectations equilibria, creating a feedback loop. Therefore actual data may be generated by the end point of such a feedback loop and relatively limited empirical shrinkage could still imply large nominal rigidities relative to comparable rational expectations equilibria.

Since estimating without shrinkage is infeasible in reality, as discussed above, I start with an extremely parsimonious specification and then consider how much it is worth shrinking the effects of additional macroeconomic variables that are added to this benchmark¹². As well as estimating the degree of shrinkage that is optimal to apply to the predictive associations of each of these variables when forecasting future inflation, I also estimate the true association between each variable and future inflation. This lets me analyse the size of the conditional biases present in the estimated optimal feasible inflation expectations.

The baseline specification that I start with is an extremely simple direct autoregressive model estimated by OLS. It simply expresses inflation at time $t + h$, π_{t+h}^r ,

¹²This also seems sensible given the importance placed on extreme parsimony by the inflation forecasting literature discussed in Section 1.1.

in terms of inflation at time t , π_t^r , and a constant:

$$\pi_{t+h}^r = \gamma_0 + \gamma_1 \pi_t^r + \varepsilon_{t+h} \quad (1.4)$$

I can then consider the optimal level of shrinkage to apply to additional macroeconomic variables¹³ using pseudo out of sample inflation forecasting performance. To do this I take many overlapping sub-samples from my sample and then in each of these training sub-samples calculate estimates of the coefficients on these additional variables with different levels of shrinkage. The optimal level of shrinkage can then be taken as the one which minimises the pseudo out of sample forecasting error from the remaining test datasets. The approach is therefore a conservative one for estimating the optimal degree of shrinkage, as a new test set is not used for every variable. This pseudo out of sample approach is a common method of setting the level of shrinkage in machine learning approaches such as those in Medeiros et al. (2020). Note that the optimal level of shrinkage will become very small as the sample becomes very large, so this approach can still produce a consistent forecast.

I implement the shrunk estimates using weighted ridge regression¹⁴, which can shrink different coefficients by different quantities and can be expressed as a linear transformation of OLS so can be calculated analytically. I only apply shrinkage to the additional variable that is included. Weighted ridge regression minimises a loss function which combines the OLS loss function with a quadratic penalisation term, so the WR loss function and the OLS loss function can be expressed as follows:

$$Loss\ Function^{WR} = (\Pi - X\beta)'(\Pi - X\beta) + \beta'\Lambda\beta \quad (1.5)$$

$$Loss\ Function^{OLS} = (\Pi - X\beta)'(\Pi - X\beta) \quad (1.6)$$

¹³These variables are only included in linear form, which is conservative as there are so many potential non-linear transformations of variables that attempting to include all of them would require the use of significant shrinkage for estimation to be feasible.

¹⁴Given the maximum number of variables considered is low, there is little difference between the forecasts produced with this method and alternatives such as lasso or elastic net shrinkage. However there is an analytical solution for weighted ridge, unlike for lasso or elastic net penalisation, making the bootstrapping process used dramatically faster.

where Λ is a diagonal shrinkage matrix in which the values corresponding to the constant and lagged inflation are zero while the value corresponding to the additional variable considered is positive or zero, Π is the vector formed by stacking the dependent realised inflation variable, X is the matrix formed by stacking independent variables and β is the vector formed by stacking coefficients.

Therefore, for each additional variable considered, I estimate the following specification by OLS for the whole period and by a series of weighted ridges with multiple levels of shrinkage over each of a series of training periods for each additional variable v^i :

$$\pi_{t+h}^r = \gamma_0 + \gamma_1 \pi_t^r + \alpha v_t^i + \varepsilon_{t+h} \quad (1.7)$$

In all specifications I shrink the coefficient on the additional variable included towards zero as a neutral choice and apply no shrinkage to the mean and autoregressive term. The out of sample forecasting results with different levels of shrunken coefficients then provide estimates of the optimal level of shrinkage to be applied to different key variables. Comparing these optimal shrunken coefficients to the OLS coefficients is then an estimate of the conditional biases imposed on the information contained in these variables¹⁵. The greater the difference between the optimal shrunken coefficients and the OLS coefficients the greater the conditional biases in estimated optimal feasible expectations. Larger conditional biases imply larger deviations of optimal feasible expectations from rational expectations and so larger nominal rigidities that arise from inflation expectations.

The deviation of optimal feasible expectations from rational expectations implies that any conditional biases in agent's forecasts are not necessarily a deviation from optimal behaviour in the real world. Therefore in the second part of my analysis I estimate if there are conditional biases in household inflation forecasts with respect to the variables considered above and compare these conditional biases to those in the estimated optimal feasible expectations. To do this, I compare the OLS and weighted ridge estimates from the previous specification with OLS estimates of the

¹⁵Note they will include both the direct information included in the variable itself and the information included through its correlations with all other variables that have not been included.

equivalent specification with household inflation forecasts as the dependent variable as follows:

$$\pi_{t+h}^f = \gamma_0 + \gamma_1 \pi_t^r + \alpha v_t^i + \varepsilon_{t+h} \quad (1.8)$$

Differences between the estimated OLS coefficients with realised inflation and household inflation forecasts as the dependent variables imply that there are conditional biases in household forecasts¹⁶, so household forecasts deviate from rational expectations. Note that this is true even if one includes a subset of the data available to agents (Sheffrin, 1996). I formally test the differences between these coefficients from Equations 1.7 and 1.8 and obtain confidence intervals for the difference using a joint block-bootstrap. Specifically, I use a bias-corrected version of Hall’s empirical bootstrap approach, which allows for auto-correlated errors and parameter distributions which are skewed and incorrectly centered.

I also compare the similarities between the OLS coefficients with household inflation forecasts as the dependent variables and the estimated optimal weighted ridge coefficients. This is because similarities suggest that the conditional biases considered are consistent with the conditional biases in estimated optimal feasible expectations, whereas differences suggest they are not. Finally, I also check whether households’ forecasts are consistent with being optimal feasible expectations by comparing their out of sample forecast performance with that given by my parsimonious benchmark, as this benchmark is feasible and approaches like this have long been known to forecast reasonably well (Gordon, 1982). If the household forecast performance is as good as or better than the estimated forecast performance of this benchmark then this is consistent with households using optimal feasible expectations. Although one should remember that there clearly may be better feasible alternatives to my benchmark available, so this a necessary and not sufficient con-

¹⁶It is theoretically possible that ‘peso problems’ could explain such differences in short samples, however this seems unlikely to be important in a sample which includes the financial crisis, the dot-com bubble and many other extreme events. Additionally the effects of large infrequent events seem especially unlikely to be something that households could estimate perfectly and so incorporate in line with rational expectations.

dition. However if households' forecast perform worse than my benchmark then this strongly implies that households make behavioural mistakes that cause their expectations to deviate from optimal feasible expectations.

1.3 Macroeconomic data and factors

My two primary dependent variables are household inflation forecasts and realised inflation. The household inflation forecasts are taken from the Michigan Survey of Consumers: they are one year ahead inflation forecasts and the questionnaire aims for quantitative responses with prompts provided if necessary¹⁷. I choose an annual horizon as this is long enough for many shocks to have some inflationary effects, but is not long enough for the Federal Reserve to have resolved these effects, due to lags in the effect of monetary policy (Havranek and Rusnak, 2013) and a preference for gradual monetary policy action (Coibion and Gorodnichenko, 2012b). The household forecasts are usually based on a sample of approximately 500 people, of which up to 20% give non-quantitative answers. These non-quantitative answers are hard to reconcile with rational expectations or optimal feasible expectations, so strongly suggest that household forecasts may not be optimal even before any formal analysis is conducted.

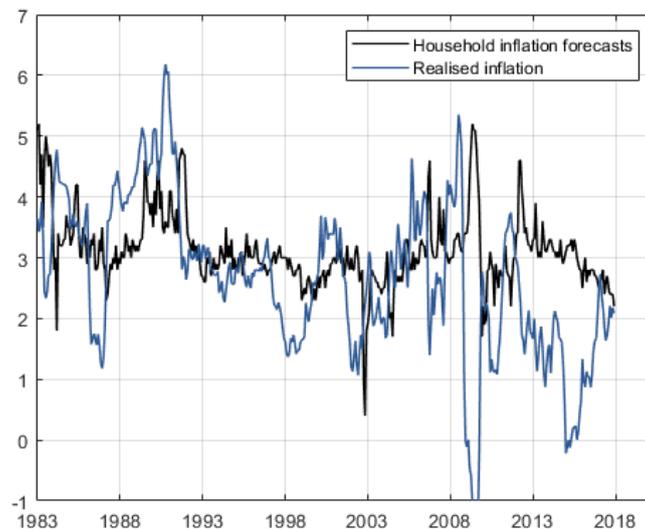
I take the consumer price index as my measure of realised inflation. This is because its definition is methodologically most suitable, as it aims to capture the experienced inflation of consumers, which is not true of alternatives such as the personalised consumption expenditures index. Its mean level is also closer to the mean household inflation forecast than the mean level of alternatives such as the personalised consumption expenditures index, which supports it being the appropriate inflation index. Inflation is widely considered to be stationary in the absence of structural breaks, as it does not seem feasible that the central bank would allow significant deviations from its goals and its long-term mean has remained similar over recent decades. Breitung and Eickmeier (2011) find that there is a structural break in inflation at the start of Paul Volcker's chairmanship of the Federal Reserve

¹⁷The specific questionnaire can be accessed online through <http://www.sca.isr.umich.edu/>

and several of the variables in my sample are only available from close to this point onwards, so I begin my sample near this point. If there is any additional structural change over time this should be picked up by the inflation factor, so will not cause spurious results.

The main sample for inflation as a dependent variable therefore runs from the annual price growth up to January 1983 to the annual price growth up to December 2017. The sample for dependent inflation expectations necessarily covers the expectations for the same period and the sample for control variables is lagged by a year.

Figure 1.1: Realised inflation and household inflation forecasts



Notes: Plots of the median household annual inflation forecast for the past year and annual consumer price index inflation. The vertical axis is in percentage points and the horizontal axis is in years.

Figure 1.1 plots household inflation forecasts and realised inflation over the sample. Household forecasts are generally of a similar approximate level to realised inflation, however there are several features which may seem surprising if one were expecting inflation forecasts to be formed by rational expectations. Spikes in household forecasts often follow, instead of precede, spikes in realised inflation and there

are long periods of divergence between forecasts and realised inflation. These features suggest that inflation expectations are not formed rationally, although this will be examined in much more detail in Section 1.5.

The additional macroeconomic variables I consider adding to the forecasting procedure include some of the most important potential transmitters of shocks to inflation and a narrative measure of aggregate demand shocks available in real time. These potential predictors include six series: corresponding to business cycles, financial cycles, broad inflation, wages, exchange rates and real-time narrative monetary shocks. They therefore include equivalents of the macroeconomic variables suggested as potential predictors of inflation in Stock and Watson (2008) as well as one of very few narrative measures of shocks available in real time. I do not include the measures from specific financial markets that Stock and Watson (2008) suggest including as proxies for expectations themselves, as in this chapter the expectations being formed are viewed as the dependent variable to be explained, so including proxies for them as an independent variable would not be helpful. In all cases I take the variables from the Fred MD database or the broader FRED database except for the monetary shocks which are taken from Gertler and Karadi (2015).

In the case of the first three series (business cycles, financial cycles and broad inflation) many monthly measures are available so I combine them using a factor approach, whereas for the latter three series (wages, exchange rates and real-time narrative monetary shocks) there are few series available so I simply use the corresponding series in Fred MD or Gertler and Karadi (2015). I produce the factors using the principal components approach of Stock and Watson (2002). This is applied separately to different groups of variables, so each variable only loads on one factor, as suggested by Bernanke et al. (2005). This ensures statistical identification and also gives each factor a clear economic interpretation. Bai and Ng (2006) show that the factors converge at rate $\min(N, T)$, whereas if the factors were known then the coefficients would converge at rate \sqrt{T} . Therefore it is a reasonable approximation to treat the factors as known if N is reasonably large compared to \sqrt{T} , which is the case here. Indeed, it may well still be an improvement over using specific variables to proxy for each factor, which would remove the estimation issue but

potentially introduce significant measurement error.

I take the majority of the variables from the Fred MD database. For the business cycle factor I take 15 variables from the output and income section and 21 variables from the labour market section, which are all in real terms. For the price factor I take 19 variables from the prices section and add 16 extra price series from the broader FRED database. For the financial cycle factor I take 7 credit series from the money and credit section and add 2 extra credit series and 31 house price series from the broader FRED database. For the exchange rate series I take the trade weighted US Dollar index against major currencies, where a rise implies an appreciation of the dollar, and for the wage series I take the average hourly earnings of goods producing workers. This gives 36 business cycle series, 35 inflation series and 40 financial cycle variables in a sample where \sqrt{T} is approximately 20. Therefore in each case \sqrt{T}/N is small, so any estimation error in the factors will be limited relative to estimation error of the coefficients.

The narrative monetary shocks are taken from Gertler and Karadi (2015). They are constructed as the high frequency changes in federal funds futures markets around federal reserve announcements and are discussed and contrasted to other shocks in Ramey (2016). It is important to note that I do not necessarily give the narrative shocks a causal interpretation, as Miranda-Agrippino (2016) shows that they respond to Federal Reserve forecasts. I instead simply view them as one of the most widely-used and reliable measures of monetary shocks available in real time. The sample period is shorter when monetary policy shocks are used, as data is not available for the earlier part of the sample and not usable for the latter part of the sample due to the zero lower bound (Gertler and Karadi, 2015). The sample for inflation as a dependent variable when monetary shocks are used runs from the annual price growth up to July 1990 to the annual price growth up to June 2012. The sample for dependent inflation expectations necessarily covers the expectations for the same period and the sample for control variables is lagged by a year.

All variables are transformed to stationarity, which is primarily by using the FRED MD recommended transformations expressed in annual terms. However inflation is considered stationary over my sample, as discussed above, since it is shorter

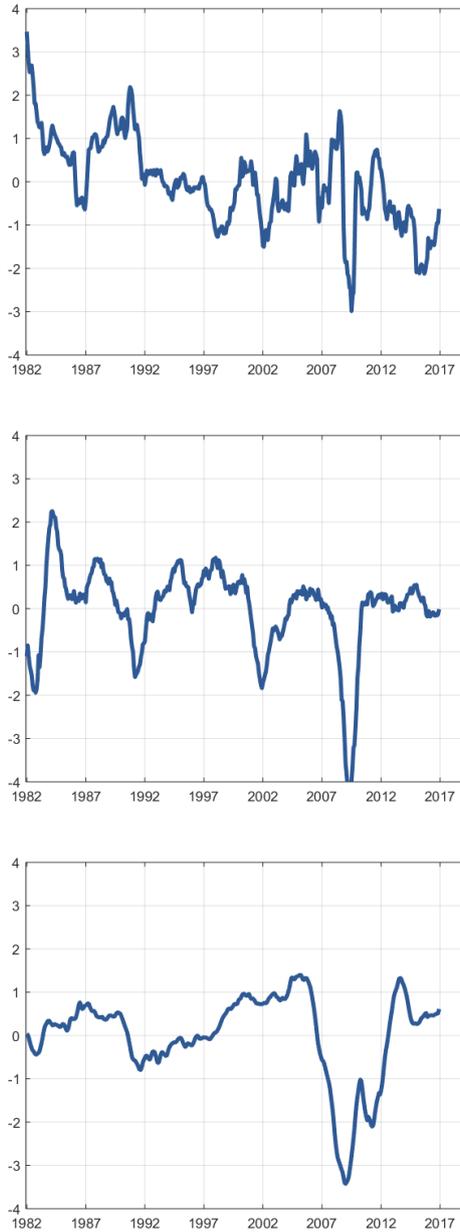
than the FRED MD sample, so I do not take the second difference of nominal series. All the variables used in the factors are normalised to have zero mean and unit variance before factors are extracted from them. While this uses data from the whole sample in a forecasting exercise it does not change any variable substantially, but just makes them easier to combine and compare. They are also normalised to load positively on a measure of employment, inflation and house prices respectively. I also transform all six additional series that shrinkage is applied to so they have zero mean and unit variance for comparability. A full table of the variables used and the transformations applied is available in Appendix B.

Figure 1.2 plots the inflation, business cycle and financial cycle factor over the sample. For the inflation cycle, the early parts of the sample contains the large effects of the supply side crises which the Federal Reserve was starting to control, such as the oil price surge at the end of the 1970s. The latter part of the sample has more short-term volatility, although one can see the dip associated with the aftermath of the financial crisis and the subsequent dip associated with the global economic slowdown in 2015 to 2016. The four recessions in the sample are all clearly visible in the business cycle factor and are marked by increases in growth in the recovery after each one. Indeed, this factor could proxy fairly well for the NBER business cycle dating. The financial cycle factor is loosely similar, however the effects of the first two recessions are small and the third is virtually absent, whereas the latter part of the sample is dominated by the huge effects associated with the global financial crisis.

The three factors all load sensibly on their underlying features. In fact, every single variable loads on its factor with the expected sign: all positive for the inflation factor, all positive for the financial cycle factor, negative for the unemployment series and positive for all other series for the business cycle factor. The magnitudes of the factor loadings are also sensible: most are between 0.2 and 0.8 and none are dramatically outside this range¹⁸. Therefore the factors appear to capture the information in the inflation, business cycle and financial cycle indicators well.

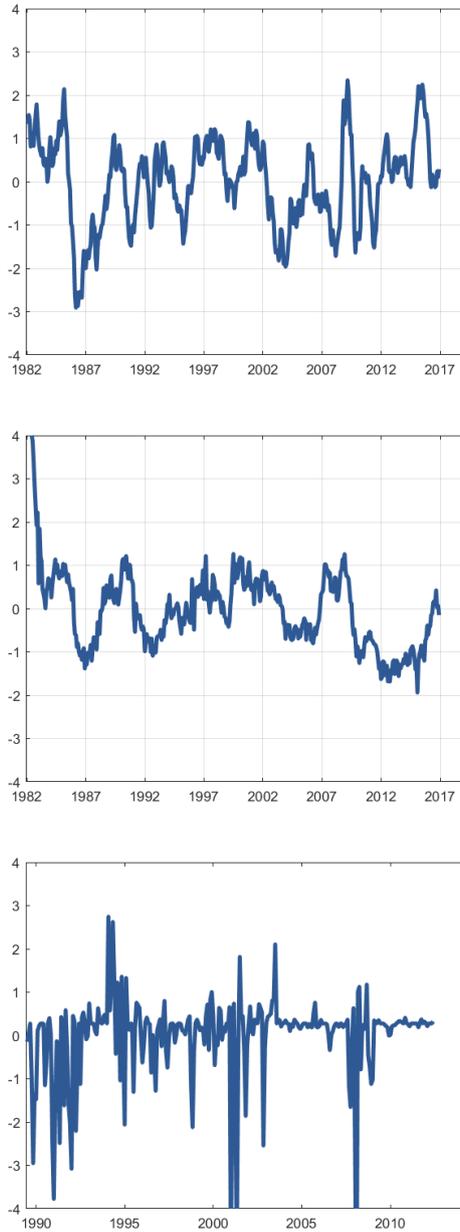
¹⁸Note that this does not imply that the underlying variables move less in absolute terms than the factor, as they have been normalised to have unit variance.

Figure 1.2: Inflation factor, business cycle factor and financial cycle factor



Notes: Plots of the factors extracted with the principal components method from transformed data. The vertical axis is in units and the horizontal axis is in years. Since the underlying series are transformed to have zero mean and unit variance and most factor loadings are between 0.2 and 0.8, a one unit change in each factor causes changes in most of its underlying series of between 0.2 and 0.8 standard deviations.

Figure 1.3: Exchange rates, wages and narrative futures markets monetary shocks



Notes: Plots of the transformed trade weighted exchange rate index, transformed average hourly earnings and transformed narrative futures markets monetary shocks. The vertical axis is in standard deviations of each variable units and the horizontal axis is in years.

Exchange rates, wages and narrative monetary shocks are plotted in Figure 1.3. There are few clear patterns in the exchange rate graph, as its movements are quite

volatile. However one can pick out certain large movements, such as the large depreciation in the latter part of the 1980s following the Plaza Accord and the large appreciation in 2014/2015 following monetary divergence between the Federal Reserve and many other developed market central banks. The wage series is also relatively volatile, although one can again notice several large movements, such as the very high wages at the start of the sample as the inflation-wage spiral was being brought under control and the large and sustained declines in wages that occurred in the period following the global financial crisis. The narrative futures market monetary shocks series is the most volatile of all. However one can see especially high volatility in the earlier part of the sample, as well as in the period around the 9/11 attacks and in the period around the global financial crisis.

1.4 Conditional biases in optimal feasible inflation expectations

I now turn to estimating the optimal degree of shrinkage to apply to each additional variable in inflation forecasts. As discussed in Section 1.2, I do this using pseudo out of sample inflation forecasting performance. For each variable considered I take many overlapping sub-samples from my sample and then in each of these training sub-samples calculate estimates of Equation 1.7 with many different levels of shrinkage. I only apply shrinkage to the additional variable added to the baseline specification in each case, so the most shrunken specification corresponds to OLS estimation of the auto-regressive specification in Equation 1.4 while the least shrunken case corresponds to OLS estimation of the specification in Equation 1.7, with other levels of shrinkage giving estimates between the two. The estimated optimal level of shrinkage to apply to each variable can then be taken as the one which minimises measures of pseudo out of sample forecasting error from the remaining test data. The training sample sizes are set to 70% of the total sample size in the baseline case, which is relatively typical¹⁹ and ensures the financial crisis period can be in

¹⁹This gives approximately the same probability of any one observation being in the sample as would be the case if one took a sample with replacement of the same size as the original sample.

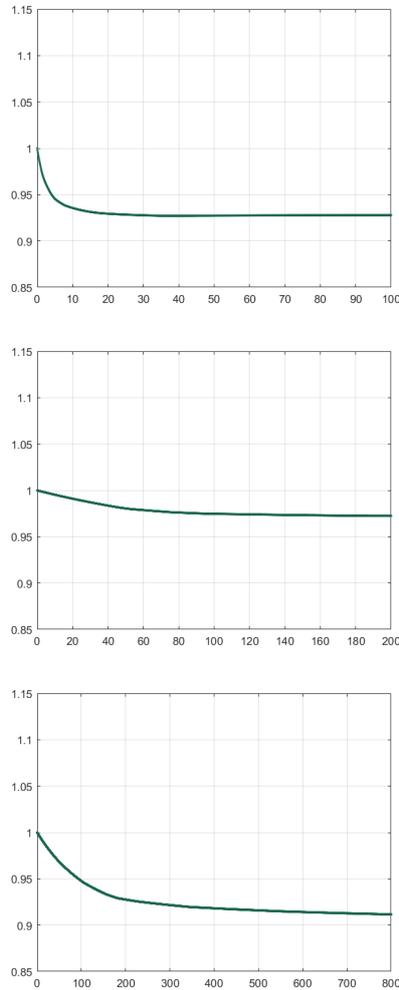
both types of sub-sample. However robustness checks based on increasing or decreasing this sample size are available in Appendix B and the results do not change dramatically in either case.

Figures 1.4 and 1.5 plot the out of sample forecast performance for different levels of shrinkage applied to each variable. The forecast performance measure used is mean absolute forecast error, although analysis which uses the mean square forecast error is also available in Appendix B and is very similar²⁰. The forecast performance is expressed relative to the forecast error using no shrinkage, i.e. that obtained using OLS with the additional variable in question. Therefore values lower than one imply superior performance and values above one imply inferior performance. The optimum shrunken value of each coefficient is reported and compared to the OLS estimate of each coefficient in Table 1.1.

The first variable I consider is the inflation factor, which captures simultaneous changes in a broad range of the price components of inflation. Therefore including information on this variable allows broad price changes, such as caused by rising consumer confidence, to have different effects on the forecasts produced than the effects of a change in inflation driven by large changes in a small number of prices, such as change in the price of food or oil. The results in the top graph of Figure 1.4 make it clear that using some shrinkage improves forecast performance: it can reduce the forecast error measure by over 5%. The figures in Table 1.1 actually show that partial shrinkage is optimal: this is also plotted in the top graph of Figure 1.4 but is hard to see clearly. This implies that the optimal coefficient on broad inflation when forecasting should be lower than the OLS coefficient of its association with future inflation, but should not necessarily to zero. Although using complete shrinkage and setting the coefficient equal to zero barely reduces the forecast performance from its optimal level. The OLS estimates of the true association between broad inflation and future inflation shows is positive and both economically and statistically significant. Therefore this strongly suggests that optimal feasible expectations should incorporate large conditional biases with respect to information of broad vs

²⁰These are chosen as they are two of the most common forecast measures (Diebold, 2017) and because the mean square forecast error is the equivalent measure for rational expectations.

Figure 1.4: Relative forecast error from shrinking information on inflation cycles, business cycles and financial cycles



Notes: Plots of the out of sample mean absolute forecast error of the shrunken estimates of Equation 1.7 with an additional variable included, presented relative to the out of sample mean absolute forecast error of the equivalent OLS estimate of Equation 1.7. The inflation factor (top), the business cycle factor (middle) and the financial cycle factor (bottom) are considered. Estimates are based on training sample of 70% of the total dataset. The vertical axis is in relative units, so higher values imply worse performance relative to the OLS case. The horizontal axis is in values of λ , where higher values of λ imply more shrinkage. When $\lambda \rightarrow \infty$ the specification tends to Equation 1.4.

narrow inflation, as a result of shrinkage.

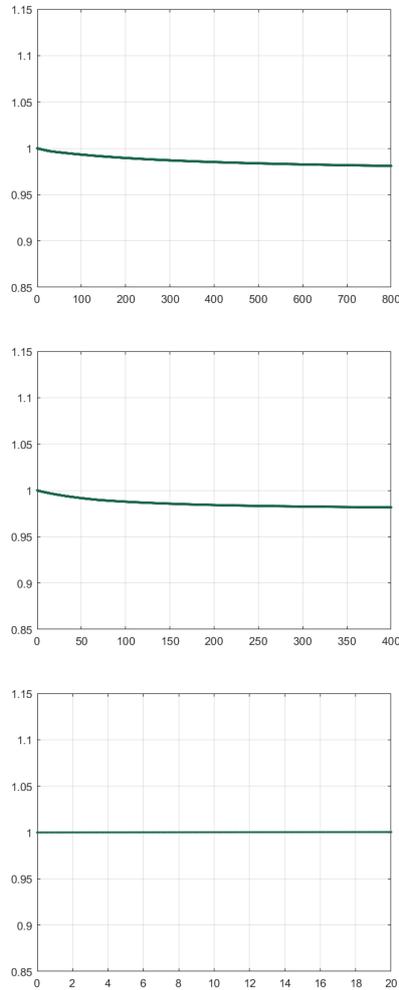
The next variables I consider are the business cycle and financial cycle factors. The first of these captures the movements of a set of macroeconomic indicators while the second captures the movements of longer term variables in financial markets. The results in the bottom two graphs in Figure 1.4 and in Table 1.1 suggest that absolute shrinkage should be applied to these factors, as this reduces the forecast error measure by around 3% and 8% respectively, so optimal feasible expectations should not incorporate information on these variables at all. This suggests that optimal feasible expectations incorporate conditional biases with respect to information on business and financial cycles, as the OLS estimates of their associations with future inflation are positive and economically meaningful, albeit just short of statistical significance.

The variables considered in the upper two graphs in Figure 1.5 and also in Table 1.1 are the exchange rate and wage series, which are changes in the trade weighted value of the dollar and hourly earnings respectively. The results again indicate that absolute shrinkage should be applied to these series, as this reduces the forecast error measure meaningfully, so optimal feasible expectations should not incorporate them at all. In the case of wages this appears to be because they have little association with future inflation²¹. However exchange rates have a negative and economically meaningful association with future inflation that is just short of statistical significance, suggesting that optimal feasible expectations contain conditional biases with respect to exchange rate information.

Finally I consider the narrative monetary shock measure taken from federal funds futures markets, which is only available over a shorter sample. The results for this series in the bottom graph of Figure 1.5 and in Table 1.1 indicate that it is technically optimal not to apply shrinkage to the effects of this series, so optimal feasible expectations could include this information. However the results also show that the associations of these shocks with future inflation is very small, so that using information on the shocks does not significantly change the forecasts produced and

²¹A result which holds even if one removes the high initial values of the wage series at the very start of the sample.

Figure 1.5: Relative forecast error from shrinking information on exchange rates, wages and narrative federal funds market monetary shocks



Notes: Plots of the out of sample mean absolute forecast error of the shrunken estimates of Equation 1.7 with an additional variable included, presented relative to the out of sample mean absolute forecast error of the equivalent OLS estimate of Equation 1.7. Exchange rates (top), wage (middle) and monetary shocks (bottom) are considered. Estimates are based on training sample of 70% of the total dataset. The vertical axis is in relative units, so higher values imply worse performance relative to the OLS case. The horizontal axis is in values of λ , where higher values of λ imply more shrinkage. When $\lambda \rightarrow \infty$ the specification tends to Equation 1.4.

Table 1.1: Estimated true association and optimal associations for forecasting between variables and future inflation

	OLS	Optimal WR
γ_0	1.91* (1.32 to 2.46)	
γ_1	0.27* (0.06 to 0.50)	
α^{inf}	2.08* (1.28 to 2.80)	0.69 <i>(0.05 to 1.42)</i>
α^{bc}	0.24 (-0.07 to 0.50)	0.00 <i>(0.00 to 0.21)</i>
α^{fc}	0.30 (-0.07 to 0.69)	0.00 <i>(0.00 to 0.17)</i>
α^{er}	-0.18 (-0.44 to 0.03)	0.00 <i>(-0.13 to 0.00)</i>
α^w	0.07 (-0.26 to 0.35)	0.00 <i>(0.00 to 0.06)</i>
α^{nms}	0.03 (-0.11 to 0.21)	0.03 <i>(0.00 to 0.03)</i>

Notes: Column 1 shows the OLS estimates of the coefficients from Equation 1.4 for the baseline variables and the OLS estimates of the coefficient on each new variable from its version of Equation 1.7. 95% confidence intervals are displayed in standard text beneath OLS estimates. * = statistically significant at the 10% level. Column 2 shows the WR estimate of the coefficient on each new variable from its version of Equation 1.7 that minimises out of sample error. Bands of shrunken estimates where the forecast error is within 1% of the optimal forecast error are displayed in italicized text beneath WR estimates. By definition these bands are always non-negative or non-positive.

has hardly any effect on forecast performance. Indeed any level of shrinkage produces forecasts with values of the forecast error measure within 1% of each other. It is also interesting to note that the association between the shocks and future inflation is positive, not negative as typically suggested by theory. This may be a result of the fact that these series may capture signals of the Federal Reserve's private economic information as much as they capture true monetary shocks (Miranda-Agrippino, 2016). Therefore this does not appear to suggest that optimal feasible expectations respond to true monetary shocks.

These results have important consequences. Firstly consider how optimal feasible expectations respond to macroeconomic shocks: in particular consider the case of a contractionary monetary shock, i.e. an exogenous increase in interest rates. The response to narrative monetary shocks would imply that inflation expectations would initially rise and there would be no reaction to any change in exchange rates. Inflation expectations would then also not respond to any change in financial cycle and business cycle variables as a result of the shock. They would only start to fall after inflation itself had fallen and even then this response would still be constrained. If higher inflation expectations cause higher future inflation, as seems very likely, then this strongly suggests that optimal feasible expectations would cause large nominal rigidities in the response to such shocks.

Secondly these results provide evidence that optimal feasible inflation expectations contain large conditional biases with respect to some of the most important variables in many macroeconomic models. Therefore they suggest that agents who learn optimally from data will use expectations of key macroeconomic variables that are very different from rational expectations. This undermines one of the major arguments used to try and justify the rational expectations revolution and suggests that macroeconomic models based on rational expectations may be seriously misspecified.

1.5 Conditional biases in household inflation forecasts

I now turn to estimating the conditional biases in household inflation forecasts and assessing whether these are similar to those in optimal feasible inflation expectations. As discussed in Section 1.2, testing for conditional biases with respect to a given variable is achieved by testing if there are significant differences in the OLS estimates of coefficients from Equation 1.7 and 1.8 with that variable included. Any significant differences would imply different systematic reactions of forecast and realised inflation to the variable and so would suggest conditional biases and hence deviations of household forecasts from rational expectations. I then also analyse whether these conditional biases are the same as those in estimated optimal feasible inflation expectations from the previous section and hence whether household forecasts appear to be similar to estimated optimal feasible expectations.

As discussed in Section 1.2, I estimate Equations 1.7 and 1.8 multiple times, once with each of the six additional variables as v_t^i and once in the baseline case without v_t^i . In all cases I also calculate the difference between each equivalent coefficient from Equation 1.7 and Equation 1.8 and bootstrap confidence intervals. Table 1.2 shows the abridged results of this analysis. The left column shows the results with realised future inflation as the dependent variable, the middle column shows the results with household forecasts of inflation as the dependent variable and the right column shows the difference between the two. The first two parameters are taken from the estimations in the baseline case without any additional variables. Each of the other six parameters are taken from the estimations in the case in which the corresponding variable is v_t^i .

The response to the baseline variables is similar for realised and forecast inflation with no significant or important differences. In both cases inflation has a sensible average value²² and a positive but low auto-correlation. Therefore there do not appear to be important conditional biases in the responses to the baseline information. The response to the three factors is much more interesting. Realised inflation reacts

²²Note that the average value is not just equal to the constant.

significantly and positively to realised inflation, suggesting that broad price rises are more sustained than narrow price increases. However forecast inflation reacts far less strongly to broad inflation, so there is a significant difference between the two, implying a large conditional bias. Comparing these coefficients to the equivalent estimated optimal coefficients from Table 1.1 also suggests that this bias in household forecasts is very sensible, as it is well within the band of the optimal shrunken coefficients.

The response of realised inflation to the business cycle and financial cycle factors is positive. However the household forecasts barely respond to the business cycle factor and actually respond negatively to the financial cycle factor. There are therefore meaningful differences in the reactions to both factors, although only the financial cycle difference is statistically significant, implying conditional biases in the household forecasts. The lack of response to business cycle information is completely consistent with the optimal feasible expectation estimates in Table 1.1. However the response to financial cycle information is actually overly negative, suggesting that it may arise from a behavioural mistake²³, that leads to an even greater reduction in the coefficient than that required for optimal feasible expectations.

The response of realised inflation to exchange rates is clearly negative whereas household forecasts barely respond to exchange rates, suggesting a conditional bias although the difference between the two is just short of statistical significance. This conditional bias is completely in line with the conditional bias in optimal feasible expectations, as complete shrinkage is optimal in this case. Both wages and narrative federal funds futures markets monetary shocks only have very small associations with future inflation. They also have small associations with household forecasts, so there are no large conditional biases, although it is hard to comment on whether there are any conditional biases as the scale of their associations is too small to statistically detect biases with any confidence. Both very small coefficients are close to the equivalent coefficients in optimal feasible expectations, so are consistent with optimal feasible expectations.

²³It could also be a small sample effect driven by mistaken expectations around the global financial crisis as this event was so important for this variable.

Table 1.2: Conditional biases in household inflation forecasts

	π_{t+h}^r	π_{t+h}^f	Difference
γ_0	1.91* (1.32 to 2.46)	2.30* (2.08 to 2.54)	0.39 (-0.26 to 1.12)
γ_1	0.27* (0.06 to 0.50)	0.29* (0.21 to 0.37)	0.02 (-0.26 to 0.28)
α^{inf}	2.08* (1.28 to 2.80)	0.43 (-0.07 to 0.80)	-1.64* (-2.65 to -0.81)
α^{bc}	0.24 (-0.07 to 0.50)	-0.04 (-0.18 to 0.06)	-0.28 (-0.61 to 0.07)
α^{fc}	0.30 (-0.07 to 0.69)	-0.14* (-0.23 to -0.07)	-0.44* (-0.89 to -0.02)
α^{er}	-0.18 (-0.44 to 0.03)	0.03 (-0.05 to 0.11)	0.21 (-0.03 to 0.51)
α^w	0.07 (-0.26 to 0.35)	-0.01 (-0.12 to 0.12)	-0.08 (-0.41 to 0.34)
α^{nms}	0.03 (-0.11 to 0.21)	0.04 (-0.02 to 0.09)	0.02 (-0.20 to 0.18)

Notes: Column 1 shows the OLS estimates of Equation 1.7 with realised inflation as the dependent variable, Column 2 shows the OLS estimates of Equation 1.8 with household inflation forecasts as the dependent variable and Column 3 shows the difference between the two coefficients. 90% block bootstrapped confidence intervals are in brackets and * = statistically significant at the 10% level.

These results suggest that the conditional biases in household inflation forecasts are very similar to the conditional biases in optimal feasible inflation expectations. In fact the only conditional bias that appears to be meaningfully different is that with respect to the financial cycle factor. Therefore many of the important nominal rigidities in the response of actual household inflation expectations to shocks will

be the same as those in the response of optimal feasible inflation expectations. For instance the response of household inflation expectations to a contractionary monetary shock is likely to be similar to the response of optimal feasible expectations discussed at the end of Section 1.4. Indeed, it may actually be more rigid as a result of household inflation expectations possibly reacting positively to any decline in financial cycle indicators.

Household expectations are not consistent with being entirely formed by optimal feasible expectations however, as they are clearly beaten in forecast performance by a feasible alternative. I show this by comparing the forecast performance of the forecasts produced by estimating my baseline auto-regressive forecasts on each of the training sets of data used in Section 1.3 with the equivalent household forecasts made at the end of each training set. The results show that the mean absolute forecast error of the household forecasts is 132% of that of the auto-regressive forecasts and the mean square forecast error of the household forecasts is 160% of that of the auto-regressive forecasts. This seems sensible as Figure 1.1 suggests that the household forecasts sometimes deviate persistently from realised inflation for years at a time. Therefore both optimal feasible inflation expectations and behavioural errors that reduce forecast performance seem to be important in explaining household inflation forecasts.

1.6 Conclusion

Inflation expectations have a particular importance in macroeconomics, as they affect the degree of nominal rigidities to macroeconomic shocks and so the size of their real effects. The existing empirical literature has suggested that inflation expectations contain conditional biases with respect to publicly available macroeconomic information, causing nominal rigidities. However this is usually justified by behavioural factors, such as limited attention or imperfect cognitive abilities. While I do not deny the importance of these factors, in this chapter I primarily study whether rational expectations are the optimal feasible expectations for agents, i.e. are they the expectations that are predicted to minimise a measure of forecast error

out of the set of expectations that are feasible for agents to use.

I discuss that, with data samples of realistic length, agents will have to introduce conditional biases into their forecast in the vast majority of macroeconomic settings, due to the limited number of relevant monthly observations available relative to variables that can affect how prices are set. They can do this by shrinking their forecasts towards those given by a simpler specifications. However even if an agent had included sufficient shrinkage to make estimation feasible, it may well still be worth including additional shrinkage, as this may reduce the conditional forecast variance sufficiently to outweigh the increased conditional biases. Macroeconomic theory is unlikely to typically help to set the simpler specifications used in expectations formation, as its predictions are usually conditioned on state variables such as macroeconomic shocks and output gaps that are not well observed in real time. The importance placed on parsimony by the empirical forecasting literature and the degree to which auto-regressive benchmarks are hard to substantially beat in forecasting horse races, despite the no doubt numerous effects of many macroeconomic variables on each other, suggests that the extent of this optimal shrinkage and the conditional biases it causes could be very large in most applications. Therefore rational expectations do not typically appear to be feasible for agents to learn from data and the optimal feasible expectations may be very different to rational expectations in the vast majority of macroeconomic settings. As a result I suggest optimal feasible expectations as a new class of expectations.

The precise size of the conditional biases in optimal feasible expectations in any particular setting is, however, ultimately an empirical question. I therefore empirically examine the size of the conditional biases in estimated optimal feasible expectations of US inflation. I do this by starting with a sensible auto-regressive benchmark and then consider the degree of shrinkage that it is optimal to apply to information on six important macroeconomic variables using pseudo out of sample forecast performance. The variables I consider are a combined business cycle indicator, a combined financial cycle indicator, a combined indicator of broad inflation, trade-weighted exchange rates, hourly wages and narrative monetary shocks. I find that it is optimal to apply a very high degree of shrinkage to the most of these

variables. Indeed it is optimal to apply absolute shrinkage to the business cycle indicator, the financial cycle indicator, exchange rates and wages, so this information is not included in the forecasts produced. It is also optimal to apply partial shrinkage to the broad inflation series but none to the narrative monetary shocks, although the shocks only have small associations with future inflation and the sign on these is not that implied by theory for true monetary shocks. However some of the macroeconomic variables having economically and statistically significant associations with future inflation, so the results imply that there are large conditional biases in optimal feasible expectations. Optimal feasible inflation expectations therefore appear to be very far from rational inflation expectations and are likely to contain large conditional biases that cause significant nominal rigidities in the reactions to macroeconomic shocks.

I also examine the conditional biases in surveys of US households' inflation forecasts and compare them to the conditional biases in estimated optimal feasible expectations. I find that household forecasts have a much smaller association with the broad inflation index than future realised inflation does and barely react to most of the other variables, although they do have a statistically significant but incorrect association with the financial cycle indicator. Therefore their conditional biases appear to be very similar to the conditional biases in optimal feasible expectations, except with regards to financial cycles. As a result the household inflation expectations are likely to produce similar, or even greater, levels of nominal rigidities in response to macroeconomic shocks than optimal feasible expectations would. However I also confirm that household expectations are not consistent with being entirely formed by optimal feasible expectations by showing that they are clearly beaten in a pseudo out of sample forecasting horse race by a feasible alternative: my auto-regressive benchmark. This may be caused by several persistent and seemingly unjustified deviations of household inflation forecasts from realised inflation that can last for years at a time as well as a mistaken reaction to financial cycle information. Therefore household forecasts of inflation expectations appear to be well explained by a combination of optimal feasible inflation expectations and behavioural mistakes that reduce forecast performance.

Optimal feasible expectations are therefore likely to cause important nominal rigidities, however they may also have important implications for many areas of economics that could be explored in interesting future research. For instance, they may imply that agents forecasts of future asset returns are conditionally biased towards the long-term average return of similar assets and hence suggest why agents fail to adjust asset demand to remove small associations between current variables and future risk-adjusted returns, providing an empirically optimal justification of bubbles. They may also imply that agents forecasts of the probabilities of being the pivotal voter between each combination of two plausible candidates in an election are biased towards a single probability in electoral systems that make these probabilities hard to predict, limiting the applicability of Arrow's impossibility theorem.

Optimal feasible expectations also strongly suggest that agents learning optimally from data will not use rational expectations in the vast majority of macroeconomic settings, contradicting one of the arguments often put forward for the use of rational expectations. Since agents' expectations in the real world may often be formed by a combination of optimal feasible expectations and behavioural deviations from optimal behaviour they will often contain large conditional biases, suggesting models based on rational expectations may be fundamentally mis-specified.

Chapter 2

Asset price convergence, international asset holdings and the quality of financial integration

2.1 Introduction

Financial integration matters. At its best it has the potential to channel capital to where it is most productive and enable greater diversification in asset holdings, but at its worst it can drive financial instability by enabling rapid capital flight after negative shocks. It is therefore highly important to understand changes in financial integration as well as to measure them.

The benefits of financial integration are usually suggested as occurring due to increases in the policy definition of increased financial integration, i.e. a removal of frictions on the basis of location, nationality or other irrelevant characteristics which affect agents' access to and investment of capital (Coere, 2013). These suggested benefits have their roots in the general microeconomic suggestion that the permanent removal of frictions from markets leads to more efficient market outcomes. Specifically, the removal of such frictions has been suggested to channel capital to where it is most productive and so lead to more efficient capital allocation, as well as to diversify investors' risk exposures and so increase risk sharing (Baele et al., 2004).

It is extremely hard to measure financial integration as defined in its policy definition. Certain frictions, such as current de jure capital controls, can be measured but many others, such as the probability of future exchange rate controls, cannot, so it is not generally possible to measure the level of frictions directly. It is also not generally possible to measure if access to and investment of capital varies due to differences in location, as there are many unobserved differences in assets and investors that are correlated with nationality and different countries have many unobserved underlying differences that are correlated with location.

Therefore, in practice, empirical measures of financial integration are made without controlling for underlying characteristics. The two most common types of measure are price based measures, based on the spread of returns on assets between different countries, and quantity based measures, based on the percentage of overseas assets in investors' portfolios. These measures capture concepts that are closer to the dictionary definition of financial integration, which is simply closer links between international financial markets. These measures of financial integration will capture differences in the underlying characteristics of different assets and aversion to these characteristics as well as the effect of international frictions. Therefore increases in these measures do not necessarily stem from reductions in financial frictions. They may instead stem from changes in underlying asset risks and aversion to them that vary over economic cycles and in response to shocks. As a result, increases in these measures that are not driven by reductions in financial frictions will not necessarily cause benefits from improved capital allocation or risk sharing.

However even changes in policy financial integration might also have economic costs. Stiglitz (2010) explains that, given there are other imperfections in financial markets relating to financial contagion, reducing international frictions may increase the probability of financial contagion and instability, causing economic harm from exposing countries to more severe financial shocks. This is a specific application of the microeconomic principle that removing one friction may exacerbate the effects of others. Changes in underlying asset risks and aversion to them that increases measured integration may also increase the probability of financial contagion and instability in a similar way, especially if the changes are cyclical and unsustainable.

These issues have led policymakers to begin discussing the quality of financial integration, i.e. whether integration is having the potential economic benefits and whether it is also avoids economic costs, as well as simply measuring the level of empirical financial integration indicators (European Central Bank, 2016; European Investment Bank, 2017). We therefore use the quality of financial integration to mean the extent to which financial integration is likely to have net economic benefits. The main aim of this chapter is to statistically analyse several features that the literature suggests are associated with the quality of financial integration in the case of the EU in the 21st Century. We focus on analysing three features.

Firstly we test whether increases in financial integration jointly affect both price and quantity indicators of integration, as discussed in European Investment Bank (2017). This is because reductions in financial frictions should lead to increases in both price and quantity integration and increases in both will be needed to obtain the full economic benefits: quantity integration to improve risk sharing and price integration to encourage new capital investment (Coeure, 2013). On the other hand, convergence in underlying asset risks may increase measured price integration, as risks are priced into assets, but may have little effect on either the efficiency of capital allocation or risk sharing.

Secondly we test the permanence of increases in financial integration, as discussed in European Investment Bank (2019). This is because temporary and unsustainable increases in financial integration create a danger of becoming risks to financial stability when they unwind, despite only having the potential to offer short-term benefits. On the other hand, permanent changes in financial integration will not, in themselves, create a financial instability risk and have the potential for perpetual risk sharing and capital efficiency benefits, the latter of which can often only be obtained over the long term as real capital formation can be a very slow process (European Investment Bank, 2017).

Thirdly we test how robust financial integration is to shocks which affect macroeconomic conditions, as discussed in European Central Bank (2018). This is because financial integration that is very vulnerable to shocks may decline dramatically after macroeconomic shocks and become an additional transmission mechanism for

the shock, whereas resilient financial integration will not. Additionally individuals and companies are most vulnerable after negative shocks, so it is important that financial integration enables risk sharing in these periods of vulnerability¹.

To perform this analysis we want to produce a financial integration index that captures the primary changes in financial integration and allows us to test the three features associated with the quality of financial integration. Constructing such an index is not econometrically easy. Approaches based on the cross-sectional means of indicators, such as those in Hoffmann et al. (2019), imply the result of the first test based on the weights chosen for price and quantity indicators in the construction of the index, so cannot be used to test whether there are joint movements in price and quantity integration measures. Factor approaches resolve this issue. However simple factor approaches, such as those based on the principal components analysis used in European Investment Bank (2017), are only valid with stationary data and so are constructed to give a stationary factor. Therefore they imply an answer to the second test, so cannot be used to test the permanence of changes in the financial integration index. Maximum likelihood factor approaches, such as those based on the Kalman filter, avoid these issues, but are computationally challenging given the number of parameters involved.

Therefore we use a Bayesian approach based on Markov chain Monte Carlo that is both theoretically valid and computationally feasible. Specifically we generate a financial integration factor in a Bayesian factor-augmented vector auto-regression (FAVAR) that also includes macroeconomic and financial variables. We then test whether this factor jointly drives changes in price and quantity integration by testing the signs of price and quantity factor loadings, whether the changes in the factor are permanent by analysing if there are deterministic or stochastic trend movements in the factor and whether financial integration is robust to macroeconomic shocks by using the results of the second test and testing the response of financial integration

¹Note that it is possible to share the risks of shocks that are large enough to have aggregate European effects, as well as idiosyncratic shocks. This is because large shocks have heterogeneous effects across different countries, which can be seen by considering the effect of the financial crisis on Ireland relative to its effect on Poland, or the effect of the Eurozone Debt crisis on Greece relative to its effect on the UK.

to sign identified shocks that affect macroeconomic conditions.

Applying this approach yields a financial integration factor for the EU that shows clear increases in the early 2000s, large but uneven falls in the latter 2000s and 2010s and has recently begun to recover, albeit at a relatively slow rate. The results of the first test indicate that both price and quantity measures of integration tends to load positively on this factor, so there are important joint movements in price and quantity integration. However the results of the second and third tests indicate that the movements in the factor are primarily cyclical and vulnerable to shocks that affect macroeconomic conditions. Therefore the primary changes in financial integration in Europe in the 20th Century do not appear to have been high quality changes that caused large net economic benefits. Indeed the changes in financial integration measures appear most related to cyclical changes in the underlying risks of European assets and aversion to these risks. This contrasts with the description of the changes in financial integration the 1990s accompanying the preparation for European Monetary Union and the enlargement of the EU provided by the existing literature (Lane, 2008).

This chapter builds on the existing empirical financial integration literature, particularly the part of the literature that focuses on financial integration in Europe. Several papers have produced empirical measures of financial integration in Europe based on the unconditional convergence of asset prices (Abascal et al., 2015) and/or unconditional increases in the proportion of internal assets held in portfolios (Hoffmann et al., 2019). These measures are constructed using the weighted cross-sectional means of underlying factors or principal component factors taken from the underlying series, or both, and are regularly updated by European institutions (European Central Bank, 2018; European Investment Bank, 2019)².

There have also been attempts to measure policy financial integration more directly in specific settings. Baele et al. (2004) attempts to measure price convergence conditional on the underlying characteristics of particular assets and Schindler (2009) measures the underlying frictions caused by de jure capital controls. Other

²Parts of the analysis from a preliminary version of this Chapter were referenced and repeated in European Investment Bank (2019).

strands of the literature focus on measuring potential benefits of financial integration, such as risk sharing, and correlating them with changes in financial integration. For instance Kalemli-Ozcan et al. (2014) take this approach to the financial crisis and Ferrari and Rogantini Picco (2017) take this approach to the adoption of the Euro. These approaches are potentially interesting in specific settings but cannot be used more generally, as typically many characteristics of assets, financial frictions and that affect the benefits of financial integration cannot be observed. Therefore policy papers typically limit themselves to discussing the extent to which they believe changes in empirical financial integration measures are driven by high quality changes in financial integration (European Central Bank, 2016; European Investment Bank, 2017).

This chapter contributes to this literature in two ways. Firstly we introduce a Bayesian methodology for testing three aspects of the quality of financial integration changes that are discussed in the existing policy literature: whether there are jointly driven movements in price and quantity financial integration, whether financial integration is long-term and whether financial integration is resilient to negative macroeconomic shocks. This approach yields a financial integration factor that is not automatically stationary and does not automatically load on particular indicators as a result of its construction. Secondly we apply this methodology to test the quality of financial integration changes in the EU in the 21st century. In doing so we also provide new stylised facts on EU financial integration. We find that changes in empirical financial integration are driven by joint changes in price and quantity integration, however we also find that they are primarily cyclical and the majority of their changes are vulnerable to shocks that affect macroeconomic conditions. Therefore increases in financial integration in the EU in the 21st Century do not appear to have been particularly high quality: they seem more likely to result from cyclical changes in the underlying risk of assets and aversion to them, rather than genuine reductions in international financial frictions.

The rest of the Chapter is organised as follows: Section 2.2 sets out our econometric methodology, Section 2.3 describes our dataset and presents the financial integration factor we produce, Section 2.4 presents the results of the three tests of

the quality of financial integration changes and Section 2.5 offers some concluding remarks.

2.2 Econometric methodology

The econometric setup we use for our analysis is a Bayesian FAVAR. We define y_t as the dataset at period t , which can be split into an auxiliary and a main dataset, so $y_t = (y_t^a \quad y_t^m)'$. y_t^a is an $n \times 1$ vector containing variables which contain information about financial integration at period t and y_t^m is an $(m-1) \times 1$ vector containing key macroeconomic and financial variables at period t . We define f_t as the unobserved level of financial integration at period t , $z_t = (f_t \quad y_t^m)'$ and $Z_t = (z_t' \cdots z_{t+1-L}')'$, where L is the maximum lag. The measurement equation is then given by:

$$y_t = \Gamma + \Lambda Z_t + U_t \quad (2.1)$$

$$\Gamma = \begin{pmatrix} \gamma \\ 0 \end{pmatrix}, \Lambda = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & I & 0 \end{pmatrix}, U_t = \begin{pmatrix} u_t \\ 0 \end{pmatrix}, u_t \sim N(0, \omega), \Omega = \begin{pmatrix} \omega & 0 \\ 0 & 0 \end{pmatrix}$$

where γ is a vector of constants, λ is a vector of factor loadings and u_t s a set of idiosyncratic errors, so ω is diagonal and Ω is the measurement equation covariance matrix. Note that one factor loading will have to be set to one to ensure uniqueness of the factor, however with one factor this simply becomes a scaling constant.

The state equation is given by:

$$Z_t = \Theta + \Phi Z_{t-1} + V_t \quad (2.2)$$

$$\Theta = \begin{pmatrix} \theta \\ 0 \end{pmatrix}, \Phi = \begin{pmatrix} \phi \\ I & 0 \end{pmatrix}, V_t = \begin{pmatrix} v_t \\ 0 \end{pmatrix}, v_t \sim N(0, \sigma), \Sigma = \begin{pmatrix} \sigma & 0 \\ 0 & 0 \end{pmatrix}$$

where θ is a vector of constants, ϕ is a matrix of VAR coefficients and v_t s a set of reduced form VAR errors so σ is a non-diagonal matrix and Σ is the state equation covariance matrix.

The structural errors are given by:

$$\eta_t = chol(\sigma)Qv_t = Dv_t \quad (2.3)$$

$$\eta_t \sim N(0, I)$$

where $chol$ = cholesky decomposition, Q is an orthogonal rotation matrix and η_t is a vector of structural errors.

Minnesota style priors are used for Θ and Φ so their prior distribution is jointly multivariate normal, with a mean which implies a univariate random walk for each variable and a diagonal variance-covariance matrix with hyperparameters similar to typical values in the VAR literature reported by Canova (2007). These are a hyperparameter of 0.2 on own variable lags, a hyperparameter of 1 on other variable lags, a hyperparameter of 2 on lags greater than 1 and a hyperparameter of 100 on constants. The prior distribution for σ is inverse-wishart, with a mean set by taking the covariance of the residuals from random walks for each variable and a very low degrees of freedom: 9, which is one more than the dimension of sigma.

The prior distribution for Γ and Λ is jointly multivariate normal. The mean is set by regressing each indicator on the cross sectional mean of the indicators after normalising so the identification condition is met. The variance-covariance matrix is diagonal and set by bootstrapping the above regression, to allow for the possibility that financial integration is non-stationary. The prior distribution for each diagonal element of ω is inverse gamma, with a mean set as the variance of the residuals from each of the above regressions and a very low degrees of freedom: 2. The prior for the initial lagged values of the financial integration factor, F_0 , is normally distributed with the mean set as the value of the lagged cross sectional means of the indicators after normalising so the identification condition is met and the variance-covariance matrix diagonal with values given by the cross-sectional variance of the normalised mean at each lagged value.

These independent conjugate priors ensure that the conditional distribution of each group of parameters has a known distribution. Therefore one can construct a Gibbs sampler where at each step we draw one group of parameters from their

posterior distribution, conditional on all parameters not in their group. The steps are as follows: draw f from its multivariate normal distribution using the algorithm of Carter and Kohn (1994), draw Θ and Φ from their multivariate normal distributions, draw σ from its inverse-wishart distribution, draw Γ and Λ from their multivariate normal distribution and draw each element of γ from its inverse-gamma distribution. See Appendix C for details of this procedure. We set the lag length to two: since we work with variables in levels this ensures that the model can capture cyclical behaviour in a parsimonious way.

We initialise the sampler using the prior means of the parameters and take 30,000 draws from this distribution then burn the first 10,000. We also then discard those draws which give explosive estimates³. This leaves sufficient draws for excellent convergence of the statistics, which is confirmed in Appendix C by showing that the same results are obtained using an arbitrary initialisation and a greater number of draws.

We can then test the three aspects of the quality of changes in the financial integration factor as follows. We test whether there are joint movements of price and quantity financial integration by testing whether price and quantity indicators load on the financial integration factor with the same sign. Specifically, since there are many indicators, we compare the sign of the mean factor loading on the price indicators with the sign of the mean factor loading on the quantity indicators. We also compare the mean factor loading on the price indicators within each market to the mean factor loading on the quantity indicator within that market. If the signs of the mean factor loadings on price indicators and quantity indicators are the same then this suggests that there is an important joint driver. However if the mean factor loadings on price indicators and quantity indicators either have opposite signs or only one of them is non-zero then this suggests that there is not an important joint driver.

We test whether the movements in the financial integration factor have been driven by permanent components as follows. We take the value of Z in a period p as

³This implies that we are actually using restricted versions of the above priors which place zero weight on explosive solutions.

given, $\tilde{Z}_p = Z_p$, and then iterate forwards but only using lagged variables and linear trend components, i.e. without including future shocks: $\tilde{Z}_t = \Theta + \Phi\tilde{Z}_{t-1}$. This gives a series \tilde{Z} which only includes the effects of shocks which occurred up to period p . Since we ensure that the estimated parameters are such that the time series they generate are either $I(1)$ or $I(0)$, the changes in these series ultimately converges, so the series themselves converge to grow along a linear trend⁴ for any plausible parameter values. We can therefore test whether there is a positive linear trend, i.e. deterministic trend growth in financial integration, and whether any shocks cause permanent increases in the level of the linear trend, i.e. stochastic trend growth. Either of these implies permanent increases in the financial integration factor.

Table 2.1: Sign classifications

	Financial integration	Output	Other variables
Joint shocks	-	-	
Separating shocks	+	-	

Notes: Sign classifications used to classify the reduced form shocks in Equation 2.3 into two types of structural shock, based on their effects on the variables on impact. There are eight variables in the FAVAR, so in each Gibbs draw there must be more than one of at least one type of shock.

We test how vulnerable financial integration is to shocks that affect macroeconomic conditions as follows. First we note that any deterministic trend changes in financial integration indicated by the second test will be independent of macroeconomic shocks, so part of this test is conducted in the test above. However we also test if other movements in financial integration are vulnerable to shocks that affect macroeconomic conditions. We begin by calculating the correlation between the annual changes in financial integration and output caused by shocks in the sample.

⁴Where the variables are included in log form this implies that the underlying variable converges to grow exponentially.

If this correlation is positive and large then this implies that financial integration is generally high after shocks that cause positive macroeconomic conditions and vice versa. However, since the overall correlation might be driven partly by the response of macroeconomic conditions to changes that increase policy financial integration, we also test if there is a meaningful subset of shocks to macroeconomic conditions to which financial integration is not vulnerable. We do this by splitting the shocks into those that are very likely and those which are not very likely to produce highly positively correlated effects on financial integration and output. We identify the first category as shocks that move output and financial integration in the same direction on impact, which we call ‘joint shocks’, and the second category as shocks that move output and financial integration in opposite directions on impact, which we call ‘separating shocks’. We can then test the correlations between the annual changes in financial integration and output caused by each type of shock and test if both are meaningfully positive.

We obtain the sign-restricted structural form using the approach of Rubio-Ramírez et al. (2010)⁵. This usually involves repeatedly drawing Q matrices using the qr decomposition of a normal matrix for each Gibbs draw until the instantaneous responses satisfy the sign restrictions. However here the sign restrictions are automatically met so only one Q matrix is needed. We can then test the correlation between the non-trend changes in financial integration and output caused by separating shocks. If even the correlation produced by separating shocks is positive then this implies that the vast majority of increases in financial integration occur with improvements in macroeconomic conditions.

⁵This implies a prior on Q which is not diffuse for structural analysis, as Baumeister and Hamilton (2015) show, however it is technically agnostic in the sense that it only depends on the reduced form parameters, as Kilian and Helmut (2017) discuss.

2.3 Dataset and estimated financial integration factor

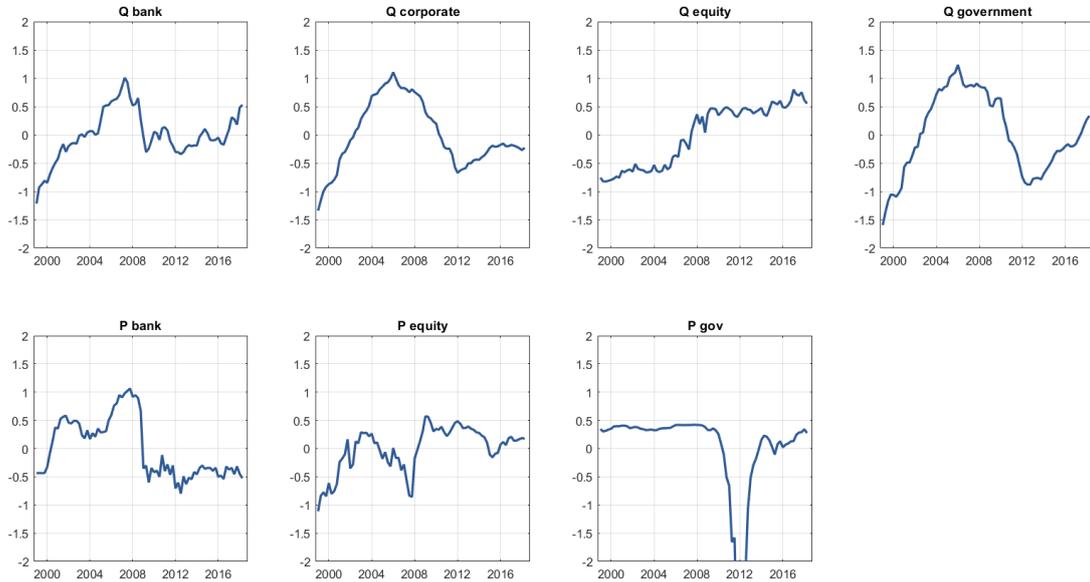
The dataset we use covers the first quarter of 1999⁶ to the second quarter of 2018 and consists of an auxiliary dataset and a main dataset. The auxiliary dataset includes both price (asset price convergence) and quantity (cross-border asset holding) measures of financial integration and was constructed at the European Investment Bank (EIB) based on data from European Union institutions. The countries covered are Germany, France, Italy, Spain, Belgium, Netherlands, Austria, Portugal, Ireland, Greece, Finland, UK, Sweden, Denmark, Poland, Czech Republic and Hungary. The data included per country consists of up to one bank quantity series, one corporate bond quantity series, one equity quantity series, one government bond quantity series, two bank lending price series (one to households and one to businesses), two equity price series (one of financial businesses and one of non-financial businesses) and one government bond price series. However not all series are available for all countries, so in total the dataset consists of 11 bank quantity series, 11 corporate bond quantity series, 11 equity quantity series, 11 government bond quantity series, 16 bank price series, 28 equity price series and 14 government bond price series⁷. This gives 44 quantity series and 58 price series, so a total of 102 series.

The quantity series are the shares of non-domestic bank debt, corporate debt, government debt and equity, held by domestic monetary financial institutions in different EU countries, measured monthly. We then process these by taking the three month moving average of the underlying series to transform them to quarterly frequency, then subtract the mean and divide by the standard deviation of each series to make them more easily comparable and give the resulting indicators. The price series are based on the price to book ratio for the equity series and the implied interest rate for other series. Each underlying series is the negative of the

⁶There is limited data available before the is period, however a longer but much less rich dataset is considered in Appendix C. Extending the dataset forward into the 2020s once data is available on the effects of Brexit and Covid-19 would be an interesting area for future research.

⁷Corporate bond price series cannot be used as there is no available data source on regular corporate bond rates in the EU over the sample.

Figure 2.1: Mean auxiliary indicators

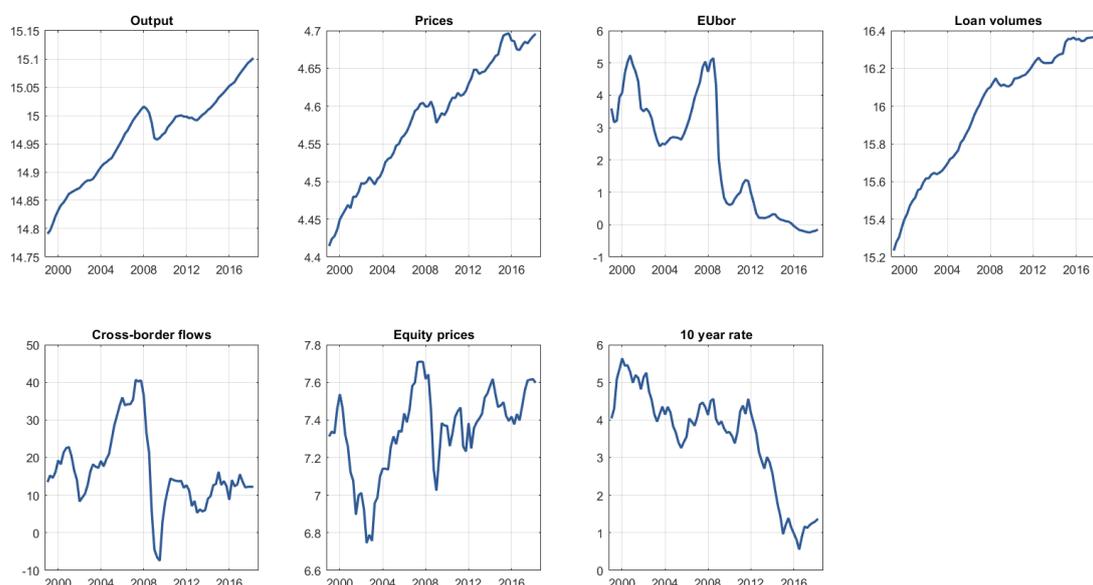


Notes: Plots of the mean values of the transformed auxiliary indicators in each category: see the text for the definitions of the auxiliary indicators. The vertical axes are in units and the horizontal axes are in years. Q denotes quantity indicators, P denotes price indicators, bank denotes the bank lending market, corp denotes the corporate debt market, eq denotes the equity market and gov denotes the government bond market.

squared deviation of the ratio or rate from its cross-sectional average. The squared deviation captures dispersion and we use the negative so a higher value corresponds to more financial integration. We then also process these series by taking the three month moving average of monthly underlying series to transform them to quarterly frequency, then subtract the mean and divide by standard deviation to make their more easily comparable and give the resulting indicators. The cross-sectional mean value of the indicators in each market segment is plotted in Figure 2.1.

The main dataset is constructed from the following sources and is plotted in Figure 2.2. Output is the seasonal and calendar adjusted chain-linked GDP volume for the EU 28 in log form from Eurostat. The price level is calculated by dividing the seasonal and calendar adjusted GDP volume by the equivalent figure in euros for

Figure 2.2: Main variables



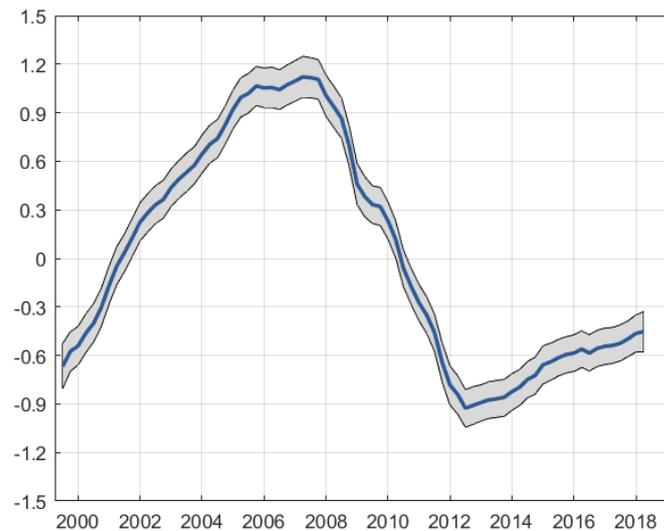
Notes: Plots of the values of each of the main variables included directly in the VAR. The horizontal axes are in years and the vertical axes are in various units: see the text for details of the variables and their units.

the EU 28 from Eurostat and is in log form. The weighted 3 month rate (‘EUbor’) is constructed as 80% the 3 month Euro libor rate and 20% the 3 month GDP libor rate from IBE bench-marking⁸. Loan volumes are the loan liabilities of non-financial corporations in EU28 countries that have data available over the whole sample from Eurostat. It is deflated with the GDP deflator and in log form. Cross-border gross flows are the sum of financial inflows and financial outflows from the financial account for the EU28. It is based on EIB internal calculations from IMF data and is expressed as a percentage of GDP. The equity price index is the weighted average of non-financial equity indices for EU28 countries. It is based on EIB internal calculations from Thomson-Reuters data. It is deflated with the GDP deflator and in log form. The 10 year rate index is the weighted average of 10 year government bonds for EU28 countries. It is based on ECB internal calculations with ECB data.

⁸These proportions are approximately equivalent to the size of the relative economic areas.

It is not generally possible to perform robustness checks using different data sources, as there are far fewer data series available for the whole EU than there are for a typical developed country, however many of these series are widely used across the EU institutions.

Figure 2.3: Financial integration factor



Notes: Plot of the posterior median of the baseline financial integration factor, with the area between the 5th and 95th posterior percentiles shaded. The vertical axis is in units and the horizontal axis is in years. The factor is normalised such that a one unit change implies a one standard deviation change in the first auxiliary series, which is an approximately 7.7 percentage point increase in the Austrian banking quantity series.

The posterior median of the estimated financial integration factor is plotted and the area between the 5th and 95th percentiles shaded in Figure 2.3. The factor is defined so that an increase implies greater integration and is normalised to the banking quantity indicator for Austria, the first series in the auxiliary dataset, however normalising to another variable would simply scale the indicator by a different constant. Therefore a one unit increase in the factor causes a one standard deviation, which is a 7.7 percentage point, increase in the share of non-domestic bank loans held by Austrian monetary and financial institutions. This value is not unusual. A

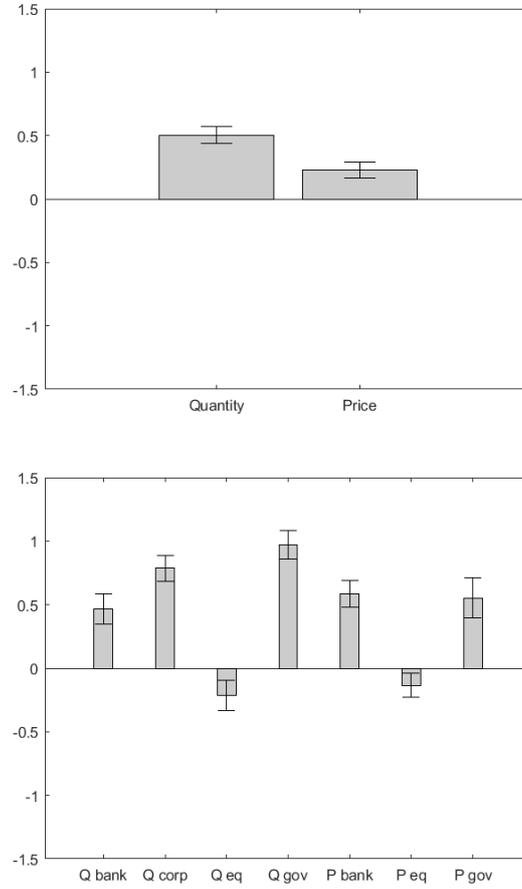
one unit increase in the factor also increases the share of non-domestic bank loans held by monetary and financial institutions in most countries by between 5 and 15 percentage points.

The factor broadly rose through the early 2000s, before approximately plateauing around 2006 and then starting to decline early in the financial crisis. This decline continued unevenly for several years: it was rapid just after the financial crisis but then appeared to be slowing before becoming rapid again in the period of the Eurozone debt crisis. Since 2013 there has been a slow and slightly uneven period of recovery. The decline during the period associated with the financial crisis and sovereign debt crisis is very large relative to the gains before and after, so that financial integration in the EU at the end of the sample is only just higher than its level in 2000. The movements of the factor and its relationship with the underlying indicators and the macroeconomic series are discussed in much more detail in Section 2.4 when the tests of quality are conducted. The percentile bands are also usually fairly tight around the factor, showing that it is reasonably well observed but that there is some uncertainty in the level of the factor in parts of the sample.

2.4 The quality of EU financial integration

We now turn to conducting the three tests of the quality of financial integration in the EU in the 21st Century. We begin by testing whether there is a joint driver of price and quantity financial integration indicators by studying whether price and quantity indicators load with the same sign on the financial integration factor. To do this we plot the posterior median, 5th percentile and 95th percentile of the mean factor loading on price indicators and the mean factor loading on quantity indicators in the upper panel of Figure 2.4. To break down these results further we also plot the posterior median, 5th percentile and 95th percentile of the mean factor loading on price and quantity indicators in each of the market segments in the lower panel of Figure 2.4. This ensures that the results are not just driven by a few large factor loadings and lets us compare the signs of the factor loadings on price and quantity indicators within each market.

Figure 2.4: Mean factor loadings



Notes: The upper panel plots the posterior median of the mean factor loading on price and quantity indicators with error bands showing the 5th and 95th posterior percentiles. The lower panel plots the posterior median of the mean factor loading in each market category with error bands showing the 5th and 95th posterior percentiles. The vertical axis is in units and the loading on the first indicator, the Austrian banking quantity indicator, is constrained to be one. Q denotes quantity indicators, P denotes price indicators, bank denotes the bank debt market, corp denotes the corporate debt market, eq denotes the equity market and gov denotes the government bond market.

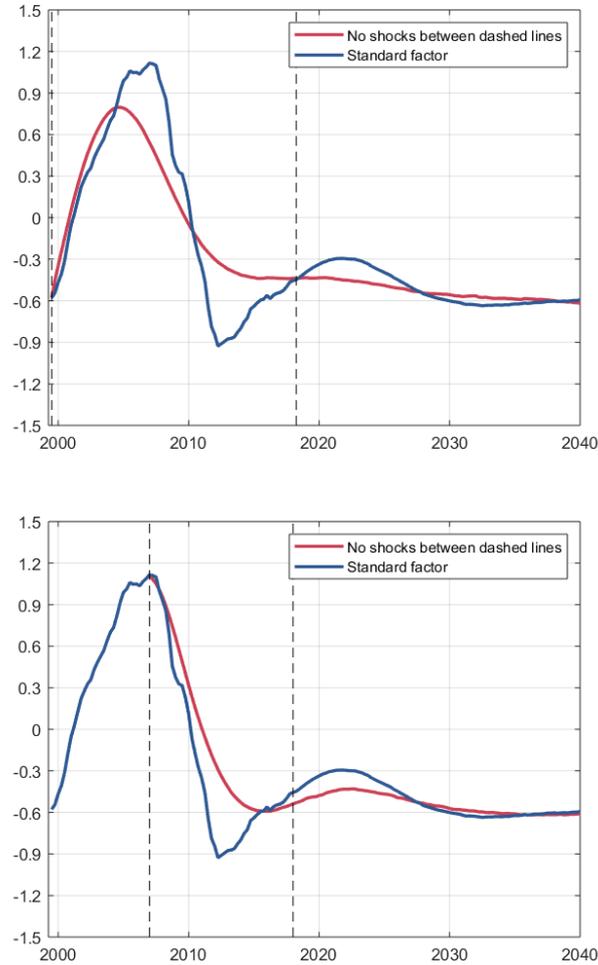
The mean factor loadings on price and quantity indicators are both clearly positive, with even their 5th percentiles above zero. This is not just driven by a small subset of the average factor loadings, as the mean factor loadings by market seg-

ment show. The mean factor loadings on quantity indicators are strongly positive in three out of four market segments and the mean factor loadings on price indicators are strongly positive in two out of three market segments. The factor also appears to load slightly more positively on quantity indicators than price indicators. The only market segments in which the mean factor loading is negative are the equity quantity and equity price segments, although in both cases the size of the mean negative factor loading is small. Therefore this implies that the results of the first test of the quality of financial integration in the EU are broadly positive: there are strong joint drivers of price and quantity measures of financial integration. These results also show that the financial integration factor captures most developments in EU financial integration well in a parsimonious manner. However it is important to remember that the factor does not capture most improvements in equity market integration, as changes in integration in this market appear to be statistically distinct, or in some cases even opposite, from other changes in EU financial integration.

We now move on to testing the permanence of changes in financial integration. We test whether there is a positive linear trend in financial integration and whether shocks in the sample cause permanent increases in the level of the linear trend to which the financial integration factor converges to. To do this we calculate the impact of these shocks by comparing the linear trends to which the standard series and the series cleaned of shocks, as described in Section 2.2, converge to. The series cleaned of different sets of shocks are plotted in Figure 2.5, with the uncleaned shocks always plotted for comparison. The sets of shocks we strip out include all shocks in the sample and all shocks from 2007 onward. The latter date was chosen as it is around the period where one might intuitively expect a financial cycle to peak.

The empirical results are very clear, although their interpretation is less absolute. All the series converge to have a growth rate that is close to zero and have reasonable posterior probabilities of being either positive or negative. Therefore there does not appear to be any substantial deterministic trend growth in the financial integration factor. The two series with different sets of shocks removed also converge to a level which is extremely similar to, and with reasonable posterior probabilities of being

Figure 2.5: Extrapolated financial integration with different sets of shocks



Notes: Plots of the posterior median of the estimated financial integration factor during the sample and the extrapolated financial integration factor after the sample. Both panels show the standard case and the case where shocks between the dashed lines have been removed. In the upper panel this corresponds to all shocks in the sample and in the lower panel this corresponds to all shocks after 2007. The vertical axes are in units and the horizontal axes are in years. See Figure 2.3 and the accompanying text for the interpretation of the factor's units.

either above or below, the level to which the standard factor converges. Therefore there does not appear to be any substantial stochastic trend growth in the financial integration factor either. As a result there are virtually no permanent increases

in the financial integration factor in the sample. This suggests that there were not large permanent changes in financial integration in the sample. However one cannot conclude that there were no permanent changes in financial integration in the sample, as the factor either does not load or loads with negative signs on the vast majority of equity integration indicators. Since Figure 2.1 shows that there appear to be positive trends in this series this suggests that there may have been some trend increases in financial integration in the sample, but that these were small in comparison to the scale of the cyclical changes. Therefore the results of the second test of the quality of financial integration in the EU are negative, but not necessarily totally negative.

It is also interesting that there appear to be long-lasting cyclical effects in financial integration. Our evidence suggests that pre-sample shocks explain a meaningful proportion of the build-up and decline in financial integration in the 2000s. Also shocks before 2008 can explain a reasonably large proportion of the subsequent decline and recovery of financial integration up to the end of the sample. These long-term cycles appear to be the equivalent for financial integration of the long-term cycles in credit and other macrofinancial variables that have previously been reported in the literature (Borio, 2014). The conclusion that long-lasting cyclical effects, rather than permanent changes, explain most of the movements of European financial integration since 2000 also contrasts with the description of the events of the 1990s, the period in which much of the activity accompanying the process of European Monetary Union was undertaken, given by Lane (2008).

We now move on to testing whether financial integration is robust to shocks that affect macroeconomic conditions. Part of this test is contained in the test of the permanence of changes in financial integration, as any deterministic trend changes in financial integration will by definition be robust to shocks. However any deterministic trend changes in the sample are relatively small compared to other changes, as explained in the discussion of the results of the previous test, so we can focus on testing the robustness of other changes in financial integration. We calculate the correlation between the annual changes in financial integration and output caused by all shocks and the correlations caused just by either joint shocks or

separating shocks, which are particularly likely and unlikely respectively to produce changes in financial integration that are positively correlated with changes in output. The results are displayed in Table 2.2. Impulse response functions (IRFs) for the two types of shocks are also available in Appendix C, however these are not very informative in this setting. This is because we classify the eight structural shocks into two types of shocks, rather than into individual specific shocks, so the IRFs do not necessarily represent how the effects of specific shocks on different variables are related.

Table 2.2: Correlations between annual changes in financial integration and output driven by shocks

Shocks	Correlation
All	0.53 (0.30 to 0.72)
Joint shocks	0.62 (0.19 to 0.86)
Separating shocks	0.34 (-0.32 to 0.75)

Notes: Posterior median correlations between the annual changes in financial integration and output around their initial predicted paths. Row 1 shows the correlations driven by all shocks, Row 2 shows the correlations driven by joint shocks only and Row 3 shows the correlations driven by separating shocks only. Posterior 5% and 95% percentiles are also shown.

The results show positive correlations between the changes in financial integration and output in all three cases. The correlation coefficient of 0.53 resulting from

all shocks is high for macroeconomic series in growth rates and the correlation coefficient of 0.34 resulting from only separating shocks is still relatively high given that the shocks are defined to try to capture potentially negative correlations. Figure 2.3 also shows that changes in the financial integration factor correspond to fairly large changes in the underlying indicators. This suggests that the vast majority of shocks to output also cause meaningful corresponding changes in the financial integration factor. Therefore financial integration in the sample does not appear to have been very robust to shocks that affect macroeconomic conditions. However this conclusion cannot be absolute, as the financial integration factor does not capture some aspects of equity market integration, which may have had a different level of resilience to shocks. Hence the results of the third test are negative, but not necessarily totally negative.

Therefore, the result of the first test is fairly positive while the results of the second and third tests are fairly negative, although none of the three results is absolute. This suggests that there is a joint driver of most price and quantity indicators of financial integration but it is primarily cyclical and vulnerable to shocks that affect macroeconomic conditions. Thus it appears that most of the changes in empirical financial integration measures are not high quality. Cyclical changes in the underlying risk of assets or aversion to them, such as the behavioural counter-cyclical risk aversion demonstrated in Cohn et al. (2015), seem the most likely drivers of such financial integration.

2.5 Conclusion

Financial integration that results from reductions in international financial frictions has the potential to improve capital allocation and risk sharing, but may also increase the possibility of financial contagion and instability. Measuring this kind of integration is also challenging, as empirical financial integration measures also capture changes in the underlying risks of assets and aversion to them that will not necessarily have risk sharing and capital allocation benefits. It is therefore very important to assess the quality of changes in financial integration, i.e. how likely the

changes are to cause net economic benefits, as well as simply measuring them. In this chapter we statistically analyse three aspects of the quality of changes in financial integration indicators that have been discussed in the policy literature. These are whether there is a joint driver of changes in price and quantity integration, whether the changes in financial integration are permanent and whether financial integration is robust to macroeconomic shocks.

We suggest a new methodology to overcome the econometric difficulties associated with producing an index of financial integration from many non-stationary financial integration indicators and conducting the three tests. Our methodology, which is based on a Bayesian FAVAR, allows us to produce a non-stationary financial integration factor. It also lets us conduct the three tests by testing the signs of the factor loadings on price and quantity indicators, testing whether there are important deterministic or stochastic trend movements in the factor and testing the correlations between changes in the factor and changes in output driven by structural shocks identified with sign restrictions.

We then apply this methodology to the case of financial integration in the EU from 1999 to 2019. The results suggest that there is an important common factor that has jointly driven price and quantity indicators of financial integration in the EU in this period. This common driver loads strongly on the vast majority of banking, corporate bond and government bond indicators, but not most equity indicators. However financial integration has been very cyclical: any permanent changes have been small relative to long-term cyclical changes. It has also been vulnerable to shocks that affect macroeconomic conditions, as virtually all shocks to macroeconomic conditions also cause important corresponding effects in financial integration. It is possible, however, that there has been some limited robust and permanent financial integration, particularly in European equity markets.

We conclude that most changes in EU financial integration in the 21st Century thus far do not appear to have been of particularly high quality, so may not have produced large net economic benefits. Instead, they seem more likely to have been primarily driven by cyclical changes in the underlying risks of European assets and aversion to these risks.

Chapter 3

Behavioural finance at home: house price cycles in the USA

3.1 Introduction

Housing is the largest asset on most households' balance sheets, as well as one of the largest items of consumption, in many developed countries (Piazzesi and Schneider, 2016; Musso et al., 2011), so it is important to understand house price¹ movements in themselves. However the cyclical movements in house prices have also been shown to be intricately linked to financial cycles (Borio, 2014) and their downturns are associated with some of the most serious recessions (Jorda et al., 2015b), so it is also important to understand them because of their wider implications.

One key consideration is whether changes in house prices are the result of changes in rational expectations of fundamentals, because otherwise this implies the potential for housing bubbles that could cause serious macroeconomic shocks when they burst (Glaeser and Nathanson, 2015). There are very different approaches to the issue of housing bubbles in the existing literature. The vast majority of theoretical macroeconomic models that incorporate housing surveyed by Piazzesi and Schneider (2016) do so in settings in which agents have full information rational expectations²

¹Throughout this chapter and in Appendix D house prices, housing values, housing rents and other related variables are all considered in real terms unless stated otherwise.

²Full information rational expectations are defined as agents expectations being equal to the true

and there are few frictions in housing markets. As a result house prices in the models surveyed are equal to the full information rational expectations of the present value of the stream of future rents. This is referred to as the fundamental value of a housing asset³. Therefore changes in house prices in these models are always due to changes in the fundamental value of housing. In reality there are frictions in housing markets, such as search costs, and it is possible to extend this approach to construct an adjusted fundamental value that also incorporates these features, as shown by Dusha and Janiak (2018). However there is also a large behavioural housing literature, surveyed in Salzman and Zwinkels (2017), that suggests that house prices may deviate from fundamental values, even after adjusting for transaction costs and frictions. These deviations may result from behavioural features such as non-rational expectations, so changes in house prices need not result from changes in the fundamental value of housing.

Providing empirical evidence on whether house prices are equal to their fundamental value is very hard. Testing the proposition directly is limited by the impossibility of estimating the fundamental value of housing without making extreme approximations, such as those in Mayer and Sinai (2007). Testing the proposition indirectly is more plausible but has its own problems. If asset prices are equal to their standardly defined fundamental value then the asset market is efficient, so tests of efficiency can be used as tests for consistency with asset prices being given by their fundamental value (Fama, 1970). Even so, testing market efficiency in aggregate house prices remains difficult. Aggregate house price series are not available at high frequencies, so it is hard to rule out the possibility that time series and event study predictability of house prices and excess housing returns reflect changes in risk premia rather than deviations from efficiency. Additionally after adjusting for transaction costs and rigidities in housing markets the adjusted fundamental value relation no longer necessarily implies market efficiency. These rigidities include the time spent searching for a house, property transaction taxes and the institutional conditional expectations of a future variable, conditioning on complete knowledge of the economy: see the survey in Coibion et al. (2018) for a detailed discussion.

³This definition can be taken from Santos and Woodford (1997) under the condition that there are no ‘rational bubbles’ or from Glaeser and Nathanson (2015) who use it in a housing context.

feature of US housing markets that prices are often committed to before exchange takes place. Therefore a lack of efficiency in housing markets does not necessarily show an inconsistency with house prices being equal to their adjusted fundamental value. The test is based on the reaction of house prices to monetary shocks at time horizons shortly after the shock, but long enough to allow for contractual rigidities. Monetary shocks are relatively unique as a variable in that the adjusted fundamental value of housing should have an unambiguously signed reaction to monetary shocks as soon as the contractual rigidities in housing markets no longer bind. I illustrate why this is likely to be the case using two conceptual frameworks that build on the fundamental value of housing: one that also considers the roles of consumption demand and housing supply and one that also considers the role of search frictions, as well as relevant empirical work on the transmission of monetary shocks. This is a test of consistency with house prices equalling their adjusted fundamental value, as it is possible that house prices could still react to monetary shocks as soon as contractual rigidities no longer bind but still not be given by their fundamental value. However if house prices do not react at the specified horizon then this rejects house prices being always equal to their fundamental value, but is consistent with several alternate behavioural explanations of house prices.

I implement this test using narrative shocks in a local projections specification. This lets me obtain potentially consistent estimates of the timing of the effects of the shocks, which could be significantly distorted in an auto-regressive model if the model is mis-specified. The narrative shocks I use are based on the natural experiment approach in Romer and Romer (2004), however I build on this approach by additionally controlling for information related to housing and financial markets to ensure that the estimates are as accurate as possible. The results show that house prices barely react to monetary shocks at a horizon of one to two months, the length of contractual rigidities, but have clearly statistically and economically significant reactions at horizons greater than a year. This provides strong evidence against house prices always being equal to the fundamental value of housing, even adjusting for the rigidities in housing markets. These results could be explained by imperfect information, for instance agents in housing markets not observing macroe-

conomic shocks perfectly, or behavioural features such as non-rational expectations, for instance agents in housing markets not understanding the functioning of housing markets perfectly. I suggest that both of these features are likely to be important.

As a secondary piece of work I use an approach inspired by the conceptual frameworks to provide estimates of the relative importance of the consumption demand, asset demand and supply channels in driving US house price cycles. I use band-pass filters to capture the cycle in US housing variables and then use sign analysis based on the linear and rank correlations between the cycles in different series to assess the relative importance of different channels. I find that consumption demand is the most important channel. Asset demand also appears to be relatively important, particularly in some time periods, whereas the supply channel appears to be by far the least important of the three.

Overall, I conclude that aggregate house prices do not appear to be always equal to their fundamental value. Limited information on shocks, non-rational expectations and behavioural features may all be important features of housing markets. This suggests that macroeconomic models that are based on full information rational expectations may be seriously mis-specified. It also suggests that changes in consumption and asset demand, which are the most important proximal drivers of house price cycles, may not have their origins in changes in fundamentals, so investors and policymakers should view them with caution.

This chapter is primarily related to two main strands of the existing literature. The first of these is the empirical literature on whether housing markets can be explained as the full information rational expectation net present value of a stream of rents and hence whether housing markets are efficient. Many papers surveyed in Ghysels et al. (2013) show strong predictability in house prices and excess returns to housing, so the predictability does not just stem from movements in rents or discount rates. This predictability has been known since at least Case and Shiller (1989) and has been confirmed with more modern econometric methods for variables including past returns (Schindler, 2013), valuation ratios (Campbell et al., 2009) and housing wealth to income ratios (Balcilar et al., 2019). This predictability has also been confirmed with respect to specific policy changes in event study approaches, such as

that in Jung and Lee (2017). The predictability also gives rise to clear cycles in house prices, the stylised facts of which are summarised with a US focus in Sinai (2015), in the context of financial cycles in Drehmann et al. (2012) and in the context of bubbles in Glaeser and Nathanson (2015). It is hard to rule out completely time-varying risk premia as an explanation of these results, but the strength and patterns of the predictability lead the authors of most of these studies to conclude that housing markets are inefficient, so deviate from their fundamental values. In a survey of the literature Glaeser and Nathanson (2015) suggest that deviations of house prices from their fundamental values could theoretically be explained by search and institutional frictions in housing markets, but the size of deviations leads them to conclude that additional factors like non-rationality are also likely to be needed to explain the deviations. However there do not appear to have been attempts in this literature to empirically test whether housing markets deviate from their fundamental values after taking into account search costs and contractual rigidities.

I primarily contribute to this literature by introducing an empirical method to study whether house prices are consistent with always being equal to their adjusted fundamental value, instead of just their fundamental value, which is also robust to changes in time-varying risk premia. Specifically I study whether there is a reaction of house prices to monetary shocks as soon as contractual rigidities no longer bind, as these shocks should have clearly signed effects on adjusted fundamental values at this horizon. This approach is in the spirit of event studies but focuses on whether there is a reaction of house prices as soon as prices can respond to the event, rather than focusing on whether there are lagged reactions which may result from changes in risk premia or search costs. This is a test for consistency with adjusted fundamental values, so positive results can theoretically prove that house prices are not always equal to their adjusted fundamental value, but negative results cannot prove that they are always equal to their adjusted fundamental value. I implement this test using narrative monetary shocks in a local projections approach to provide clear empirical evidence on the timing of the reaction of house prices to monetary shocks. The results show virtually no reaction of house prices to monetary shocks at the horizons when contractual rigidities stop binding, despite statistically and

economically significant reactions with sensible signs at much longer horizons. This strongly suggests that house prices are not equal to their adjusted fundamental value, so housing market inefficiencies are not simply the result of search frictions and contractual rigidities.

I make an additional contribution by providing new stylised facts on house price cycles. Specifically I provide empirical estimates of the relative importance of the consumption demand, asset demand and supply channels in driving US house price cycles, on the basis of a sign decomposition suggested by the conceptual frameworks. I find that the consumption demand channel is the most important and the supply channel is by far the least important.

The second strand of work that this chapter is related to is the empirical monetary shocks literature. This literature aims to estimate the effects of monetary policy shocks on a number of different macroeconomic variables, of which house prices are usually just one. Individual papers take different approaches to this. Older papers, such as Fratantoni and Schuh (2003), tended to use recursive restrictions to identify monetary shocks in auto-regressive models. Concerns over the identification of shocks using this method then led to a series of papers which either augmented or replaced these zero restrictions with sign restrictions (Del Negro and Otrok, 2007; Jarociński and Smets, 2008; Musso et al., 2011). Most recent papers have used narrative shocks, which can be used in either autoregressive models (Miranda-Agrippino, 2016), local projections (Jorda et al., 2015a) or both (Coibion et al., 2017). The broad conclusion of these empirical papers is that expansionary monetary shocks increase house prices, although these increases often only occur with a lag and the extent and significance of the increases are fairly variable between papers⁴. These results are also broadly consistent with empirical partial equilibrium estimates based on regulatory changes, such as Bhutta and Ringo (2017), and surveys, such as Fuster and Zafar (2015). This is because these approaches suggest that changes in interest rates cause little reaction in house prices at short horizons but may ultimately increase house prices as the quantity of purchases and willingness to pay for housing

⁴Results from the literature are also compared to the equivalent empirical results produced in this chapter in Section 3.4.

rises.

I contribute to this literature in two ways. Firstly I contribute by explicitly studying the response of house prices to a monetary shock at the horizon which corresponds to contractual rigidities no longer binding, as well as producing an impulse response function (IRF) in the style of the existing literature. Secondly I also consider the addition of controls that are chosen to be specifically relevant to housing in the generation of my narrative monetary shocks and the estimation of the shocks effects on house prices. The IRFs produced including these controls are fairly similar to those produced without the controls and those from the existing literature. Therefore these results also allow one to have greater confidence in the estimated effects of monetary shocks on house prices from the existing literature, even when this literature focuses on estimating the effects of monetary shocks on a wide range of variables.

The rest of this Chapter is laid out as follows: Section 3.2 sets out how I use monetary shocks to test whether house prices are consistent with always being equal to the adjusted fundamental value of housing, Section 3.3 presents the results of this test and the effects of monetary shocks on aggregate house prices in the US, Section 3.4 presents the new stylised facts on the relative importance of different proximal drivers of US housing cycles and Section 3.5 offers some concluding remarks.

3.2 Testing whether house prices are consistent with fundamentals using monetary shocks

The fundamental value of a house in the absence of rational bubbles is defined as the full information rational expectation of the net present value of the stream of rents that are either received from renting the property or saved by living in it (Glaeser and Nathanson, 2015). Assuming or using assumptions that imply that house prices are equal to their fundamental value is extremely common across the housing literature: it is used in virtually all of the macroeconomic models of house prices surveyed by Piazzesi and Schneider (2016) and is the basis of many of the papers of housing cycles considered in Glaeser and Nathanson (2015). House prices

being equal to their fundamental value can be mathematically expressed as:

$$HP_t = E_t \left(\sum_{k=0}^{\infty} \frac{HR_{t+k}}{DR_{t+k}} \right) \quad (3.1)$$

where HP = real house prices, HR = real housing rents, DR = real discount rate and t denotes the time period.

The logic that is sometimes used for this valuation is that if the price of housing is below (above) its fundamental value then agents would demand more (less) housing, increasing (decreasing) house prices and restoring the condition. However there are rigidities in housing markets, such as search frictions and contractual rigidities, that may mean that this logic is not applicable. Papers such as Dusha and Janiak (2018) show how the fundamental value can be adjusted to also incorporate transaction costs, such as the cost of time spent searching for the right house. In this case house prices are equal to the net present value of the housing rents to sellers plus sellers transaction costs and are also equal to the net present value of the housing rents to buyers minus buyers transaction costs. One then also needs to adjust for the fact that in the US house prices are set when both parties sign a legal contract, which is usually between two and eight weeks before the closure of the deal⁵. It is not then generally possible to renegotiate the price unless property specific issues are found and penalties, such as the loss of earnest money payments, are often imposed if a party withdraws from the transaction. Therefore the expectations used in the adjusted fundamental value should be lagged to reflect this. Therefore house prices being equal to their adjusted fundamental value can be mathematically expressed as follows:

$$E_{t-1} \left(\sum_{k=0}^{\infty} \frac{HR_{t+k}^s}{DR_{t+k}^s} + TC_t^s \right) = HP_t = E_{t-1} \left(\sum_{k=0}^{\infty} \frac{HR_{t+k}^b}{DR_{t+k}^b} - TC_t^b \right) \quad (3.2)$$

where HP = real house prices, HR = real housing rents, DR = real discount rate, TC = real transaction costs, ^b denotes for buyers, ^s denotes for sellers and t denotes the time period.

⁵See the origination insight reports from Ellie May for survey data supporting this.

The logic in this case is similar to above but indicates that more agents will search for (try to sell) housing if the price is below (above) their fundamental valuation after adjusting for transaction costs at the horizon that accounts for agents having to make legal offers prior to exchange.

The test I use for whether house prices are consistent with being equal to their adjusted fundamental value is based on their reaction to monetary shocks. Monetary shocks are relatively unique in that if house prices are equal to their adjusted fundamental value they should have a clearly signed response to monetary shocks as soon as contractual rigidities stop binding. Monetary shocks are by definition unexpected (Ramey, 2016), so their effects cause changes in the expected values of sellers and buyers in Equation 3.2 once the contractual rigidity no longer binds and new expectations can be used. By considering their likely effect through each of the channels in Equation 3.2, one can demonstrate that the effects through all channels are likely to imply that expansionary (contractionary) monetary shocks imply increases (decreases) in house prices if they are equal to the adjusted value of housing once new expectations can be used. The following paragraphs consider each channel at a time under the null hypothesis that house prices are equal to their adjusted fundamental value.

The effects of an expansionary monetary policy shock through discount rates are relatively clear. The shock will reduce base rates directly. It is also likely to reduce risk premia, illiquidity premia and other premia that drive the difference between the housing discount rate and base rates. Theoretically this is likely to be true as improved economic and financial conditions could reduce housing risk and risk aversion while increasing housing liquidity, while empirically this is likely to be true as Gertler and Karadi (2015) show that expansionary monetary policy shocks reduce the housing finance premium. Therefore the shock is very likely to reduce discount rates and so increase house prices through this channel.

The effects through housing rents and transaction costs are less simple, and so are explained with reference to two simple conceptual frameworks that build on the adjusted fundamental value of housing and are included in Appendix D. The first also includes consumption demand and housing supply to explain the rents channel

and the second also includes search frictions to explain the transaction cost channel. The first framework explains the linked markets for housing consumption and housing ownership in terms of a consumption demand curve in the housing consumption market, a supply curve in the housing ownership market and an asset demand curve in both markets. This framework helps to demonstrate why the effect of an expansionary monetary shock on house prices through the consumption channel is very likely to be positive, even though housing rents themselves may not rise. Theoretically consumption demand is very likely to rise as a result of improved economic conditions and empirically other forms of consumption rise (Coibion et al., 2017). This should place upwards pressure on rents in the housing consumption market and the value of housing in the housing ownership market. However the increase in asset demand resulting from reduced discount rates will place positive pressure on house prices in the housing ownership market but negative pressure on housing rents in the housing consumption market. Therefore the overall effect on housing rents is ambiguous, although the effect on house prices should be unambiguously positive.

The second framework in Appendix D helps to demonstrate how even though part of the adjustment to an expansionary monetary shock could occur through search costs, part of it should unambiguously occur through house price rises. Equation 3.2 suggests that the only way house prices could fall is if transaction costs for sellers fall and/or transaction costs for buyers rise sufficiently to outweigh the increases in the discounted stream of rents through the channels discussed above. However this cannot happen as a result of changed search costs because it would be inconsistent with bargaining. It would imply that there are more buyers relative to sellers, which would increase the bargaining position of sellers, as their outside option has improved, while agents' valuation of owning a house would increase due to the shock. Therefore, bargaining should ensure that house prices increase, even if the shock causes a change in market tightness that affects search costs. There are also more unusual ways that expansionary monetary shocks may act through search frictions to increase house prices: for instance the analysis in Ngai and Tenreyro (2014) suggests that there will be thick market effects if the expansionary shock increases

the quantity of housing traded, as seems likely after an expansionary monetary shock.

Therefore one can be confident that monetary shocks should have unambiguously signed effects on house prices once contractual rigidities no longer bind if they are equal to the adjusted fundamental value of housing. This is not true for most other shocks, as the reaction of base rates will act in the opposite direction from most other channels in aggregate demand shocks and the reaction of housing supply will operate in the opposite direction from most other channels in aggregate supply shocks. Therefore we cannot say that there should be a clearly signed response of house prices to most shocks: in theory it will be ambiguous and even in practice if we suspect a sign it may well be sufficiently small that we cannot detect it.

The null hypothesis of the test of whether house prices are consistent with always being equal to their adjusted fundamental value is therefore that expansionary (contractionary) monetary shocks increase (decrease) house prices at a horizon of one to two months. This is only a test of consistency with house prices being equal to their adjusted fundamental value, as alternate explanations of house prices may also imply that they respond to monetary shocks rapidly. Therefore the test can, at least in theory, reject the hypothesis that house prices are equal to their adjusted fundamental value, but it cannot prove it. However several prominent alternate explanations of house prices could imply that house prices respond much more slowly to monetary shocks. For instance agents may not have accurate information on shocks in real time, violating the full information assumption, so would not be able to react to shocks as rapidly. Agents may also have non-rational expectations as they may not use all available variables when forecasting, due to the amount of effort involved, or may not adjust forecasts to complex new information, due to not understanding it. Therefore they may not react until easily accessible and understandable information on shocks becomes available. Indeed, even if there are only some agents that have non-rational expectations then those agents that do have rational expectations may exacerbate deviations from fundamental values (Brunnermeier and Oehmke, 2013), as no agent is wealthy enough to materially affect aggregate US house prices alone. Therefore evidence against the hypothesis would

fit with these alternate explanations of house prices.

To identify monetary shocks I build on the narrative approach of Romer and Romer (2004)⁶. These shocks were originally constructed by regressing the intended change in the effective federal funds rate around Federal Open Market Committee meetings on the Federal Reserve’s internal forecasts of its macroeconomic targets. However Ramey (2016) suggests that one could also control for the broader state of financial markets, which transmit monetary policy and so are also taken into account by the Federal Reserve. This is likely to be especially important in the case of an asset price like house prices, so I also include the Chicago Fed financial conditions index, the effective federal funds rate, the 5 year term spread and the 30 year mortgage rate as additional financial controls in the baseline case. However I also repeat the estimation without these additional controls to assess the changes they may cause. The baseline generating regression is therefore:

$$\begin{aligned} \Delta i_t^{int} = & \phi^1 + \phi^2 i_t^{int} + \sum_i \phi_i^3 \Delta E_{FR}(\pi_{t+i}) + \sum_i \phi_i^4 E_{FR}(\pi_{t+i}) + \sum_i \phi_i^5 \Delta E_{FR}(gr_{t+i}) \\ & + \sum_i \phi_i^6 E_{FR}(gr_{t+i}) + \sum_i \phi_i^7 E_{FR}(un_{t+i}) + \sum_j \phi_j^8 fin_t^j + \epsilon_t \quad (3.3) \end{aligned}$$

where i_t^{int} = the intended federal funds rate, $E_{FR}(x)$ = the federal reserve’s expectations of variable x , π = inflation, gr = GDP growth, un = unemployment and fin = financial control variables, t denotes the time period, i denotes the quarter of the forecast and ranges from -1 to 2, j ranges from 1 to 4.

To estimate the effect of monetary shocks on real house prices, I use these narrative monetary shocks in the local projection method of Jorda (2005). This approach is robust to the exact data generating process of real house prices and directly estimates their response to the narrative shocks at each forecast horizon. This is crucial when the timing of effects is of particular interest, as it allows me to be confident

⁶I don’t use high frequency shock measures based on futures markets for two reasons. Firstly, Miranda-Agrippino (2016) as well as Romer and Romer (2000) suggest that Federal Reserve forecasts are better than private sector forecasts. Secondly, I want to include the largest natural experiment in the modern era of US monetary policy: the period of non-borrowed reserve targeting, and data on high frequency measures is only available for more recent periods.

that the timing results from the genuine correlations between real house prices and monetary policy shocks at different lags and not simply from a mis-specified model. Ramey (2016) also suggests that the small sample estimates with narrative monetary shocks could be improved by including additional lagged controls, so I also include short (one year) and long (four year) real house price changes, real housing rents changes and housing starts. I also repeat the estimation without these additional controls to assess the changes they may cause. The baseline local projections regressions are therefore:

$$\Delta_h \ln H P_{t+h} = \beta^h s_t + \sum_k \gamma^k H V_{t-1}^k + \varepsilon_t \quad (3.4)$$

where $\ln H P$ = log real house prices, s = the narrative monetary shocks, $H V$ = housing control variables, t denotes the time period, h denotes the horizon and k ranges from 1 to 6.

I bootstrap across both stages of the process to account for the generated regressors in the standard errors. Specifically, I use a block normal bootstrap with block length of a year to allow for the additional variance produced by any remaining auto-correlation in the errors. This bootstrap approach is computationally efficient, as it only requires bootstrapping moments of the distribution instead of its extreme values, so I use 1000 repetitions at each forecast horizon.

I can then formally test whether there is a statistically and economically significant reaction of real house prices to monetary shocks at a horizon of one to two months. I can also calculate estimated impulse response functions of real house prices to monetary shocks. These let me compare my results to those from the existing literature and check that the overall sign of the response of real house prices is consistent with the narrative shocks not simply capturing endogenous movements in base rates and that the shocks are not so rapidly reversed that an identified expansionary shock is effectively a contractionary shock.

3.3 Conducting the test with narrative monetary shock data

I now move on to conducting the test of whether US real house prices are consistent with always being equal to their adjusted fundamental value by estimating the response of aggregate US house prices to narrative monetary shocks at different horizons.

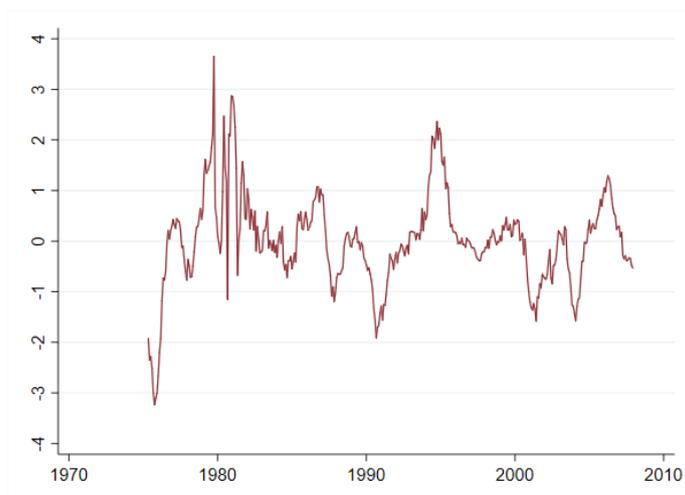
As discussed in Section 3.1, I produce the narrative shocks by using the approach of Romer and Romer (2004) but also including additional financial control variables. I therefore estimate Equation 3.3 and take the residuals as the narrative monetary shock series. This approach yields shocks which have very low levels of auto-correlation and are highly variable, so to present them graphically I take a one year moving sum. The resulting series is plotted in Figure 3.1. There are several periods in which the existing literature identifies clear narratives for exogenous loosening or tightening based on a combination of political and operational reasons. The first and largest of these is the period of non-borrowed reserve targeting from 1979-1982. However political pressure for loose policy in the late 1970s and the early 2000s as well as the desire to gain credibility in the 1990s are also commonly suggested⁷. All of these are visible in the shock series.

To estimate the effect of monetary shocks on real house prices, I use the narrative monetary shocks in the local projection method of Jordà (2005) with additional lagged housing controls, as discussed in Section 3.1. I therefore estimate Equation 3.4 at each horizon considered to produce an IRF of real house prices to a monetary shock. The data used is the monthly repeat transactions house price index produced by Freddie Mac deflated with CPI index. In Appendix D I also repeat the main analysis in this chapter using Case-Shiller data and the results are very similar.

Figure 3.2 shows the IRF of real house prices to a one percentage point decline in the narrative shock measure, which is a large expansionary shock but less than the largest absolute value observed in the shock series. The IRF is plotted to a horizon of eight years and a 95% confidence interval is shown along with the central estimates.

⁷See Romer and Romer (2004) and the sources cited therein for details

Figure 3.1: Moving sum of monetary shocks

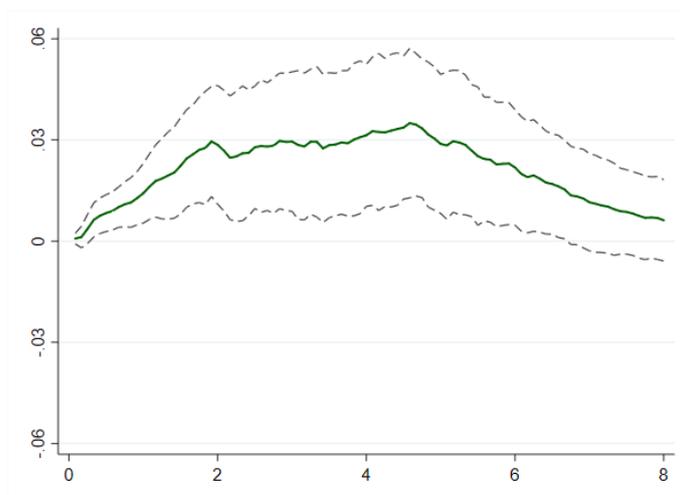


Notes: Plot of the one year centered moving sum of the narrative monetary shocks generated from Equation 3.3. The vertical axis is in percentage points and the horizontal axis is in years.

It is clear that the effects of the shocks near impact is virtually zero. This is not just because the effects of the narrative monetary shocks are always small, which could be the case if the monetary shocks measure partly capture endogenous influences or if monetary shocks are rapidly reversed and overcompensated for. The cumulative effect slowly rises from approximately zero to economically and statistically significant responses that are maintained over horizons of several years before declining to ultimately have effects close to zero again. The maximum effects occur at horizons of approximately two to five years, where a one percentage point expansionary shock causes an approximately three percentage point increase in real house prices, which is generally statistically significant at the 95% level. Therefore the monetary shocks do have significant effects, they just do not appear to occur near impact.

Table 3.1 plots the effects of a one percentage point decline in the narrative monetary shock measure at a horizon of one and two months to more formally examine the results at the horizons indicated by the test. The effect of the shocks at these horizons is virtually zero, both economically and statistically, quantitatively confirming what can be seen visually in Figure 3.2. The change corresponds to

Figure 3.2: Impulse response function of real house prices to a monetary shock



Note: Plot of the estimated cumulative log change in real house prices in response to a one percentage point decrease in the narrative federal funds shock measure with bootstrapped 95% confidence intervals. The vertical axis is in units, so 0.05 corresponds to a 5% increase, and the horizontal axis is in years since the shock.

only about a tenth of a percentage point in real house prices after a large (one percentage point) monetary shock and is not statistically significant at any reasonable level. There is, therefore, effectively no response near impact. The scale of the actual changes caused by the shocks in the data reinforces the messages above. The estimates suggest that even the approximately 3 percentage point increase in the narrative shocks at the height of the period of non-borrowed reserve targeting, easily the largest shock in the series, caused virtually no change in real house prices near impact, but caused a fall in real house prices of approximately 8% after two to three years.

The result that monetary policy shocks have no effect on house prices near impact is also consistent with visual inspections of IRFs of real house prices to monetary shocks near impact in existing empirical work. The vast majority of work appears to find essentially no effects near impact, but significant effects of with a lag measured in years. This is true for older papers, such as Fratantoni and Schuh (2003), Del Negro and Otrok (2007), Jarociński and Smets (2008) and Musso et al. (2011), that use

Table 3.1: Response of real house prices to monetary shocks near impact

Cumulative real house price growth	h = 1	h=2
Monetary shocks	-0.001 (-0.002 to .001)	-0.001 (-0.004 to .002)
House price growth (1 year)	.118* (.069 to .167)	.235* (.135 to .334)
House price growth (4 years)	-0.011 (-0.033 to .010)	-0.025 (-0.068 to .019)
Housing rents growth (1 year)	.014 (-0.105 to .133)	.015 (-0.232 to .262)
Housing rents growth (4 years)	-0.005 (-0.048 to .037)	-0.008 (-0.094 to .079)
Log housing sales (1 year)	-0.000 (-0.001 to .000)	-0.001 (-0.002 to .001)
Log housing sales (4 years)	.000 (-0.000 to .000)	.000 (-0.001 to .001)
Constant	.015 (-0.089 to .120)	0.024 (-0.187 to .236)
Adjusted R^2	0.382	0.413

Notes: The left column shows estimates of the effects of a one percentage point increase in the narrative federal funds shock measure at a one month horizon and the right column shows estimates at a two month horizon. Bootstrapped 95% confidence intervals are shown and * = significant at the 5% level.

timing and sign restrictions in auto-regressive models. It is also true of newer papers, such as Jorda et al. (2015a), Ungerer (2015), Coibion et al. (2017) and Alessi and

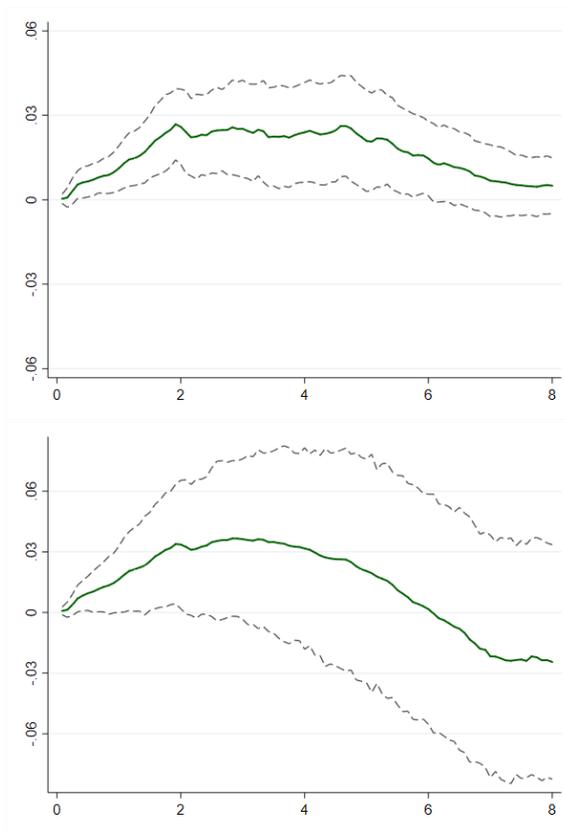
Kerssenfischer (2019) that mainly use narrative shocks. These papers find that a one percentage point expansionary monetary shock has a maximum positive impact on house prices of between one and ten percent. Therefore the existing empirical literature implicitly provides strong support for the results of my test.

It is also necessary to consider the effects that the control variables have on my results, as some of these controls are an addition to those commonly used in the literature to estimate the empirical effects of monetary shocks. Figure 3.3 shows the IRFs produced using no additional controls in the shock generation estimation (upper panel) or the local projection estimation (lower panel). The results are broadly similar to those in the baseline case, especially for the upper panel. All three IRFs have virtually no effects near impact, then slowly rise to have economically and statistically significant effects of approximately three percent at horizons of several years before declining again. The only clearly noticeable differences between the IRFs are whether real house prices are just above or just below their initial level at horizons over six years and even these differences are not large. Therefore the controls do not appear to be responsible for changing the results of the estimation dramatically.

This increases the confidence one can have in the results presented here, as the results in Figure 3.3 function as robustness checks. However it also may increase the confidence one could have in the results from the literature that estimate the effects of narrative monetary shocks on a wide variety of variables without including variable specific controls because, at least in the case of housing, these additional controls do not appear to be necessary.

Taken together, these results provide strong evidence against the hypothesis that house prices are always consistent with being equal to their fundamental value, even after this fundamental value is adjusted for search frictions and contractual rigidities. The results are, however, consistent with agents not observing monetary shocks well in real time or agents' expectation not reacting rationally to available information on monetary shocks. I would suggest that both of these factors are likely to play a role in explaining the results. Identifying monetary shocks in real time is clearly hard, as even experts struggle to identify monetary shocks (Ramey, 2016). Also

Figure 3.3: Impulse response function of real house prices to a monetary shock with fewer controls



Note: Plots of the estimated cumulative log change in real house prices in response to a one percentage point decrease in the narrative federal funds shock measure with bootstrapped 95% confidence intervals. The vertical axis is in units, so 0.05 corresponds to a 5% increase, and the horizontal axis is in years since the shock. The upper panel uses no additional controls in the shock generation estimation and the lower panel uses no additional controls in the local projection estimation.

many agents in housing markets are clearly not experts and so survey evidence, such as that in Case et al. (2012), suggests that some homebuyers do not have rational expectations. Monetary shocks do not seem relatively harder to measure or understand than many other shocks. Indeed, identifying monetary shocks benefits from the Federal Reserve having very clearly specified aims, presenting its policy tools in quantitative form and publishing information which Ericsson (2016) shows

can be used to essentially infer its information. The equivalent is not necessarily true for other shocks, such as productivity shocks or financial shocks. This suggests that agents may struggle to observe or use expectations that respond rationally to many macroeconomic shocks. Therefore models of housing markets based on full information rational expectations, which are common in macroeconomics, may be seriously mis-specified so their implications could be very misleading.

3.4 New stylised facts on the drivers of housing cycles

The two conceptual framework in Appendix D are primarily introduced to explain why monetary shocks are very likely to have an unambiguously signed effect on house prices once contractual rigidities no longer bind. However the first framework also implies a sign decomposition⁸ that can be used to analyse the relative importance of proximal drivers of house price movements. These proximal drivers are some of the channels through which the ultimate shocks that drive housing cycles are transmitted. This first conceptual framework explains the linked markets for housing consumption and housing ownership with a consumption demand curve that directly affects the housing consumption market, a supply curve that directly affects the housing ownership market and an asset demand curve that directly affects both markets. This implies the following effects of shifts in the curves, where all effects are normalised to increase real house prices: shifts in housing supply, such as those generated by housing regulation or building cost changes will reduce the quantity of housing and increase real rents. Shifts in consumption demand, such as those generated by higher incomes or consumer confidence, will increase the quantity of housing and real rents. Shifts in housing asset demand, such as those generated by increased real house price expectations or easier housing credit will increase the quantity of housing and reduce real rents. This implies the correlation structure for movements driven by each type of shift in Table 3.2.

⁸This sign decomposition is also very likely to remain valid if the search frictions in the second framework are included too, as discussed in Appendix D.

It is important to note that these frameworks aim to capture some of the main mechanisms in housing markets and I do not claim to capture all mechanisms in housing markets. This is why I limit the models to sign analysis that seems unlikely to be invalidated by including other variables. One aspect of this simplification is that I use one variable for each series, rather than including separate variables for each expected future value of the series. Therefore it is appropriate to apply the sign restrictions to components of the variables that have enough persistence⁹ that there will not be lagged effects that more than offset the effects of an initial change in a variable. As a result of this, and the interest in housing cycles due to their macroeconomic consequences, I focus on the cyclical components of housing variables. Burns and Mitchell (1946) historically suggested that business cycles occur at lengths of 1.5 to 8 years, however Drehmann et al. (2012) suggests that more recent housing cycles can have longer lengths, so in the baseline case I extract cyclical components with cycles of between 3 years and 40 years. However, I extract cyclical components with longer maximum cycle lengths and shorter minimum cycle lengths in Appendix D and there are no meaningful changes in my main results.

The correct filter to extract cycles from data depends upon the exact data being analysed. Parametric approaches such as unobserved components models require a model with very specific assumptions over the form of the components to be kept or removed. Therefore the cycle obtained can reflect the assumed model as much as the data in question. Non-parametric approaches, such as band-pass filters, overcome this main problem. These filters aim to isolate the components of a variable driven by cycles at a specified range of frequencies and so fit with the definition of housing cycles used here. It is necessary to difference any data used to remove unit roots before applying these techniques (Murray, 2003) and use a finite sample approximation to the ideal filter (Baxter and King, 1999). Band-pass filters include the desired frequencies from all components of the series: signal components and

⁹This is one of the reasons why it was necessary to confirm in the previous section that the narrative monetary shocks do have sensibly signed effects on house prices at reasonable horizons, as otherwise they might be so rapidly more than reversed that what was identified as an expansionary shock would effectively be a contractionary shock.

Table 3.2: Cyclical correlations implied by the first conceptual housing market framework

	HP	HQ	HR
HP	1		
HQ	-	1	
HR	+	-	1

	HP	HQ	HR
HP	1		
HQ	+	1	
HR	+	+	1

	HP	HQ	HR
HP	1		
HQ	+	1	
HR	-	-	1

Notes: Cyclical correlations based on shifts in the three curves in the first conceptual housing market framework. The top correlation matrix denotes a shift in housing supply, the middle correlation matrix denotes a shift in housing consumption demand and the bottom correlation matrix denotes a shift in housing asset demand. HP = real house prices, HR = real housing rents and HQ = the quantity of housing.

noise components. Therefore it is important to check that there really are strong auto-correlations in the data before being using a band-pass filter. Therefore, in

practice, one has to understand the order of integration and persistence of the data before deciding exactly how, and whether it is appropriate at all, to use them. To aid this decision in the case of real house prices, Table 3.3 displays the results of applying integration tests to real house prices in levels and growth rates and the auto-regressive parameters from the augmented Dickey-Fuller test.

Table 3.3: Integration and persistence properties of monthly real house prices

	Level	Growth rate
Augmented Dickey-Fuller test	-	*
Phillips-Perron test	-	*
Elliot et al. test	*	*
Dickey-Fuller persistence parameter	1.00	0.88

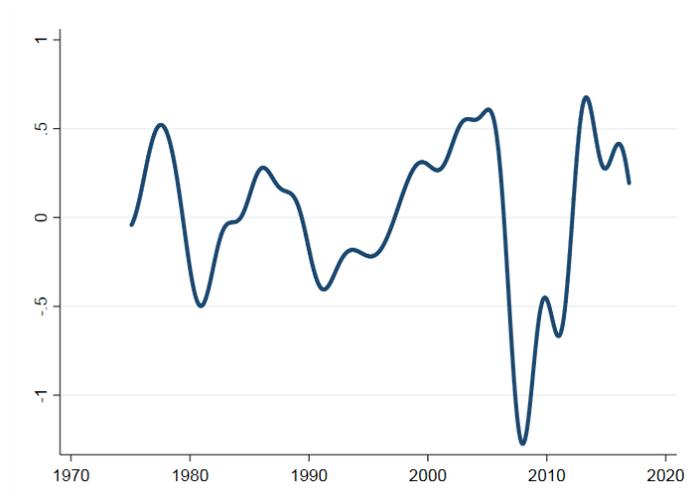
Notes: Results of applying integration tests to real house prices in levels and growth rates. The null hypothesis of the tests is integration of order of at least one and * = significant at the 5% level.

Despite the low power of unit root tests, the results support the assumption used in previous work, such as Drehmann et al. (2012), that real house prices contain a unit root in levels but not in growth rates. Therefore there is no non-stationary issue when applying a band-pass filter to real house price growth. The real house price growth series is highly persistent in growth rates, so I use the Christiano and Fitzgerald (2003) filter. This filter is fairly accurate for persistent processes, as it is derived from a random walk, and has the advantage of allowing filtered values to be constructed for all time periods in the sample. The high persistence of the series and visual inspection of the raw data suggest that the filter will primarily capture cyclical signal, rather than certain frequencies from noise.

Figure 3.4 shows the extracted cyclical component of monthly real house price growth data. A very clear cyclical component emerges, which is similar to that

produced in Drehmann et al. (2012). It is associated with US recessions: it has small troughs alongside the 1980s Volcker recessions and the early 1990s recession and a significant trough alongside the great depression. In all these cases it also peaks towards the end of the preceding booms. The only exception to this is that there is little decline around the dot-com recession.

Figure 3.4: Cyclical element of real house price growth

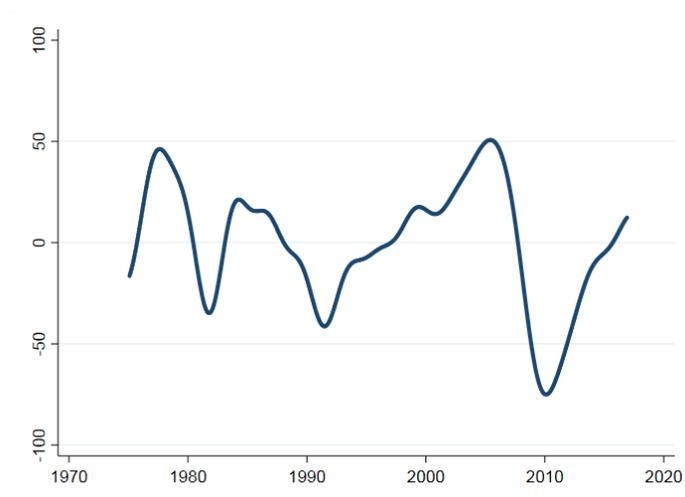


Notes: Plot of the component of monthly real house price growth rates with frequencies that give cycles of between 3 and 40 years in length, extracted with the Christiano-Fitzgerald filter. The vertical axis is in percentage points and the horizontal axis is in years.

I now produce equivalent cycles for rents and the quantity of housing supplied. Since real house prices are in growth rates I would ideally also take the growth rates of real rents and the housing stock. For real rents I use the growth rate of the CPI rent of primary residence series. Since there is no complete time series data on the US housing stock available, I use housing starts from the Census Bureau as a proxy. These series also conform to typical guides for stationarity. I apply the same Christiano-Fitzgerald filter to these series as applied to real house price growth to housing starts and real rental price growth. These cyclical components of these two variables are plotted in Figures 3.5 and 3.6 respectively. They both appear reasonable: the cyclical component of housing starts appear to be fairly similar to

the real house price cycle and so is also related to the US business cycle, as measured by NBER recessions. There does not appear to be as clear a cycle in rents as the other two variables, but some downturns are visible.

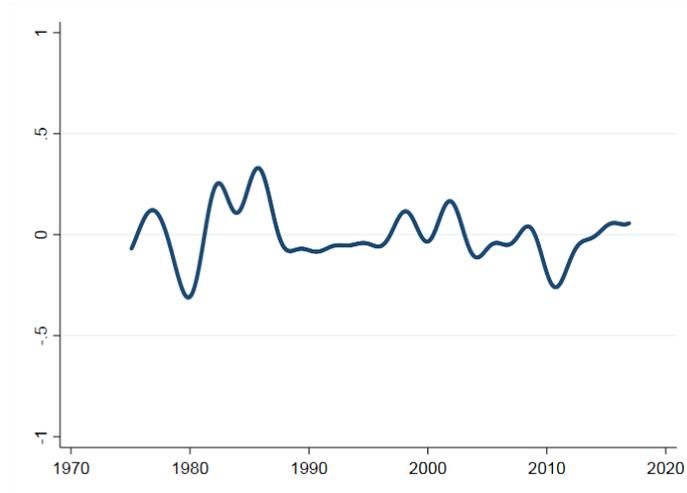
Figure 3.5: Cyclical element of housing starts



Notes: Plot of the component of monthly log housing starts (scaled up by 100 in construction) with frequencies that give cycles of between 3 and 40 years in length, extracted with the Christiano-Fitzgerald filter. The vertical axis is in units and the horizontal axis is in years.

I now have data series which correspond to the cyclical changes in the three main variables from the conceptual framework. Therefore their empirical associations can be used to judge the relative importance of housing consumption demand, housing asset demand or housing supply, based on which of the predicted associations in Table 3.2 are closest to the empirical associations. This may seem similar in principle to a sign restricted VAR or FAVAR, however there are important differences. Firstly, in my setting I do not assume that the changes I am capturing are independent and exhaustive, as I am clear in my conceptual framework that I am identifying transmission mechanisms, not structural shocks. For instance, a credit supply shock is extremely likely to influence both the asset and consumption demand for housing over time. Secondly, I don't use a very specific statistical model to try to remove all past correlations and trends, but instead aim to keep particular cyclical components

Figure 3.6: Cyclical element of real housing rental price growth



Notes: Plot of the component of monthly real housing rental price growth rates with frequencies that give cycles of between 3 and 40 years in length, extracted with the Christiano-Fitzgerald filter. The vertical axis is in percentage points and the horizontal axis is in years.

of interest, so have far more flexibility and robustness. On the other hand this flexibility implies that I cannot conduct the formal decompositions associated with sign restricted VARs or FAVARs and can only compare relative importance. Thirdly, sign restricted VARs impose their restrictions, so give no evidence as to the validity of the restrictions themselves, whereas I do not impose my restrictions when calculating correlation structures, so could obtain a result which simply supports the rejection of the conceptual framework.

To study the empirical associations between the three variables I calculate standard linear correlations and Spearman rank correlations, in case there are important non-linear associations. In both cases I produce confidence intervals using a block normal bootstrap with block length of a year, which is asymptotically valid as both statistics are asymptotically normal. The linear correlation matrix and 95% confidence intervals are displayed in Table 3.4. The results show a very strong positive and significant correlation between the house price series and the housing starts series. Even at the lower end of the confidence interval this correlation is still eco-

nomically significant and at the top end it is close to perfect positive correlation. The correlations between the housing rents series and the other two series are also positive, but they are weaker and only significant in one of the two cases.

Table 3.4: Housing market cyclical linear correlations

	HP	HN	HR
HP	1		
HN	0.62*	1	
	(0.40 to 0.83)		
HR	0.27*	0.27	1
	(0.06 to 0.49)	(-0.03 to 0.56)	

Notes: Linear correlations between the cyclical components of real house price growth, real rental price growth and housing starts. Bootstrapped 95% confidence intervals are also shown and * = significant at the 5% level. HP = real house prices, HR = real housing rents and HN = the quantity of new houses started.

Table 3.5 shows the Spearman rank correlation matrix. The rank correlations are extremely similar to the linear correlations: there is a very strong positive and significant correlation between the house price series and the housing starts series and the correlations between the housing rents series and the other two series are also positive, but they are weaker and less significant. Therefore there do not appear to be important non-linear associations between the variables and the results of the rank correlations simply support the results of the linear correlations.

The strong and significant positive correlations between the house prices and housing starts variables shows that housing supply shifts are clearly the least important of the three cyclical transmission mechanisms. This does not imply that there are no supply effects, but only that they are strongly outweighed by consump-

Table 3.5: Housing market cyclical rank correlations

	HP	HN	HR
HP	1		
HN	0.64*	1	
	(0.41 to 0.87)		
HR	0.31*	0.27	1
	(0.04 to 0.57)	(-0.03 to 0.57)	

Notes: Linear correlations between the cyclical components of real house price growth, real rental price growth and housing starts. Bootstrapped 95% confidence intervals are also shown and * = significant at the 5% level. HP = real house prices, HR = real housing rents and HN = the quantity of new houses started.

tion and asset demand shifts as a driver of housing market cycles. The positive correlations between the housing rents variable and each of the other two variables suggests that consumption demand is the most important driver of housing cycles. However, the correlations between housing rents and the other two variables are weaker than those between house prices and housing starts and only two of the four correlations are statistically significant. Therefore, asset demand changes may also have been important and in particular unusual periods could even have been the most important driver. For instance it is interesting that the only period in which cyclical real house price growth and housing starts appear to grow significantly accompanied by a decrease in cyclical real rental price growth is during the early 2000s housing boom, suggesting that shifts in asset demand may have outweighed shifts in housing consumption in this period. However over most of the sample shifts in consumption demand were the most important, shifts in asset demand were also relatively important and shifts in housing supply were the least important proximal

drivers of US housing cycles.

3.5 Conclusion

It is important to understand if house price fluctuations are driven by changes in fundamentals, because if they are not then this implies the potential for housing bubbles that could have serious macroeconomic consequences. However assessing whether changes in house prices are driven by changes in the fundamental value of housing is extremely difficult, especially once the definition of fundamental value is adjusted to account for the search frictions and contractual rigidities in housing markets.

In this chapter I introduce a new test of whether real house prices are consistent with always being equal to the adjusted fundamental value of housing. This test is based on the idea that monetary shocks should cause clearly signed reactions in the adjusted fundamental value of housing as soon as contractual rigidities no longer bind: an idea that I support with two conceptual frameworks and references to existing empirical work on the transmission mechanisms of monetary policy. Therefore I test whether real house prices have a significant reaction to monetary shocks at the horizon when contractual rigidities no longer bind, which is within two months.

I implement this test using narrative monetary shocks, which are constructed with the Romer and Romer (2004) approach, but augmented with finance and housing specific controls. The reaction of real house prices to these shocks is estimated using local projections to accurately capture the timing of the reactions. The results show that there are no statistically significant or economically meaningful reactions of real house prices to monetary shocks within two months of a monetary shock. However IRFs show that there are statistically significant and economically meaningful responses at horizons greater than a year and are similar to IRFs in the existing literature. Therefore the results provide strong evidence that house prices are not consistent with always being equal to the fundamental value of housing, even after this fundamental value has been adjusted for search frictions and contractual rigidities.

Instead the results are consistent with explanations such as agents' inability to observe monetary shocks precisely in real time or agents expectations not reacting rationally to available information about monetary shocks. Monetary shocks are not necessarily harder to observe than other shocks, and may well be easier, as the Federal Reserve publishes a large amount of information on its aims, expectations and policy tools in quantitative form. No similar information is available for other shocks such as productivity or financial shocks, suggesting that agents are unlikely to precisely observe and adjust their expectations rationally to these shocks either. As a result the many recent macroeconomic models that assume agents in housing markets use full information rational expectations may be seriously mis-specified and so may be of little practical relevance.

In additional work I also use a sign decomposition, suggested on the basis of the conceptual frameworks, to analyse the relative importance of different proximal drivers of US housing cycles. I implement this decomposition using linear and rank correlations between the cyclical components of different housing variables, where the cyclical components are extracted using an appropriately chosen band-pass filter. I find that the consumption demand channel has been the most important proximal driver of housing cycles, the asset demand channel has also been relatively important and the housing supply channel has been clearly the least important.

These results suggest that changes in house prices do not always result from changes in the fundamental value of housing, even after adjusting for the frictions in housing markets. The results are instead consistent with agents not observing many shocks and not adjusting their expectations rationally in reaction to information on shocks. Therefore housing cycles are likely to arise from the partially behavioural reaction of housing markets to shifts in the demand for housing consumption and housing assets.

Appendix A

Supplement: Behavioural financial cycles as the cause of perpetual business cycles

Currently the role of financial fluctuations in driving business cycles is mainly modelled using one of two broad approaches. The first approach, which is dominant in mainstream academia, uses dynamic stochastic general equilibrium models with financial frictions and is summarised in Brunnermeier et al. (2011), whereas the second approach is based on Minsky cycles and is summarised in Nikolaidi and Stockhammer (2017).

The first approach suffers from potentially serious deficiencies. It is built on rational expectation maximisation of utility functions, or very minor departures from such behaviour, which is very unrealistic given that most people in most countries do not display basic financial literacy (Lusardi and Mitchell, 2014). As a result, models based on this approach may also be very unrealistic, despite being relatively complex and opaque. The second approach does not suffer from these deficiencies. However it is built on Ponzi financing regularly occurring and driving financial cycles, which is not an accurate characterisation of most financial cycles: for instance it is not mentioned in surveys such as Borio (2014). It also implies the existence of a predictable cycle in asset or credit returns from which pure risk-adjusted profits could be made, despite even experts in the financial services industry struggling to

predict such a cycle well enough to act upon it (Cheng et al., 2014).

In this supplement I develop a simple conceptual framework of how financial cycles drive continual macroeconomic cyclicity: a result which can also be found in the Minsky approach but is not a feature of the DSGE approach. However my approach does not require Ponzi financing or an easily predictable cycle from which pure risk-adjusted profits could be made. The conceptual framework is built upon realistic psychological preferences, such as fear that reacts to events (Guiso et al., 2018) and optimal feasible expectations, which are introduced in Chapter 1 of this thesis.

This framework is intended to be a parsimonious conceptual framework to explain the ideas presented, rather than a technical and complex mathematical model. It therefore focuses on the key elements in financial markets needed to explain how simple and realistic behaviour could lead to a permanent financial cycle, which could in turn drive a permanent business cycle. My core analysis is conducted under the assumption that no excluded variables or features have sufficiently large counteracting effects to change the core analysis, although I do discuss how several additional features could be added to the framework. Given this, the framework does not attempt to derive exact quantitative relationships between variables, but focuses on giving signed and relative effects. As a result its implications are not dependent on any specific mathematical formulations and can be understood relatively easily.

I first introduce the three main areas of the framework: the determination of aggregate risk aversion, the determination of expected and realised excess returns on risky assets and the determination of macroeconomic variables. I then explain how these features could generate a permanent behavioural financial cycle and how this could be transmitted to generate a permanent business cycle. Finally I discuss the implications of these ideas.

I start with the determination of aggregate risk aversion. Aggregate risk aversion rises after large unexpected negative returns on assets, as these events provoke fear, however because the salience and memory of these events fades over time these events have less impact if they are further in the past. To capture this mathematically I specify that aggregate risk aversion is a negative function of the number of periods

since the last large unexpected decline in asset prices.

$$ra_t = f_1(\bar{ts}) \quad (\text{A.1})$$

where ra = aggregate risk aversion, ts = the number of periods since the realised excess return on risky assets, er_i , was sufficiently less than the expected excess return on risky assets, $E_{i-1}^{of}(er_i)$, to provoke an increase in fear-based preferences, t denotes in time period t and f_x denotes a function with the sign of differentials given by + or - above variables.

Next I consider the determination of the excess returns on risky assets, defined as the difference between the return on these assets and the base rate set by the central bank. These assets could include equities, property, higher risk corporate bonds and higher risk bank assets. Investors will purchase assets until the expected return is equal to the individual compensation that they would require for bearing the asset risk implied by holding a particular quantity of assets given their level of individual risk aversion. This implies that their expected excess return is equal to the compensation required by investors in aggregate for bearing the asset risk given aggregate risk aversion.

$$E_t^{of}(er_{t+1}) = f_2(ra_t^+) \quad (\text{A.2})$$

where E^{of} = optimal feasible expectation, er = excess return on risky assets, ra = aggregate risk aversion, t denotes in time period t and f_x denotes a function with the sign of differentials given by + or - above variables.

I now consider the difference between realised excess returns and agents' previous expectations of those real returns. Realised excess returns also contain the effects of shocks from economic and financial changes that are unrelated to the economy when expectations were formed, so by definition could not have been included in agents' expectations.

However the effects of these shocks are not the only difference between realised excess returns and agents' previous expectations of those real returns, as agents form their expectations using optimal feasible expectations based on common measures

of forecast error¹. Optimal feasible expectations are predicted to minimise the relevant measure of forecast error out of the set of expectations that agents can feasibly estimate, so capture the optimum trade-off between conditional forecast bias and conditional forecast variance². It is hard for agents to quantitatively measure aggregate risk aversion, as it is only weakly related to each individual's risk aversion and its dependency on past excess returns is complex and time-varying. It is also hard for agents to precisely estimate the effect of aggregate risk aversion on excess returns, as these effects are complex and time varying while excess returns are also subject to regular shocks. As a result it is optimal to shrink expected excess returns towards statistically simple forecasts. Therefore expected excess returns react less to aggregate risk aversion than realised excess returns do and may react in a more simplified form.

$$E_{t-1}^{of}(er_t) = er_t - f_3(ra_{t-1}^+) - \epsilon_t \quad (\text{A.3})$$

where E^{of} = optimal feasible expectation, er = excess return on risky assets, ra = aggregate risk aversion, ϵ = the effects of shocks that are unrelated to past events, t denotes in time period t and f_x denotes a function with the sign of differentials given by + or - above variables. The true conditional expectation of excess returns, $E_{t-1}^{tc}(er_t)$, is therefore equal to $f_2(ra_{t-1}) + f_3(ra_{t-1})$.

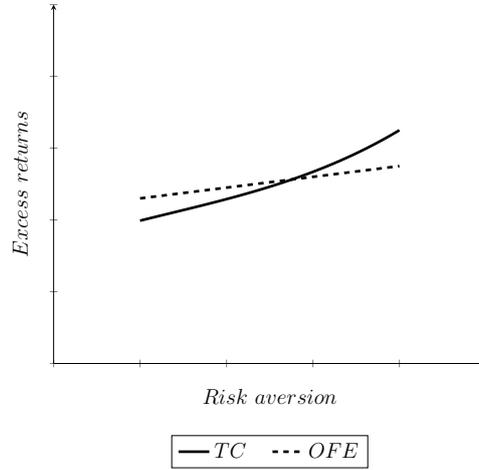
Illustrative functions of agents' expected excess returns produced with optimal feasible expectations and the true conditional expectation of excess returns, both in terms of risk aversion, are plotted in Figure A.1. The general results do not rely on the exact shape of the illustrative functions used in the figure, as only the general patterns of optimal feasible expectations described above are required.

I now consider the determination of the macroeconomic variables. The central bank sets base rates using a Taylor rule, based on current inflation and unemployment, to try and achieve its inflation target. The base rate reduces consumption and

¹I assume that it is not psychologically possible for agents to trick themselves into not becoming afraid by deliberately using expectations that have worse forecast performance but are less likely to produce results that raise risk aversion.

²I also assume that for similar reasons agents do not perfectly adjust other conditional expectations, such as the conditional expected dispersion in excess returns, to account for this effect.

Figure A.1: Expectations of excess returns



Notes: Plot of how the illustrative functions of agents' current expectations of future excess returns, formed with optimal feasible expectations, and the true conditional expectation of future excess returns vary with current aggregate risk aversion. TC = current true conditional expectations of future excess returns and OFE = current optimal feasible expectations of future excess returns.

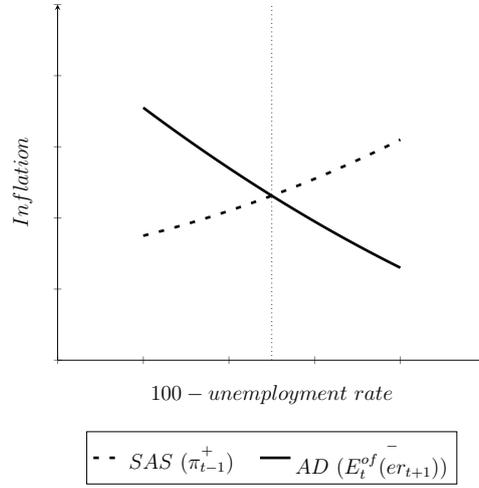
investment through both its effect on risky asset returns and its effect on non-risky asset returns, so reduces the demand for goods and workers and hence increases unemployment. However higher expected excess returns also reduce consumption and investment, so increase unemployment, at a given level of the base rate. Therefore these effects combine to give an aggregate demand curve. Unemployment decreases if inflation increases, due to the response of the central bank to inflation, and decreases if expected excess returns increase, due to its direct effects.

$$un_t = f_4(\pi_t^+, E_t^{of}(er_{t+1}^+)) \quad (\text{A.4})$$

where un = unemployment, π = inflation, E^{of} = optimal feasible expectation, er = excess return on risky assets, $_t$ denotes in time period t and f_x denotes a function with the sign of differentials given by $+$ or $-$ above variables.

In each period lower unemployment raises firms unit production costs but non-

Figure A.2: Aggregate supply and demand



Notes: Plots of the aggregate demand and short-run aggregate supply curves in terms of unemployment and inflation. A dotted line showing the flexible price level of unemployment is also included. *SAS* = short-run aggregate supply, *AD* = aggregate demand and π_{t-1} = inflation in period $t - 1$ and $E_t^{of}(er_{t+1})$ = agents' expectations in period t of excess returns in period $t + 1$.

inal rigidities, from wage and price rigidities, limit the extent of price increases and cause price changes to also depend on past price changes. Therefore this gives a short-run aggregate supply curve, as inflation falls with unemployment but rises with past inflation. This curve is defined so that inflation can only be constant at the level of unemployment that would prevail if there were no price rigidities and prices were perfectly flexible³.

$$\pi_t = f_5(un_t^-, \pi_{t-1}^+) \quad (\text{A.5})$$

where un = unemployment, π = inflation, t denotes in time period t and f_x denotes a function with the sign of differentials given by + or - above variables.

³This level is interpreted as being fixed here however hysteresis after financial crises, as suggested by the results in Cerra and Saxena (2011), could increase the macroeconomic impacts of the financial cycle described below.

Illustrative functions for aggregate supply and aggregate demand and the flexible price level of unemployment are plotted in Figure A.2. As with the illustrative plots of excess returns the results do not require the specific shapes of the functions used in the figure.

These components can generate a realistic continual financial cycle that drives a continual business cycle. In the description of this process below the following notation is used: ra = risk aversion, er = excess returns on risky assets, $E^{of}(er)$ = agents' expectations of the future excess return on risky assets, $E^{tc}(er)$ = the true conditional expectation of the future excess return on risky assets, ϵ = a default or payoff shock, un = unemployment and π = inflation. All variables except inflation are expressed in real terms and exclude any long-term structural changes. The process is continuous, so the description could be started at any point in the cycle. It operates by cycling through the three following categories of periods:

1. *Standard periods.* Fear falls as memories of large unexpected negative returns slowly fade over time, so behavioural risk aversion decreases over many years. Therefore individual investors require less compensation for bearing asset risk, so increased asset demand reduces agents' expected excess returns. However agents expectations of excess returns decrease by less than the true conditional expectation of excess returns as a result of optimal feasible expectations. Therefore the probability of excess returns well below agents expectations rises. The reduction in agents expected excess returns places upwards pressure on aggregate demand, however it occurs over a long enough time period that the effects of this on unemployment and inflation can mainly be counteracted by monetary policy. The movements of the variables are therefore typically: $ra : \downarrow$, $E^{of}(er) \downarrow$, $E^{tc}(er) - E^{of}(er) \downarrow$, $un - / \downarrow$, $\pi - / \uparrow$.
2. *Distress periods.* After enough time a negative shock of sufficient size to act as a trigger will occur. While the direct effects of this shock might not be very large they will cause an excess return well below agents' expectations, as these were previously conditionally biased towards their long-term average. This will then trigger a rise in fear-based risk aversion and

so a large rise in required expected future excess returns as a result of decreased asset demand. This will then cause a large reduction in aggregate demand that significantly increases unemployment and, over time, significantly reduces inflation. The movements of the variables are therefore typically: *initial* ϵ \downarrow , ra $\uparrow\uparrow$, $E^{of}(er)$ $\uparrow\uparrow$, $E^{tc}(er) - E^{of}(er)$ $\uparrow\uparrow$, un $\uparrow\uparrow$, π $\downarrow\downarrow$.

3. *Recovery periods.* After the downturn caused by this financial distress, aggregate demand is slowly restored by monetary and possibly other macroeconomic policies, so unemployment declines and inflation rises. Risk aversion will also fall slightly as a result of a slight reduction in fear, however this reduction will only be fairly small as the financial distress will still be salient in investors' memories since it will still be quite recent. Therefore agents' required compensation for bearing risk, their expectations of excess returns, only decrease slightly. The movements of the variables are therefore typically: ra \downarrow , $E^{of}(er)$ \downarrow , $E^{tc}(er) - E^{of}(er)$ \downarrow , un \downarrow , π \uparrow .

Therefore the financial cycle is primarily driven by slow declines in risk aversion after periods of financial distress reducing agents' expected excess returns, but by less than they reduce true conditional expected excess returns due to optimal feasible expectations, so increasing the probability of a new period of financial distress being triggered. However this process occurs stochastically and is too unpredictable for it to be optimal in terms of forecast performance to try and include it fully in expectations. It is possible for large exogenous shocks to cause financial distress even relatively soon after the recovery from previous periods of distress: it is just that a smaller trigger would be needed longer after the recovery so such a trigger is more likely to occur. This perpetual financial cyclicity causes perpetual macroeconomic cyclicity through the effects of agents' expected excess returns on aggregate demand⁴. Therefore the macroeconomy will be characterised by an uncertain but perpetual cycle.

This conceptual framework could also be widened to include other features. For

⁴This does not however exclude the possibility of significant macroeconomic shocks that do not cause periods of financial distress.

instance one could incorporate the zero lower bound by changing the sign of the effect of inflation on unemployment below a particular level of inflation. One could also incorporate forecast-performance reducing mistakes in expectations when considering markets where this seems likely, such as the housing market, by making agents' expected returns also depend upon features such as behavioural optimism and pessimism. These features would be interesting, but they are not required to explain the key ideas presented here so are not included in the main conceptual framework for the sake of parsimony.

This framework has three main implications. Firstly, we need to be very aware of the perpetual cycle in financial and macroeconomic conditions, due to realistic behavioural risk aversion and expectations, and embed this in current thinking and analysis. These concepts are missing in the DSGE models used by many central banks, so these models will not provide accurate explanations of financial cycles. These models should be replaced with frameworks that incorporate realistic behavioural risk aversion and expectations.

Secondly, we need as good real-time measures of financial cycles as possible. The suggestion to use house prices and credit quantities as measures of financial cycles in Drehmann et al. (2012) is a good place to start. These capture the aspects of financial cycles that could have particularly large real effects, as credit and housing markets are particularly powerful transmitters to the real economy (Jorda et al., 2015b), and have been shown to forecast financial distress in real time (Borio, 2014). However more work in this area is needed, as financial cycle indicators should receive a comparable level of attention to that received by business cycle indicators.

Thirdly, we need macro-prudential policy tools to counter the issues presented by continual financial cyclicalities. Monetary policy will respond indirectly to these issues, as they potentially affect the future paths of inflation and unemployment as well as the transmission of monetary policy. However, it does not generally seem sensible for monetary policy to be used to directly counteract potential financial cycles, as this could impose huge costs in terms of macroeconomic stability (Williams, 2015). Therefore it is very important to develop cyclical macro-prudential policy tools that particularly target financial stability and keep these distinct from the

conventional and unconventional monetary policy tools that are directly used for inflation targeting.

The conceptual framework in this supplement suggests that macro-financial cyclicity is an inherent part of our economic system. However, with better understanding, monitoring and policy responses we might be able to reduce the extent of this cyclicity and the damage that it causes.

Appendix B

Appendix of ‘Optimal feasible expectations in our uncertain economy’

B.1 Variables used to construct factors

The following tables contain a complete list of the variables used in Chapter 1 and the transformations applied to them. APC stands for annual percentage change and the sources and transformations are discussed in Section 1.3.

Table B.1: List of dependent, non-factor and price factor variables

	Transformation	Source
CPI: Index	APC	Fred MD
Annual inflation expectations	None	Michigan Consumer Survey
Trade-weighted exchange rate index	APC	Fred MD
Average hourly earnings	APC	Fred MD
Narrative monetary shocks	None	Ramey (2016)
PPI: Finished Goods	APC	Fred MD
PPI: Finished Consumer Goods	APC	Fred MD
PPI: Intermediate Materials	APC	Fred MD
PPI: Crude Materials	APC	Fred MD
Crude Oil, spliced WTI and Cushing	APC	Fred MD
PPI: Metals and metal products	APC	Fred MD
CPI: Apparel	APC	Fred MD
CPI: Transportation	APC	Fred MD
CPI: Medical Care	APC	Fred MD
CPI: Commodities	APC	Fred MD
CPI: Durables	APC	Fred MD
CPI: Services	APC	Fred MD
CPI: All Items Less Food	APC	Fred MD
CPI: All items less shelter	APC	Fred MD
CPI: All items less medical care	APC	Fred MD
PCE: Chain index	APC	Fred MD
PCE: Durable goods	APC	Fred MD
PCE: Nondurable goods	APC	Fred MD
PCE: Services	APC	Fred MD
CPI: Food at home	APC	BLS
CPI: Food away from home	APC	BLS
CPI: Rent of primary residence	APC	BLS
CPI: Fuel and utilities	APC	BLS
CPI: New and used motor vehicles	APC	BLS
CPI: Motor fuel	APC	BLS
CPI: Medical care services	APC	BLS
CPI: Other goods and services	APC	BLS
PCE: Excluding food and energy	APC	BEA
PCE: Energy goods and services	APC	BEA
PCE: Food	APC	BEA
Sticky price index	APC	Atlanta Fed
Sticky price index less food and energy	APC	Atlanta Fed
Sticky price index less shelter	APC	Atlanta Fed
Flexible price index	APC	Atlanta Fed
Flexible price index less food and energy	APC	Atlanta Fed

Table B.2: List of business cycle factor variables

	Transformation	Source
Real Personal Income	APC	Fred MD
Real personal income ex transfer receipts	APC	Fred MD
IP Index	APC	Fred MD
IP: Final Products and Nonindustrial Supplies	APC	Fred MD
IP: Final Products	APC	Fred MD
IP: Consumer Goods	APC	Fred MD
IP: Durable Consumer Goods	APC	Fred MD
IP: Nondurable Consumer Goods	APC	Fred MD
IP: Business Equipment	APC	Fred MD
IP: Materials	APC	Fred MD
IP: Durable Materials	APC	Fred MD
IP: Nondurable Materials	APC	Fred MD
IP: Manufacturing	APC	Fred MD
IP: Residential Utilities	APC	Fred MD
IP: Fuels	APC	Fred MD
Civilian Labor Force	APC	Fred MD
Civilian Employment	APC	Fred MD
Civilians Unemployed - <5 weeks	APC	Fred MD
Civilians Unemployed - 5-14 weeks	APC	Fred MD
Civilians Unemployed - >14 weeks	APC	Fred MD
Civilians Unemployed - 15-26 weeks	APC	Fred MD
Civilians Unemployed - >27 weeks	APC	Fred MD
Initial Claims	APC	Fred MD
All Employees: Total nonfarm	APC	Fred MD
All Employees: Goods-Producing	APC	Fred MD
All Employees: Mining	APC	Fred MD
All Employees: Construction	APC	Fred MD
All Employees: Manufacturing	APC	Fred MD
All Employees: Durable goods	APC	Fred MD
All Employees: Nondurable goods	APC	Fred MD
All Employees: Service-Providing	APC	Fred MD
All Employees: Transport and others	APC	Fred MD
All Employees: Wholesale Trade	APC	Fred MD
All Employees: Retail Trade	APC	Fred MD
All Employees: Financial Activities	APC	Fred MD
All Employees: Government	APC	Fred MD

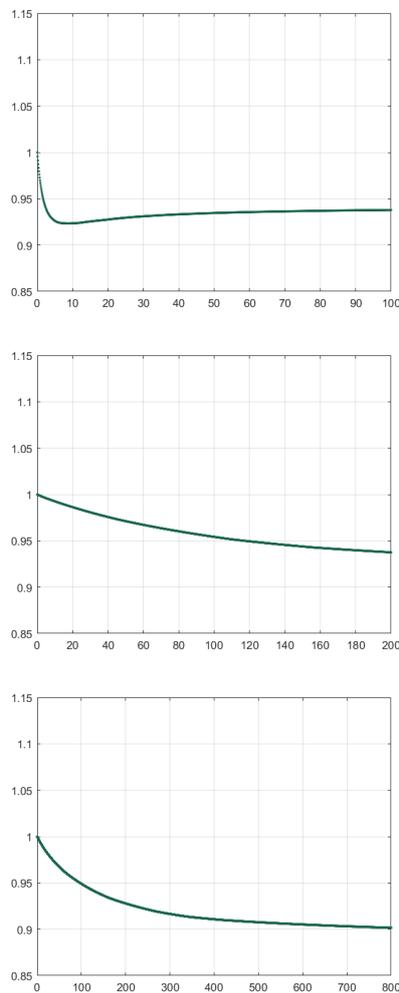
Table B.3: List of financial cycle factor variables

	Transformation	Source
Commercial and Industrial Loans	APC	Fred MD
Real Estate Loans at All Commercial Banks	APC	Fred MD
Total Nonrevolving Credit	APC	Fred MD
Nonrevolving consumer credit to Personal Income	APC	Fred MD
Consumer Motor Vehicle Loans Outstanding	APC	Fred MD
Total Consumer Loans and Leases Outstanding	APC	Fred MD
Securities in Bank Credit at All Commercial Banks	APC	Fred MD
Total Consumer Credit Owned and Securitized	APC	Fed Board
Total Revolving Credit Owned and Securitized	APC	Fed Board
House prices: National index	APC	Freddie Mac
House prices: New York index	APC	Freddie Mac
House prices: Los Angeles index	APC	Freddie Mac
House prices: Chicago index	APC	Freddie Mac
House prices: Dallas index	APC	Freddie Mac
House prices: Houston index	APC	Freddie Mac
House prices: Washington index	APC	Freddie Mac
House prices: Miami index	APC	Freddie Mac
House prices: Philadelphia index	APC	Freddie Mac
House prices: Atlanta index	APC	Freddie Mac
House prices: Boston index	APC	Freddie Mac
House prices: Phoenix index	APC	Freddie Mac
House prices: San Francisco index	APC	Freddie Mac
House prices: Riverside index	APC	Freddie Mac
House prices: Detroit index	APC	Freddie Mac
House prices: Seattle index	APC	Freddie Mac
House prices: Minneapolis index	APC	Freddie Mac
House prices: San Diego index	APC	Freddie Mac
House prices: Tampa index	APC	Freddie Mac
House prices: Denver index	APC	Freddie Mac
House prices: St Louis index	APC	Freddie Mac
House prices: Baltimore index	APC	Freddie Mac
House prices: Orlando index	APC	Freddie Mac
House prices: Charlotte index	APC	Freddie Mac
House prices: San Antonio index	APC	Freddie Mac
House prices: Portland index	APC	Freddie Mac
House prices: Sacramento index	APC	Freddie Mac
House prices: Pittsburgh index	APC	Freddie Mac
House prices: Las Vegas index	APC	Freddie Mac
House prices: Cincinnati index	APC	Freddie Mac
House prices: Austin index	APC	Freddie Mac
House prices: Kansas index	APC	Freddie Mac

B.2 Training set size robustness

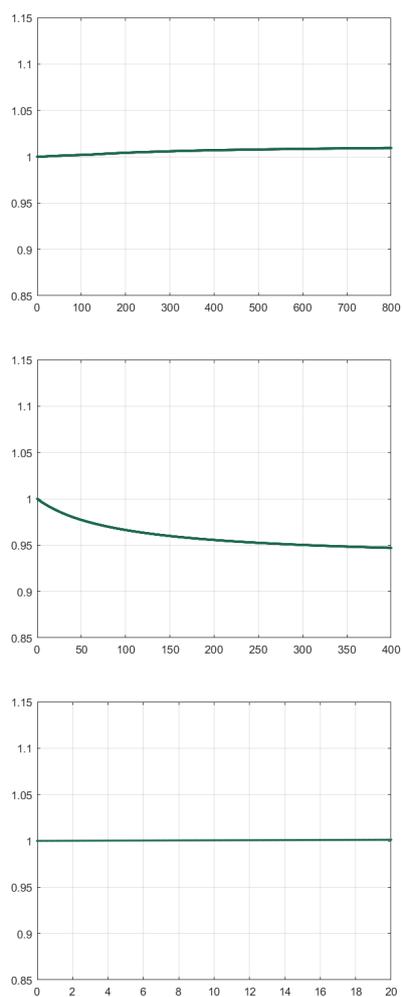
This section of the appendix contains the robustness checks from using different proportions of the sample as the training sample. In particular, it repeats analysis from Section 1.4 but takes either 60% or 80% of the full sample as the training sample, instead of the 70% in the baseline analysis.

Figure B.1: Forecast error with shrunken inflation factor, business cycle factor and financial cycle factor



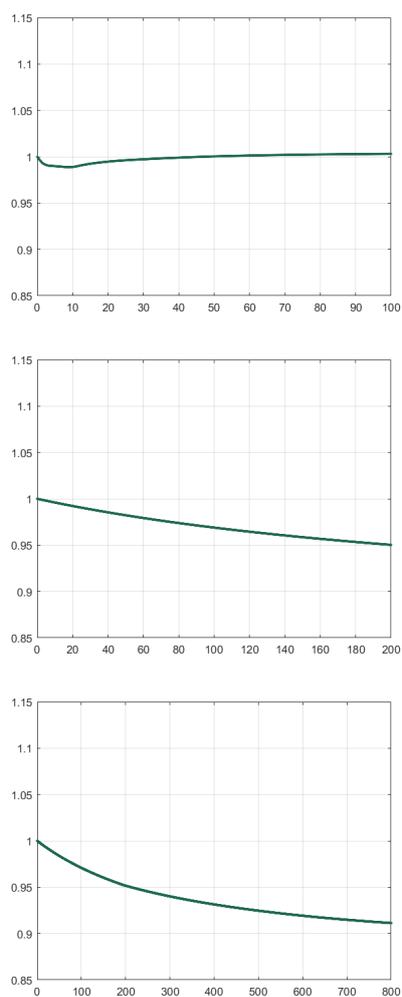
Notes: This figure repeats the graphs in Figure 1.4 but with estimates based on training sample of 60% of the total dataset.

Figure B.2: Forecast error with shrunken exchange rates, wages and narrative federal funds market monetary shocks



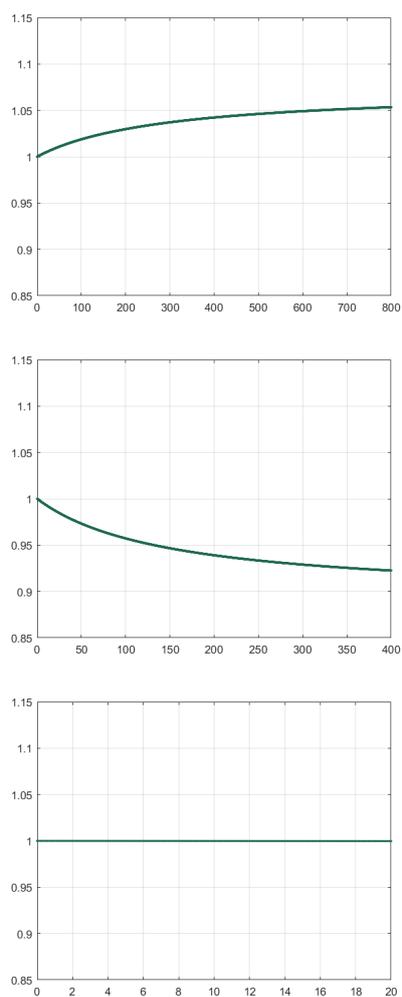
Notes: This figure repeats the graphs in Figure 1.5 but with estimates based on training sample of 60% of the total dataset.

Figure B.3: Forecast error with shrunken inflation factor, business cycle factor and financial cycle factor



Notes: This figure repeats the graphs in Figure 1.4 but with estimates based on training sample of 80% of the total dataset.

Figure B.4: Forecast error with shrunken exchange rates, wages and narrative federal funds market monetary shocks

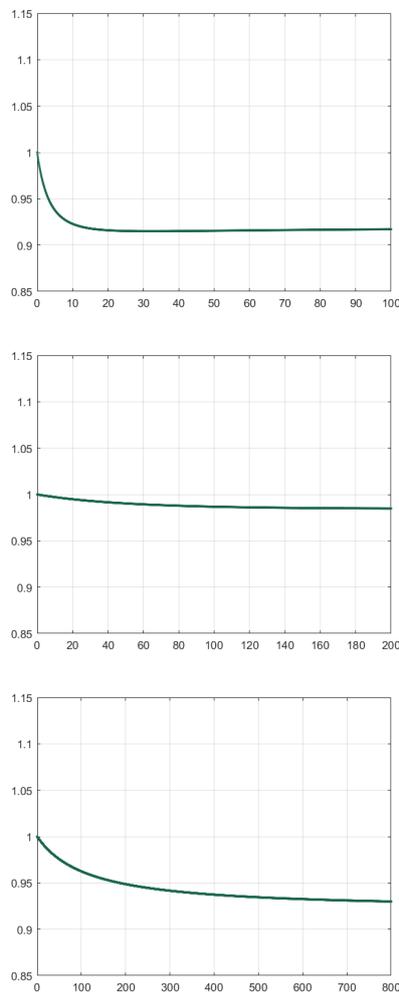


Notes: This figure repeats the graphs in Figure 1.5 but with estimates based on training sample of 80% of the total dataset.

B.3 Forecast performance measure robustness

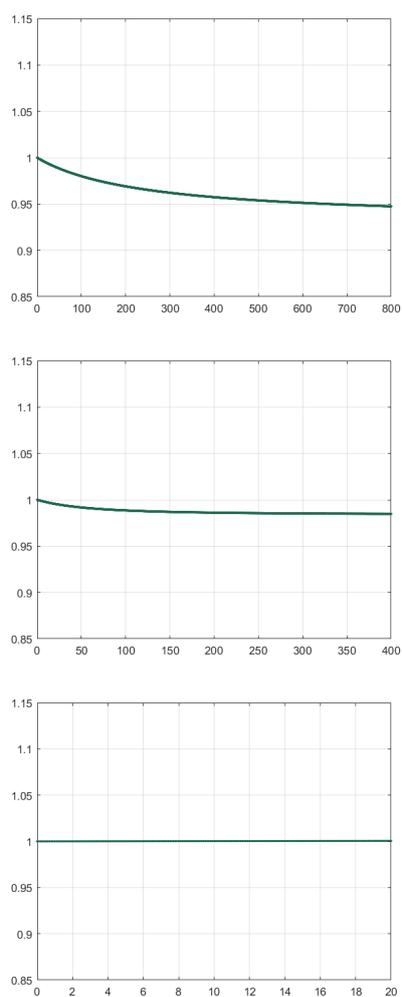
This section of the appendix contains the robustness checks from using a different measure of forecast performance. In particular, it repeats analysis from Section 1.4 but uses the mean square forecast error instead of the mean absolute forecast error as the measure of out of sample forecast performance.

Figure B.5: Forecast error with shrunken inflation factor, business cycle factor and financial cycle factor



Notes: This figure repeats the graphs in Figure 1.4 but uses the mean square forecast error instead of the mean absolute forecast error.

Figure B.6: Forecast error with shrunken exchange rates, wages and narrative federal funds market monetary shocks



Notes: This figure repeats the graphs in Figure 1.5 but uses the mean square forecast error instead of the mean absolute forecast error.

Appendix C

Appendix of ‘Asset price convergence, international asset holdings and the quality of financial integration’

C.1 Gibbs sampler

The Gibbs sampler consists of five main steps at each iteration. These steps are to draw in turn from:

1. $p(f_i|\Gamma, \Lambda, \Theta, \Phi, \Omega, \Sigma, y)$: Equations 2.1 and 2.2 describe a state-space system so the conditional posterior given the parameters and the data is normally distributed and can be calculated using the forward and backward recursions of the Carter and Kohn (1994) algorithm.
2. $p(\Theta, \Phi|f_i, \Sigma, y)$: Equation 2.2 becomes a multivariate regression once we condition on financial integration, so if we also condition on Σ then the conditional distribution of Θ, Φ is available from the natural conjugate case of the mean of a multivariate normal distribution with known variance and is normally distributed.

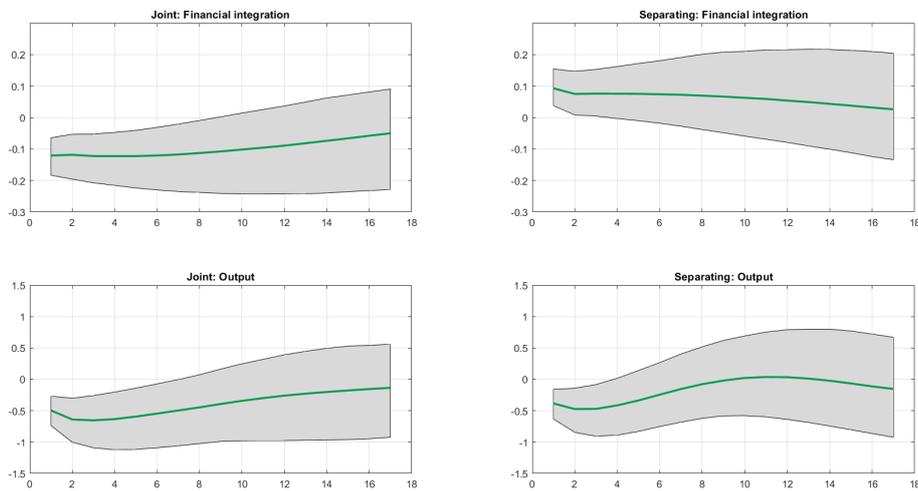
3. $p(\Sigma|fi, \Theta, \Phi, y)$: Equation 2.2 becomes a multivariate regression once we condition on financial integration, so if we also condition on Θ, Φ then the conditional distribution of Σ is available from the natural conjugate case of the variance of a multivariate normal distribution with known mean and is distributed inverse-Wishart.
4. $p(\Gamma, \Lambda|fi, \Omega, y)$: Each of the first n rows of Equation 2.1 become univariate regressions once we condition on financial integration, so if we also condition on Ω then the conditional distribution of each element of Γ, Λ is available from the natural conjugate case of the mean of a univariate normal distribution with known variance and is normally distributed.
5. $p(\Omega|fi, \Gamma, \Lambda, y)$: Each of the first n rows of Equation 2.1 become univariate regressions once we condition on financial integration, so if we also condition on Γ, Λ then the conditional distribution of each diagonal element of Ω is available from the natural conjugate case of the variance of a univariate normal distribution with known mean and is distributed inverse-Gamma.

Each step then uses the draws from the current iteration of parameters which have already been drawn and the draws from the previous iteration when they have not been drawn. For the generalised recursions of the Carter-Kohn algorithm for a state-space system and generalised expressions for the conditional posterior multivariate and univariate distributions see Blake and Mumtaz (2017).

C.2 Signed impulse response functions

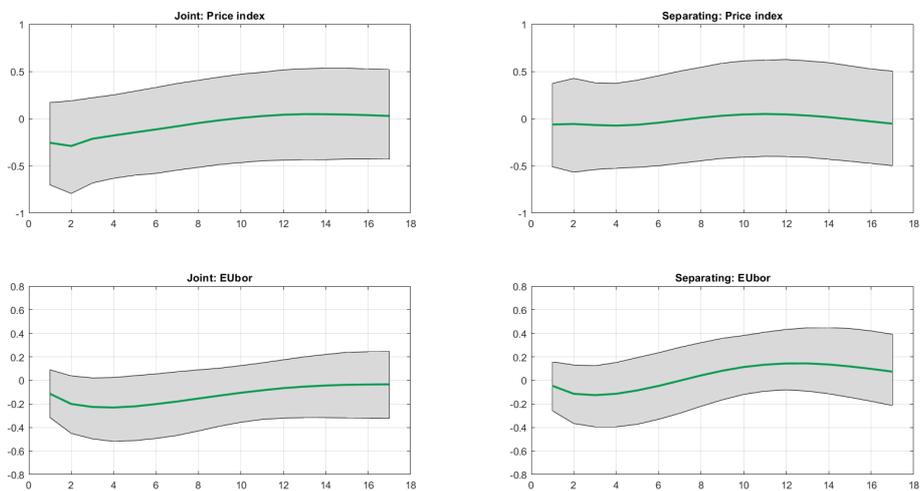
This section of the appendix contains the impulse responses of the seven macroeconomic variables and the financial integration indicator to a one standard deviation change in the joint and separating shocks. As discussed in Section 2.4 these are not particularly informative. This is because we classify the shocks into two types, rather than individual shocks, so the responses of different variables may reflect the impacts of different shocks.

Figure C.1: Main impulse response functions



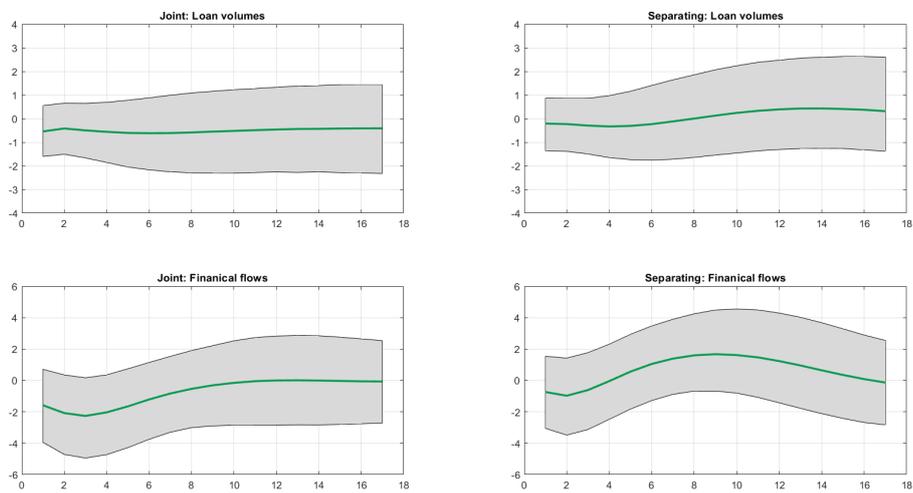
Notes: Each line plots the posterior median of the impulse response of a variable to a one standard deviation change in the joint or separating shocks. In each case the area between the 5th and 95th posterior percentiles is shaded.

Figure C.2: Other impulse response functions (one of three)



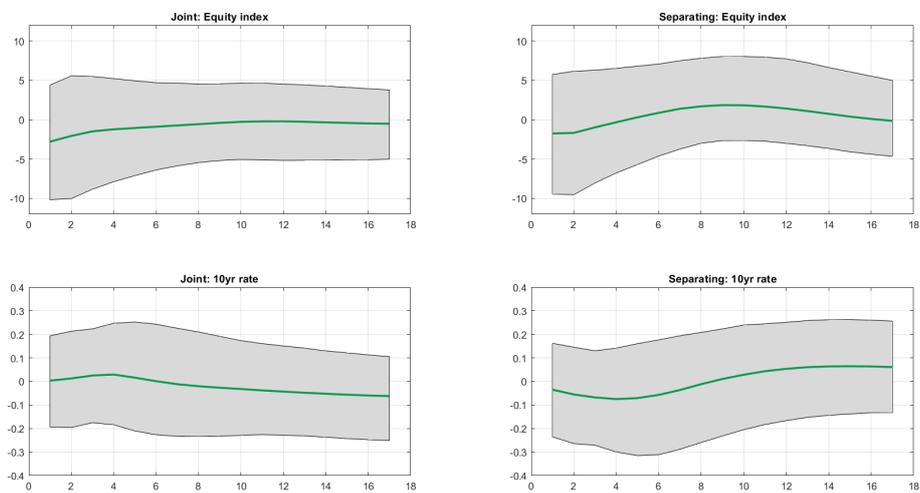
Notes: Each line plots the posterior median of the impulse response of a variable to a one standard deviation change in the joint or separating shocks. In each case the area between the 5th and 95th posterior percentiles is shaded.

Figure C.3: Other impulse response functions (two of three)



Notes: Each line plots the posterior median of the impulse response of a variable to a one standard deviation change in the joint or separating shocks. In each case the area between the 5th and 95th posterior percentiles is shaded.

Figure C.4: Other impulse response functions (three of three)



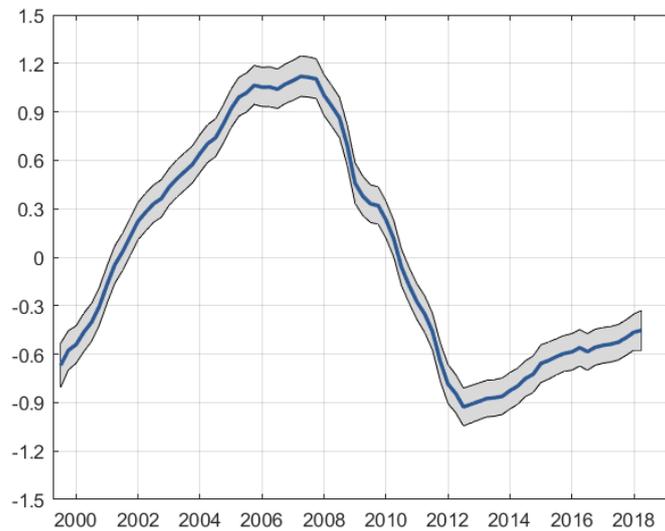
(C.1)

Notes: Each line plots the posterior median of the impulse response of a variable to a one standard deviation change in the joint or separating shocks. In each case the area between the 5th and 95th posterior percentiles is shaded.

C.3 Convergence robustness

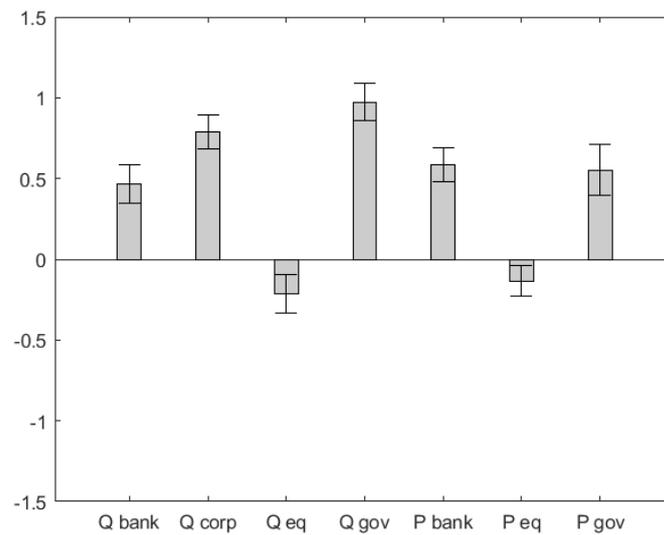
We check the numerical convergence of the markov chain monte carlo method by using a different initialisation and increasing the number of repetitions of the algorithm. Specifically we use an arbitrary initialisation based on matrices of zeros and ones and then take 100,000 draws and burn the first 50,000. The results of this procedure are shown below and are virtually identical to those in the main text, confirming that convergence is not an issue.

Figure C.5: Convergence (one of two)



Notes: This repeats the graph from Figure 2.3 but taking 100,000 draws from the MCMC algorithm and burning the first 50,000.

Figure C.6: Convergence (two of two)

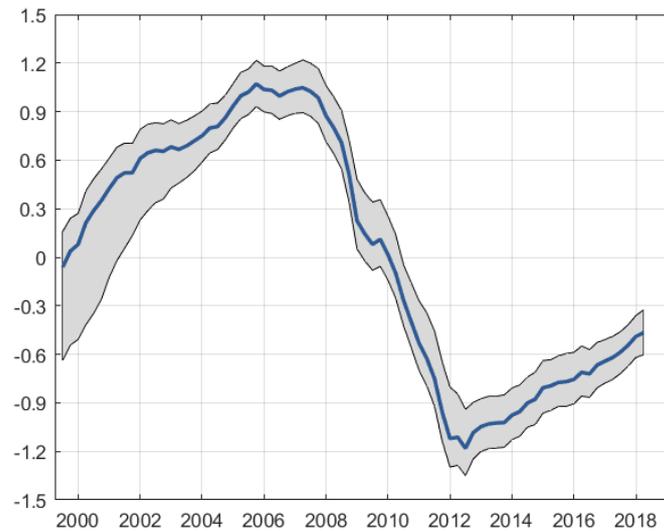


Notes: This repeats a graph from Figure 2.4 but taking 100,000 draws from the MCMC algorithm and burning the first 50,000.

C.4 Prior robustness

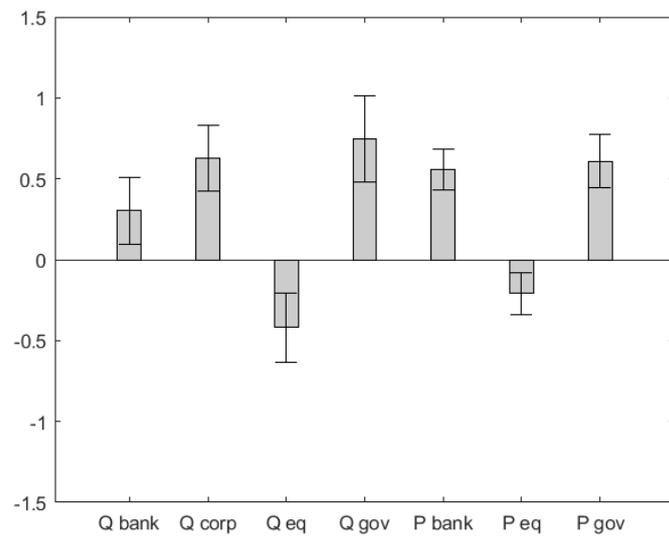
As discussed in Section 2.2, our priors are selected based on the typical approach in the literature and are set to be fairly loose. However it is still useful to check whether the results are robust to the choice of hyperparameter values. To do this we repeat the estimation with looser and tighter priors. In the first case we loosen the priors. We do this by doubling the hyperparameters where a higher value corresponds to a looser prior and halving the hyperparameters where a higher value corresponds to a tighter prior. In the second case we tighten the priors by reversing which hyperparameters we double and halve from the baseline values. It is reassuring that the results remain relatively similar given the fairly large changes in prior tightness.

Figure C.7: Prior robustness - looser priors (one of two)



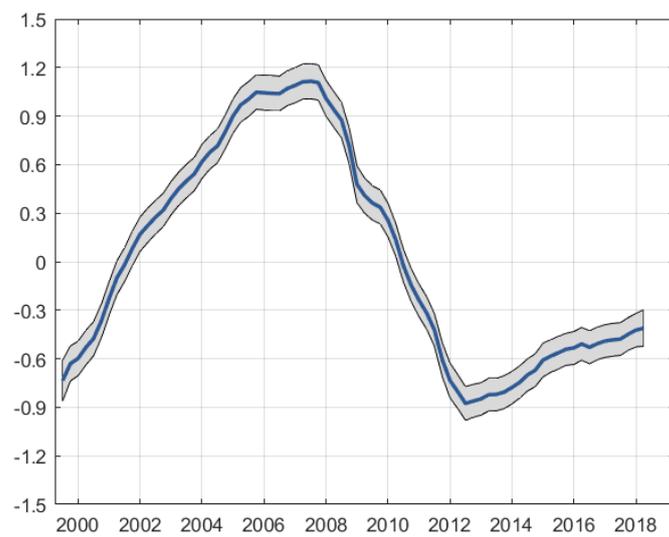
Notes: This repeats the graph from Figure 2.3 but with slightly looser priors.

Figure C.8: Prior robustness - loser priors (two of two)



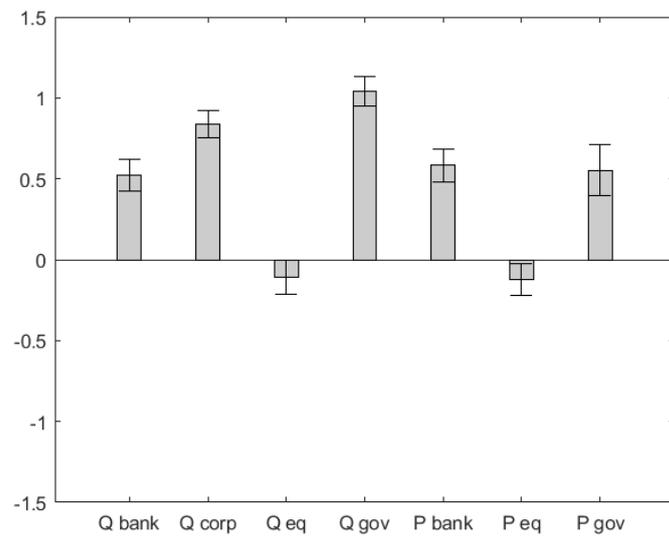
Notes: This repeats a graph from Figure 2.4 but with slightly looser priors.

Figure C.9: Prior robustness - tighter priors (one of two)



Notes: This repeats the graph from Figure 2.3 but with slightly tighter priors.

Figure C.10: Prior robustness - tighter priors (two of two)

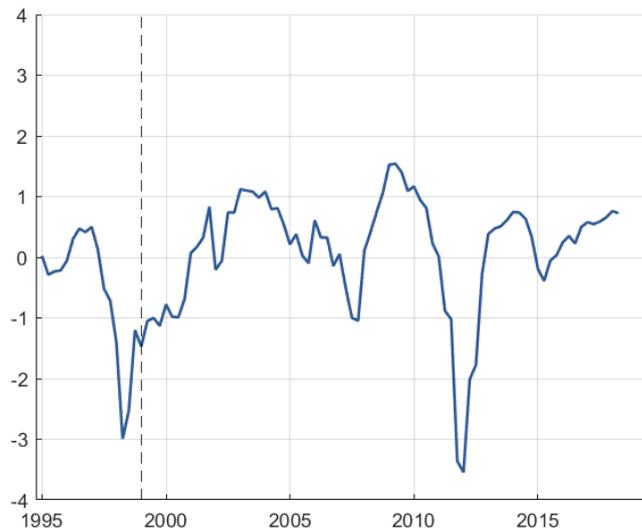


Notes: This repeats a graph from Figure 2.4 but with slightly tighter priors.

C.5 Extended dataset

There is limited comparable data on EU financial integration indicators going back before 1999, as most datasets are collected by the European Central Bank, which was only founded in 1998. There is even less comparable data on EU financial integration indicators going back before 1995, as most early datasets were collected by the European Monetary Institute, which was only founded in 1994. This issue is acknowledged in existing work, such as Hoffmann et al. (2019) and so in this section of the appendix we follow Hoffmann et al. (2019) and construct a much more basic financial integration index based only on equity price indicators and government bond price indicators going back to 1995.

Figure C.11: Basic financial integration index



Notes: Plot of the basic financial integration index produced by taking the cross-sectional mean of the subset of price indicators that are available going back to 1995. The dashed line shows the start of the sample in the analysis in the main text.

Since data to construct our macroeconomic variables is not available going back to 1995 we cannot meaningfully generate a financial integration index in a FAVAR.

This, coupled with the lack of any quantity financial integration indicators, means that we cannot conduct the tests of the three aspects of the quality of financial integration changes. Instead all we can effectively do is take a weighted combination of the subset of price indicators that are available as a financial integration index: here we do this by taking the cross-sectional mean of these indicators. The resulting indicator is plotted in Figure C.11. This index is similar to the index in the main text at the start and the end of the comparable sample, from 1999 onwards, however it has an unusual decline from approximately 2004 to 2008. This decline does not appear credible and reflects the limited nature of the data used to construct the index. The index does appear to follow a broadly similar cyclical tendency in the period before 1999 as the factor in the main text does afterwards though.

Ideally we would conduct the analysis in the main text using data that covers a longer time period however, as this section of the appendix makes clear, this is not possible by extending the dataset back to earlier periods. Indeed, the alternative approach and data used in this section of the appendix to try and do this actually highlight the usefulness of the methods and data that we use in the main text. It would however be interesting for future research to return to the results in the main text in the future and extend the dataset forwards to include the 2020s, so the sample would also include the effects of Brexit and other macroeconomic developments.

Appendix D

Appendix of ‘Behavioural finance at home: house price cycles in the USA’

D.1 Conceptual frameworks

This section of the appendix introduces two conceptual frameworks that build on the expression for the fundamental value of housing¹ introduced in Section 3.2. One which uses the basic fundamental value expression in a system with housing consumption and housing supply and one which uses the fundamental value expression with search frictions and bargaining. The first of these demonstrates the transmission of monetary shocks through housing rents, and suggests a sign decomposition for the proximal drivers of house price fluctuations, while the second demonstrates the transmission of monetary shocks through search costs.

I start with the first framework, in which the housing market effectively consists of two interrelated areas: housing consumption and housing ownership. The real prices of housing consumption are real rents and the real prices of housing ownership are real house prices. The total quantity of houses that people desire to own and the total number of houses that people live in and so consume the housing services of are

¹Throughout this section of the appendix house prices, housing values, housing rents and other related variables are all considered in real terms unless stated otherwise.

equal. Therefore there are three key variables in the market: real house prices, real housing rents and the quantity of housing. Clearly these are not the only relevant variables in reality, so I limit my analysis to the sign of likely effects and specify that this is under the assumption that no excluded variables have sufficiently large counteracting effects to change the signs implied here.

I initially consider the standard fundamental value condition and show that this is likely to be a special case of a more general expression of housing asset demand in this framework. The fundamental value condition is an expression for house prices that is a positive function of housing rents. However it is also likely to be a negative function of the quantity of housing agents choose to own, which I call housing asset demand. This is because as agents choose to own more housing they will have to bear more of the risks associated with housing, since asset classes are imperfect substitutes for one another, so the premia attached to the specific risks in housing are likely to rise. This in turn can be easily rearranged to give an expression for the quantity of housing owned, i.e. housing asset demand, that falls with real house prices and rises with real housing rents, as follows:

$$\begin{aligned}
HP_t &= E_t \left(\sum_{k=0}^{\infty} \frac{HR_{t+k}}{DR_{t+k}} \right) \\
HP &= E \left(\sum \frac{HR}{DR(HAD)} \right) \\
HP &= g(HR^+, HAD^-) \\
HAD &= f(HP^-, HR^+)
\end{aligned} \tag{D.1}$$

where HP = real house prices, HR = real housing rents, DR = real discount rate, HAD = housing asset demand and f_x denotes a function with the sign of differentials given by + or - above variables.

This expression is sensible. Housing asset demand is very likely to increase in real housing rents, as these, or the ability to forgo paying them to others, are part of the real return of holding a housing investment along with expected price changes and non-monetary benefits. Housing asset demand is also very likely to decrease with real house prices, as this is the price of purchasing the expected stream of real

returns.

I can then use similar logic to also construct equivalent signed expressions for housing consumption demand and housing supply. I start by considering the supply of housing. It will be more profitable for businesses and households to supply additional housing if the real price they receive for it is higher. Therefore housing supply is very likely to increase with real house prices. Whereas, I assume that housing supply only responds to real housing rents in so far as they affect real house prices. Next I consider the demand for housing consumption, which includes the consumption of housing services from houses that agents own. When real rents are high agents face a high price for housing consumption and are very likely to demand less of it. Therefore housing consumption demand is a negative function of real housing rents. Whereas, I assume that the demand for housing consumption supply only responds to real house prices in so far as they affect real housing rents.

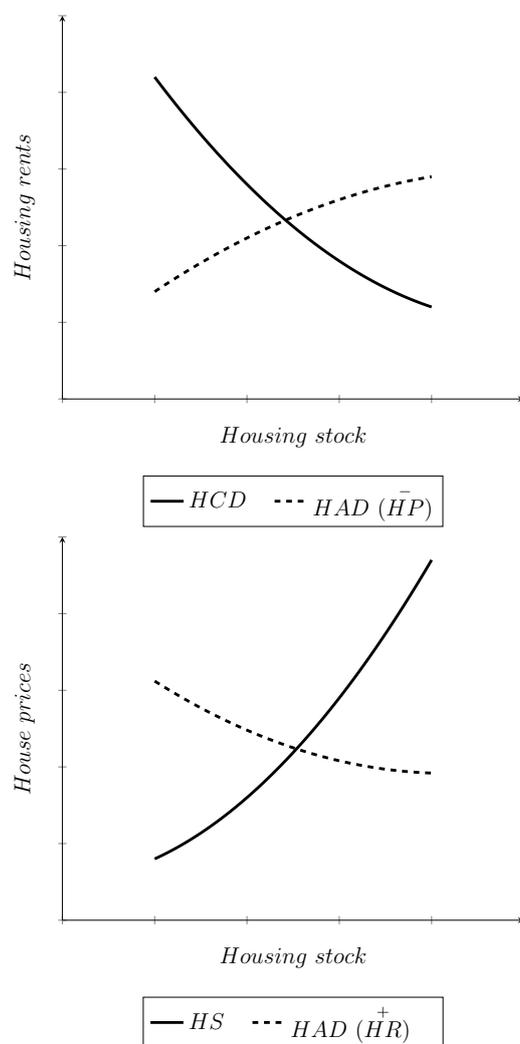
Therefore the expressions for housing supply, housing consumption demand and housing asset demand, plotted in Figure D.1, can be summarised as:

$$\begin{aligned}
 HS &= f_1(\overset{+}{HP}) \\
 HCD &= f_2(\overset{-}{HR}) \\
 HAD &= f_3(\overset{-}{HP}, \overset{+}{HR})
 \end{aligned}
 \tag{D.2}$$

where HS = housing supply, HCD = housing consumption demand, HAD = housing asset demand, HP = real house prices, HR = real housing rents and f_x denotes a function with the sign of differentials given by + or - above variables.

If the housing market clears, so demand equals supply, then shifts in each of the curves would have the effects in Table D.1, normalising the effect on real house prices to be positive. Shifts in housing supply, such as those generated by housing regulation or building cost changes will reduce the quantity of housing and increase real rents. Shifts in consumption demand, such as those generated by higher incomes or consumer confidence, will increase the quantity of housing and real rents. Shifts in housing asset demand, such as those generated by increased real house price expectations or easier housing credit will increase the quantity of housing and reduce

Figure D.1: Plots of housing ownership and consumption markets in the first framework



Notes: Plots of the conceptual functions for housing supply, housing consumption demand and housing asset demand, all measured in terms of the housing stock, as a function of either real housing rents or real house prices in the first conceptual framework. The upper panel shows the housing consumption market and the lower panel shows the housing ownership market. HS = housing supply, HCD = housing consumption demand, HAD = housing asset demand, HP = real house prices and HR = real housing rents.

real rents. This directly yields the correlation structures implied by the three shifts in Table 3.2. It is important to note that these three curves represent the main transmission mechanisms which underlie the housing market, rather than structural shocks. Structural shocks need to be independent of one another, which is not the case for transmission mechanisms.

Table D.1: Signed effects of shifts in housing market curves

	HP	HQ	HR
Supply	+	-	+
Consumption demand	+	+	+
Asset demand	+	+	-

Notes: Effects implied by the conceptual framework of a shift in each of the three curves, normalised to increase house prices. HP = real house prices, HR = real housing rents and HQ = the quantity of housing. All effects are normalised to have a positive effect on house prices.

The effects of an expansionary monetary shock on real house prices can then be analysed as follows. The reduction in expected base rates, risk and liquidity premia discussed in Section 3.2 will increase asset demand at any given levels of housing rents and house prices, shifting the asset demand curve to the right in the housing consumption and housing ownership diagrams in Figure D.1. This will place upwards pressure on house prices and downwards pressure on housing rents. The expected increase in housing consumption from expected increased economic and financial conditions discussed in Section 3.2 will increase consumption demand at any given level of housing rents, shifting the consumption demand curve to the right in the housing consumption diagram in Figure D.1. This will place upwards pressure on housing rents and so indirectly shift the asset demand curve to the right

in the housing ownership diagram in Figure D.1, placing upwards pressure on house prices. Therefore the overall effect on real housing rents is ambiguous, as the effects through consumption demand and asset demand operate in opposing directions, so it is possible for real housing rents to fall in reaction to an expansionary monetary shock. However real house prices will unambiguously rise, once contractual rigidities allow them to change, as the asset demand and consumption demand channels will both drive increases in house prices.

I now move onto the second framework, in which I only consider the market for housing ownership but now introduce search frictions in this market. Specifically I assume that it takes time for buyers and sellers to match with one another and the probability of a successful match in any period for a buyer (seller) decreases (increases) with the market tightness, i.e. the number of buyers relative to the number of sellers. I therefore only consider two main variables in this setting: real house prices and the market tightness. As with the first conceptual framework, I acknowledge that these are not the only relevant variables in reality, so I limit my analysis to the sign of likely effects and again specify that this is under the assumption that no excluded variables have sufficiently large counteracting effects to change the signs implied here.

I initially consider the adjusted fundamental value condition, introduced in Section 3.2, with a particular focus on the adjustment for transaction costs, which in this setting will be driven by the search frictions. I treat the discounted sums of rents as exogenous valuations of owning a house for buyers and sellers. Since the time spent searching for a house will be costly, as it will imply reduced leisure or working time, the longer an agent expects to spend searching the higher their transaction costs are. This implies that if the market is tighter, so search time increases for buyers and decreases for sellers, then transaction costs increase for buyers and decrease for sellers. With exogenous valuations this implies that house prices must fall to stop it being profitable for a net increase in sellers in the market to occur. Therefore the adjusted fundamental value of housing implies that house prices are a negative function of market tightness in this setting, as follows:

$$\begin{aligned}
E_{t-1} \left(\sum_{k=0}^{\infty} \frac{HR_{t+k}^s}{DR_{t+k}^s} + TC_t^s \right) &= HP_t = E_{t-1} \left(\sum_{k=0}^{\infty} \frac{HR_{t+k}^b}{DR_{t+k}^b} - TC_t^b \right) \\
E \left(HDS^s + TC^s \left(\frac{\bar{HB}}{HS} \right) \right) &= HP = E \left(HDS^b - TC^b \left(\frac{+HB}{HS} \right) \right) \quad (D.3) \\
HP &= m_1 \left(\frac{\bar{HB}}{HS} \right)
\end{aligned}$$

where HP = real house prices, HR = real housing rents, DR = real discount rate, TC = real transaction costs, HDS = the real discounted sum of rents, $\frac{HB}{HS}$ = market tightness, ^b denotes for buyers, ^s denotes for sellers and m_x denotes a function with the sign of differentials given by + or - above variables.

After a match occurs then the buyer and the seller have to bargain over the price, as the buyers willingness to pay is greater than the sellers reservation price. Their relative bargaining strength will determine which level between these two values the agreed house price is set at. This bargaining strength will be affected by the value of each negotiators reserve option of returning to search, which will be lower for the buyer and higher for the seller if the market is tighter. Therefore house price bargaining implies that house prices are a positive function of market tightness, as follows:

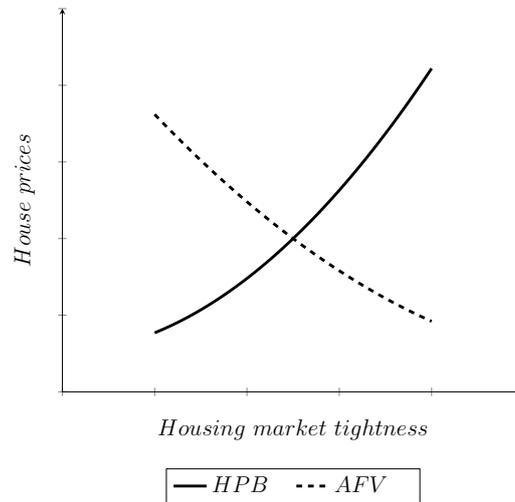
$$\begin{aligned}
HP &= HDS^s + \phi (HDS^b - HDS^s) \\
HP &= HDS^s + \phi \left(\frac{+HB}{HS} \right) (HDS^b - HDS^s) \quad (D.4) \\
HP &= m_2 \left(\frac{+HB}{HS} \right)
\end{aligned}$$

where HP = real house prices, HDS = the real discounted sum of rents, $\frac{HB}{HS}$ = market tightness, ϕ = sellers' relative bargaining strength between 0 and 1, ^b denotes for buyers, ^s denotes for sellers and m_x denotes a function with the sign of differentials given by + or - above variables.

Therefore the relationships between real house prices and housing market tightness implied by the conceptual functions for the adjusted fundamental value of housing and the house price bargaining condition have different signs. Since only the

housing ownership market is considered and the discounted sum of rents is treated as exogenous in this conceptual framework these two functions are the only ones considered. They are plotted in Figure D.2.

Figure D.2: Plot of the housing ownership market in the second framework



Notes: Plots of the relationships between real house prices and housing market tightness implied by the functions for the adjusted fundamental value of housing and the house price bargaining condition in the housing ownership market of the second conceptual framework. AFV = adjusted fundamental value of housing and HPB = house price bargaining condition.

The effects of an expansionary monetary shock in this framework can then be analysed as follows. The effects on base rates, premia and consumption demand would raise the discounted sum of housing rents for both buyers and sellers. This would cause both the adjusted fundamental value curve and the house price bargaining curve in Figure D.2 to shift upwards. Therefore the effect on market tightness is ambiguous, so it is possible that the response to the shock includes a change in market tightness. However both curves will place upwards pressure on house prices, so house prices will unambiguously rise regardless of the change in market tightness.

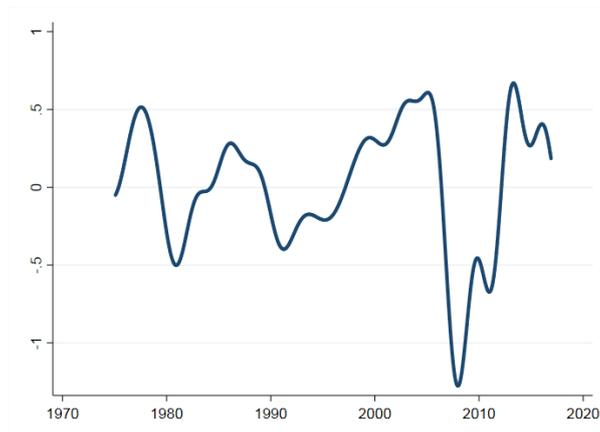
It is also worth noting that one could expand the first framework, which includes the housing ownership and consumption markets, to also include the search frictions

from the second framework in the housing ownership market. In this case the estimated signed effects from the first framework would remain true as long as the original mechanisms were not entirely counteracted by changes in search costs from the response of market tightness. Such a change in market tightness in response to shifts in housing consumption demand, housing asset demand or housing supply shocks seem extremely unlikely to be compatible with both asset demand and house price bargaining behaviour, even if these are not exactly as described in the second framework. Therefore the signed effects of the shifts in Table D.1 are very likely to remain the same in such a setup.

D.2 Cycle periodicity robustness

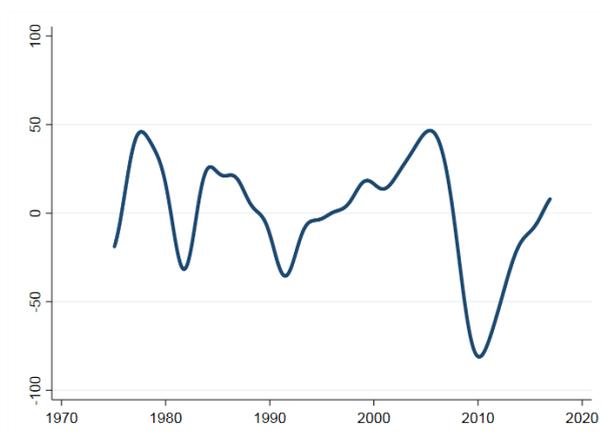
This section of the appendix contains the robustness checks from increasing the maximum cycle length and reducing the minimum cycle length of the cyclical components extracted from the housing variables. The following graphs and tables show the cyclical components and the associations between them but either extend the maximum cycle length extracted to 50 years or reduce the minimum cycle length extracted to 1.5 years. In both cases the results are similar to those in the baseline case.

Figure D.3: Cyclical element of real house price growth with longer cycles



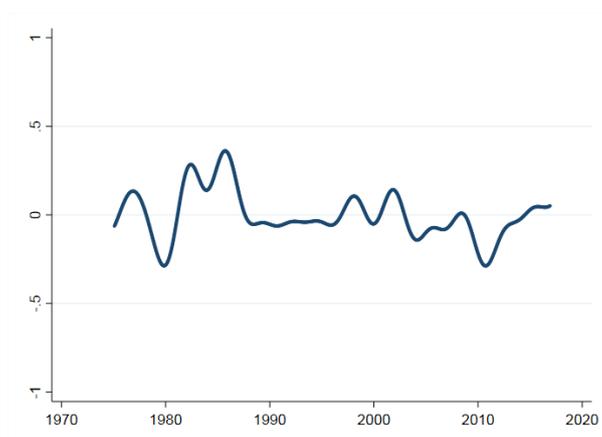
Notes: This repeats the graphs from Figure 3.4 but with frequencies that give cycles of between 3 and 50 years in length.

Figure D.4: Cyclical element of housing starts with longer cycles



Notes: This repeats the graphs from Figure 3.5 but with frequencies that give cycles of between 3 and 50 years in length.

Figure D.5: Cyclical element of real housing rental price growth with longer cycles



Notes: This repeats the graphs from Figure 3.6 but with frequencies that give cycles of between 3 and 50 years in length.

Table D.2: Housing market cyclical linear correlations

	HP	HS	HR
HP	1		
HS	0.60*	1	
	(0.39 to 0.81)		
HR	0.26*	0.31	1
	(0.03 to 0.48)	(-0.01 to 0.62)	

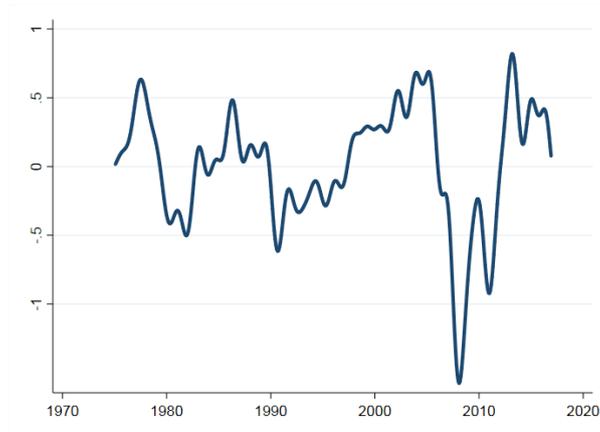
Notes: This repeats the results from Table 3.4 but with frequencies that give cycles of between 3 and 50 years in length.

Table D.3: Housing market cyclical rank correlations

	HP	HS	HR
HP	1		
HS	0.61*	1	
	(0.37 to 0.85)		
HR	0.24	0.20	1
	(-0.03 to 0.51)	(-0.14 to 0.53)	

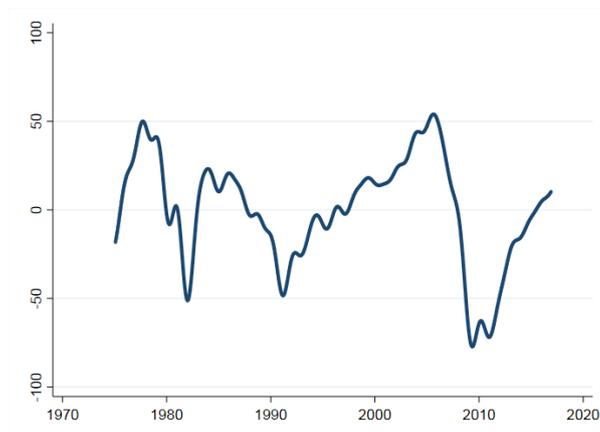
Notes: This repeats the results from Table 3.5 but with frequencies that give cycles of between 3 and 50 years in length.

Figure D.6: Cyclical element of real house price growth with shorter cycles



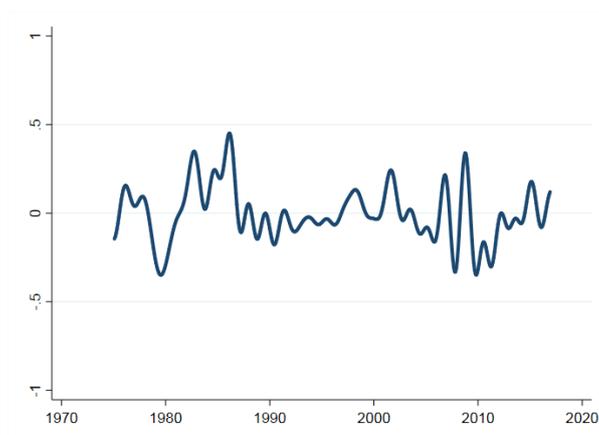
Notes: This repeats the graphs from Figure 3.4 but with frequencies that give cycles of between 1.5 and 40 years in length.

Figure D.7: Cyclical element of housing sales with shorter cycles



Notes: This repeats the graphs from Figure 3.5 but with frequencies that give cycles of between 1.5 and 40 years in length.

Figure D.8: Cyclical element of real housing rental price growth with shorter cycles



Notes: This repeats the graphs from Figure 3.6 but with frequencies that give cycles of between 1.5 and 40 years in length.

Table D.4: Housing market cyclical linear correlations

	HP	HS	HR
HP	1		
HS	0.59*	1	
	(0.40 to 0.79)		
HR	0.26*	0.20	1
	(0.03 to 0.50)	(-0.08 to 0.47)	

Notes: This repeats the results from Table 3.4 but with frequencies that give cycles of between 1.5 and 40 years in length.

Table D.5: Housing market cyclical rank correlations

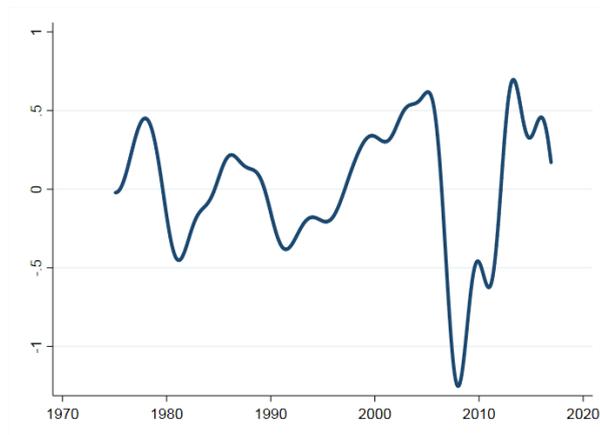
	HP	HS	HR
HP	1		
HS	0.62*	1	
	(0.40 to 0.84)		
HR	0.27*	0.23	1
	(0.03 to 0.50)	(-0.03 to 0.49)	

Notes: This repeats the results from Table 3.5 but with frequencies that give cycles of between 1.5 and 40 years in length.

D.3 Housing data source robustness

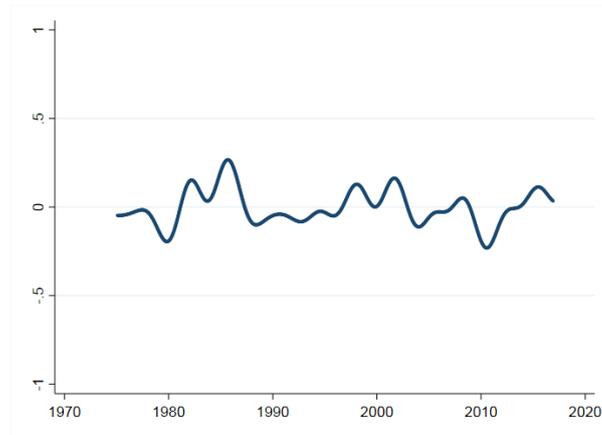
This section of the appendix contains the robustness checks from using different real house price data. It repeats the analysis in Section 3.3 and Section 3.4 but using the Case-Shiller real house price index. The following graphs and tables show the cyclical components, the associations between them and the impulse response to a monetary shock. In all cases the results are similar to those in the main text.

Figure D.9: Cyclical element of real house price growth



Notes: This repeats the graph from Figure 3.4 but uses the Case-Shiller real house price index.

Figure D.10: Cyclical element of real housing rental price growth



Notes: This repeats the graph from Figure 3.6 but deflating rents in line with the deflator used in the Case-Shiller real house price index.

Table D.6: Housing market cyclical linear correlations

	HP	HS	HR
HP	1		
HS	0.61*	1	
	(0.39 to 0.82)		
HR	0.21	0.25	1
	(-0.03 to 0.45)	(-0.05 to 0.55)	

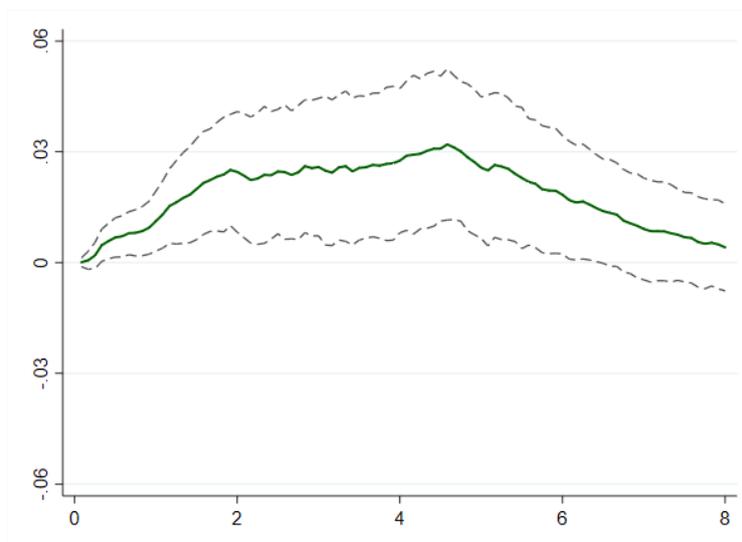
Notes: This repeats the results from Table 3.4 but with the Case-Shiller real house price index and deflator.

Table D.7: Housing market cyclical rank correlations

	HP	HS	HR
HP	1		
HS	0.63*	1	
	(0.40 to 0.86)		
HR	0.24	0.20	1
	(-0.04 to 0.52)	(-0.10 to 0.49)	

Notes: This repeats the results from Table 3.5 but with the Case-Shiller real house price index and deflator.

Figure D.11: Impulse response function of real house prices to a monetary shock



Notes: This repeats the graph from Figure 3.2 but uses the Case-Shiller real house price index.

Bibliography

- Abascal, M., Alonso, T., and Mayordomo, S. (2015). Fragmentation in the European interbank market: Measures, determinants, and policy solutions. *BBVA Working Papers*, No. 13/22.
- Ahmed, N. K., Atiya, A. F., El Gayar, N., and El-Shishiny, H. (2010). An empirical comparison of machine learning models for time series forecasting. *Econometric Reviews*, 29(5-6):594–621.
- Akerlof, G. A. (2002). Behavioral Macroeconomics and Macroeconomic Behavior. *American Economic Review*, 92(3):411–433.
- Alessi, L. and Kerssenfischer, M. (2019). The response of asset prices to monetary policy shocks: Stronger than thought. *Journal of Applied Econometrics*, 34(5):661–672.
- Baele, L., Ferrando, A., Hördahl, P., Krylova, E., and Monnet, C. (2004). Measuring European financial integration. *Oxford Review of Economic Policy*, 20(4):509–530.
- Bai, J. and Ng, S. (2006). Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica*, 74(4):1133–1150.
- Bai, J. and Ng, S. (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics*, 146(2):304–317.
- Balcilar, M., Gupta, R., Sousa, R. M., and Wohar, M. E. (2019). What Can Fifty-Two Collateralizable Wealth Measures Tell Us About Future Housing Market Returns? Evidence from U.S. State-Level Data. *Journal of Real Estate Finance and Economics*, pages 1–27.

- Baumeister, C. and Hamilton, J. D. (2015). Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information. *Econometrica*, 83(5):1963–1999.
- Baxter, M. and King, R. G. (1999). Measuring business cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics*, 81(4):575–593.
- Berge, T. J. (2018). Understanding survey-based inflation expectations. *International Journal of Forecasting*, 34(4):788–801.
- Bernanke, B. S., Boivin, J., and Eliasziw, P. (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach. *Quarterly Journal of Economics*, 120(1):387–422.
- Bhutta, N. and Ringo, D. (2017). The Effect of Interest Rates on Home Buying: Evidence from a Discontinuity in Mortgage Insurance Premiums. *Federal Reserve Finance and Economics Discussion Series*, No. 086.
- Blake, A. and Mumtaz, H. (2017). Applied Bayesian econometrics for central bankers. Technical report, Bank of England: Center for Central Banking Studies.
- Borio, C. (2014). The financial cycle and macroeconomics: What have we learnt? *Journal of Banking and Finance*, 45(1):182–198.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356):791–799.
- Breitung, J. and Eickmeier, S. (2011). Testing for structural breaks in dynamic factor models. *Journal of Econometrics*, 163(1):71–84.
- Brunnermeier, M. K., Eisenbach, T. M., and Sannikov, Y. (2011). Macroeconomics with financial frictions: A survey. *National Bureau of Economic Research Working Papers*, No. 18102.
- Brunnermeier, M. K. and Oehmke, M. (2013). Bubbles, Financial Crises, And Systemic Risk. In *Handbook of the Economics of Finance*, pages 1221–1288. Elsevier.

- Burns, A. F. and Mitchell, W. C. (1946). Measuring Business Cycles. *NBER Books*.
- Campbell, S. D., Davis, M. A., Gallin, J., and Martin, R. F. (2009). What moves housing markets: A variance decomposition of the rent-price ratio. *Journal of Urban Economics*, 66(2):90–102.
- Canova, F. (2007). *Methods for Applied Macroeconomic Research*. Princeton University Press.
- Carter, C. K. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.
- Case, K. E. and Shiller, R. J. (1989). The efficiency of the market for single family homes. *American Economic Review*, 79(1):125–137.
- Case, K. E., Shiller, R. J., and Thompson, A. K. (2012). What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets. *National Bureau of Economic Research Working Papers*, No. 18400.
- Cerra, V. and Saxena, S. C. (2011). Growth Dynamics: The Myth of Economic Recovery. *American Economic Review*, 98(1):439–57.
- Cheng, I. H., Raina, S., and Xiong, W. (2014). Wall street and the housing bubble. *American Economic Review*, 104(9):2797–2829.
- Christiano, L. J. and Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, 44(2):435–465.
- Chung, H. and Xiao, W. (2013). Cognitive Consistency, Signal Extraction, and Macroeconomic Persistence. *SUNY Binghamton Working Paper*.
- Claessens, S., Kose, M. A., and Terrones, M. E. (2012). How do business and financial cycles interact? *Journal of International Economics*, 87(1):178–190.
- Coeure, B. (2013). The way back to financial integration. In *International financial integration and fragmentation. Drivers and policy responses*, Madrid. Banco de Espana.

- Cohn, A., Engelmann, J., Fehr, E., and Maréchal, M. A. (2015). Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals. *American Economic Review*, 105(2):860–885.
- Coibion, O. and Gorodnichenko, Y. (2012a). What Can Survey Forecasts Tell Us About Informational Rigidities? *Journal of Political Economy*, 120(1):116–159.
- Coibion, O. and Gorodnichenko, Y. (2012b). Why Are Target Interest Rate Changes so Persistent. *American Economic Journal: Macroeconomics*, 4(4):126–162.
- Coibion, O. and Gorodnichenko, Y. (2015a). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78.
- Coibion, O. and Gorodnichenko, Y. (2015b). Is the phillips curve alive and well after all? Inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.
- Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018). The formation of expectations, inflation, and the Phillips curve. *Journal of Economic Literature*, 56(4):1447–91.
- Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent Bystanders? Monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89.
- De Mol, C., Giannone, D., and Reichlin, L. (2008). Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components? *Journal of Econometrics*, 146(2):318–328.
- Del Giovane, P., Fabiani, S., and Sabbatini, R. (2008). What’s behind inflation perceptions? A survey-based analysis of Italian consumers. *Bank of Italy Working Papers*, No. 655.
- Del Negro, M. and Otrok, C. (2007). 99 Luftballons: Monetary policy and the house price boom across U.S. states. *Journal of Monetary Economics*, 54(7):1962–1985.

- Diebold, F. X. (2017). *Forecasting*. University of Pennsylvania, Philadelphia.
- Drehmann, M., Borio, C., and Tsatsaronis, K. (2012). Characterising the Financial Cycle: Don't Lose Sight of the Medium Term! *Bank for International Settlements Working Papers*, No. 380.
- Dusha, E. and Janiak, A. (2018). On the dynamics of asset prices and liquidity: the role of search frictions and idiosyncratic shocks. *Society for Economic Dynamics 2018 Meeting Papers*, No. 524.
- Ericsson, N. R. (2016). Eliciting GDP forecasts from the FOMC's minutes around the financial crisis. *International Journal of Forecasting*, 32(2):571–583.
- European Central Bank (2016). Financial integration in Europe. Technical report, European Central Bank, Frankfurt.
- European Central Bank (2018). Financial integration in Europe. Technical report, European Central Bank, Frankfurt.
- European Investment Bank (2017). Investment Report: from recovery to sustainable growth. Technical report, European Investment Bank, Luxembourg.
- European Investment Bank (2019). Investment Report: accelerating Europe's transformation. Technical report, European Investment Bank, Luxembourg.
- Evans, G. W. and Honkapohja, S. (2012). *Learning and expectations in macroeconomics*. Princeton University Press.
- Fama, E. (1970). Efficient Capital Markets: A Review of the Theory. *The Journal of Finance*, 25(2):383–417.
- Faust, J. and Wright, J. H. (2013). Forecasting inflation. In *Handbook of Economic Forecasting*. Elsevier.
- Ferrari, A. and Rogantini Picco, A. (2017). International Risk Sharing in the EMU. *ESM Working Papers*, No. 17.

- Fratantoni, M. and Schuh, S. (2003). Monetary Policy, Housing, and Heterogeneous Regional Markets. *Journal of Money, Credit, and Banking*, 35(11):557–58958.
- Friedman, M. (1977). Nobel Lecture: Inflation and Unemployment. *Journal of Political Economy*, 85(3):451–472.
- Fuster, A. and Zafar, B. (2015). The Sensitivity of Housing Demand to Financing Conditions: Evidence from a Survey. *Federal Reserve Bank of New York Staff Report*, No. 702.
- Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.
- Ghysels, E., Plazzi, A., Valkanov, R., and Torous, W. (2013). Forecasting real estate prices. *Handbook of Economic Forecasting*, 2:509–580.
- Giacomini, R. (2015). Economic theory and forecasting: lessons from the literature. *Econometrics Journal*, 18(2):C22–C41.
- Glaeser, E. L. and Nathanson, C. G. (2015). Housing Bubbles. *Handbook of Regional and Urban Economics*, 5:701–751.
- Gordon, R. J. (1982). Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment. *National Bureau of Economic Research Working Paper Series*, No. 708.
- Granger, C. and Newbold, P. (1986). *Forecasting Economic Time Series*. Academic Press.
- Guiso, L., Sapienza, P., and Zingales, L. (2018). Time varying risk aversion. *Journal of Financial Economics*, 128(3):403–421.
- Havranek, T. and Rusnak, M. (2013). Transmission lags of monetary policy: A meta-analysis. *International Journal of Central Banking*, 9(4):37–76.
- Hoffmann, P., Kremer, M., and Zaharia, S. (2019). Financial Integration in Europe Through the Lens of Composite Indicators. *European Central Bank Working Papers*, No. 2319.

- Hommes, C., Mavromatis, K., Ozden, T., and Zhu, M. (2019). Behavioral Learning Equilibria in the New Keynesian Model. *De Nederlandsche Bank Working Paper Series*, No. 654.
- Jarociński, M. and Smets, F. (2008). House prices and the stance of monetary policy. *European Central Bank Working Papers*, No. 891.
- Jorda, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.
- Jorda, O., Schularick, M., and Taylor, A. M. (2015a). Betting the house. *Journal of International Economics*, 96(S1):S2–S18.
- Jorda, O., Schularick, M., and Taylor, A. M. (2015b). Leveraged bubbles. *Journal of Monetary Economics*, 76:S1–S20.
- Jung, H. and Lee, J. (2017). The effects of macroprudential policies on house prices: Evidence from an event study using Korean real transaction data. *Journal of Financial Stability*, 31:167–185.
- Kalemli-Ozcan, S., Luttini, E., and Sørensen, B. (2014). Debt crises and risk-sharing: The role of markets versus sovereigns. *Scandinavian Journal of Economics*, 116(1):253–276.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. Palgrave Macmillan.
- Kilian, L. and Helmut, L. (2017). *Structural Vector Autoregressive Analysis*. Cambridge University Press.
- Kim, H. H. and Swanson, N. R. (2018). Mining big data using parsimonious factor, machine learning, variable selection and shrinkage methods. *International Journal of Forecasting*, 34(2):339–354.
- Koop, G. (2017). Bayesian methods for empirical macroeconomics with big data. *Review of Economic Analysis*, 9(1):33–56.

- Lane, P. R. (2008). EMU and Financial Integration. *Institute for International Integration Studies Discussion Papers*, No. 272.
- Lovell, M. (1986). Tests of the Rational Expectations Hypothesis. *The American Economic Review*, 76(1):110–124.
- Lucas, R. E. (1996). Nobel Lecture: Monetary Neutrality. *Journal of Political Economy*, 104(4):661–682.
- Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of Economic Literature*, 52(1):5–44.
- Mayer, C. and Sinai, T. (2007). Housing and Behavioral Finance. In *Implications of Behavioral Economics on Economic Policy*. Federal Reserve Bank of Boston.
- Medeiros, M. C., Vasconcelos, G., Veiga, A., and Zilberman, E. (2020). Forecasting Inflation in a Data-Rich Environment: The Benefits of Machine Learning Methods. *Journal of Business & Economic Statistics (forthcoming)*.
- Miranda-Agrippino, S. (2016). Unsurprising shocks: information, premia, and the monetary transmission. *Bank of England Staff Working Papers*, No. 626.
- Murray, C. J. (2003). Cyclical Properties of Baxter-King Filtered Time Series. *The Review of Economics and Statistics*, 85(2):472–476.
- Musso, A., Neri, S., and Stracca, L. (2011). Housing, consumption and monetary policy: How different are the US and the euro area? *Journal of Banking and Finance*, 35(11):3019–3041.
- Muth, J. F. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica*, pages 313–315.
- Ngai, L. R. and Tenreyro, S. (2014). Hot and cold seasons in the housing market. *American Economic Review*, 104(12):3956–3990.
- Nikolaidi, M. and Stockhammer, E. (2017). Minsky models: a structured survey. *Journal of Economic Surveys*, 31(5):1304–1331.

- Orphanides, A. and Williams, J. C. (2007). Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54(5):1406–1435.
- Pesaran, M. H. and Weale, M. (2006). Survey Expectations. In *Handbook of Economic Forecasting*, pages 715–776. North Holland.
- Piazzesi, M. and Schneider, M. (2016). Housing and Macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 1547–1640. Elsevier.
- Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. In *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.
- Romer, C. D. and Romer, D. H. (2000). Federal reserve information and the behavior of interest rates. *American Economic Review*, 90(3):429–457.
- Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, 94(4):1055–1084.
- Romer, P. (2016). The Trouble with Macroeconomics. *The American Economist*, 20:1–20.
- Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies*, 77(2):665–696.
- Salzman, D. A. and Zwinkels, R. C. J. (2017). Behavioural Real Estate. *Journal of Real Estate Literature*, 25(1):77–106.
- Santos, M. S. and Woodford, M. (1997). Rational Asset Pricing Bubbles. *Econometrica*, 65(1):19–57.
- Schindler, F. (2013). Predictability and Persistence of the Price Movements of the S&P/Case-Shiller House Price Indices. *Journal of Real Estate Finance and Economics*, 46(1):44–90.
- Schindler, M. (2009). Measuring financial integration: A new data set. *International Monetary Fund Staff Papers*, 56(1).

- Sheffrin, S. M. (1996). *Rational Expectations*. Cambridge University Press.
- Sinai, T. (2015). House Price Moments in Boom-Bust Cycles. *National Bureau of Economic Research Working Paper Series*, No. 18059.
- Stiglitz, J. E. (2010). Risk and global economic architecture: Why full financial integration may be undesirable. *American Economic Review Papers & Proceedings*, 100(2):388–92.
- Stiglitz, J. E. (2018). Where modern macroeconomics went wrong. *Oxford Review of Economic Policy*, 34(1-2):70–106.
- Stock, J. H. and Watson, M. W. (1996). Evidence on Structural Instability in Macroeconomic Time Series Relations. *Journal of Business & Economic Statistics*, 14(1):11–30.
- Stock, J. H. and Watson, M. W. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business and Economic Statistics*, 20(2):147–162.
- Stock, J. H. and Watson, M. W. (2008). Phillips Curve Inflation Forecasts. *National Bureau of Economic Research Working Papers*, No. w14322.
- Stock, J. H. and Watson, M. W. (2016). Core inflation and trend inflation. *Review of Economics and Statistics*, 98(4):770–784.
- Ungerer, C. (2015). Monetary Policy, Hot Housing Markets and Leverage. *Federal Reserve Finance and Economics Discussion Series*, No. 048.
- Williams, J. C. (2015). Measuring the effects of monetary policy on house prices and the economy. *Bank for International Settlements Working Papers*, No. 88.